Essays in Empirical Labour Economics

Salvatore Lattanzio

Faculty of Economics
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Wolfson College

April 2021

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Salvatore Lattanzio
April 2021

Essays in Empirical Labour Economics

Salvatore Lattanzio

Summary

This thesis analyses the role of workplace heterogeneity in determining pay differentials between workers, employing a range of reduced-form tools and using detailed matched employer-employee administrative data.

The first chapter, co-authored with Alessandra Casarico, studies the contribution of differences in firm pay policy to the gender wage gap in Italy, decomposing them into a between-firm component of sorting of women in low-pay firms and a within-firm component related to differences in bargaining power between gender. Building on Card et al. (2016), we investigate the contribution of firms to the gender gap in earnings at different deciles of the earnings distribution, by age and cohort, and over time. Using a linked employer-employee dataset for Italy, covering the universe of workers in the private sector, we show that the gap in firm pay policy explains on average 30 percent of the gender pay gap in the period 1995-2015. When we decompose differences in firm pay policy into sorting and bargaining, we find that sorting of women in low pay firms dominates on average and at the bottom of the distribution, whereas bargaining prevails at the top and has increased in importance over time. We explore gendered mobility patterns towards firms with more generous pay policy as a driver of sorting and exploit exogenous variation in the gender composition of board of directors to study the impact of firm environment on gender differences in bargaining power. We find that women are less likely to move towards more generous firms, especially in the event of firm closures, and that exogenous changes in the gender balance in leadership positions reduce the gender gap in bargaining power, indicating that the latter is partly malleable to institutional changes.

The second chapter, co-authored with Long Hong, studies the contribution of coworkers on future wage growth. Using linked employer-employee data for the Veneto region in Italy, we explore coworkers' effect on wage growth in two directions. First, using a novel estimation method and accounting for the endogenous sorting of workers into peer groups and firms, we estimate the impact of average peer quality on future wages. We find that a 10 percent rise in peer quality increases one's wage in the next year by 1.8 percent. The effect decreases gradually over time and becomes about 0.7 percent after five years. Second, we delve deeper into the channels that identify the peer effect and, using an event-study specification around mobility episodes, we study how the entry and leave of high-quality and low-quality workers affect wages of movers and coworkers. We find that hiring a high-quality worker is an important driver of wage growth, as well as separating from a

low-quality worker. Movers experience an immediate gain when moving into high-quality peers. Knowledge spillover and peer pressure are likely important mechanisms in explaining our findings.

The third chapter studies the worker-, firm- and sector-level adjustment to robots. Combining detailed matched employer-employee data for Italy over the period 1994-2018 with robot counts by industry in the manufacturing sector, we show that automation adoption expands employment opportunities and reduces labour market transitions. At the worker level, those who are either high-skilled, white-collar, or employed in more productive firms experience employment and earnings gains. Meanwhile at the firm-level, sales and value added increase, while employment outcomes are highly heterogeneous between ex-ante more and less productive firms; with the former increasing employment of all workers, irrespective of their skill level, and the latter reducing it. These changes in labour demand are further inspected at the sector-level, where an event study approach following spikes in automation adoption reveals a negative effect of automation on labour market sorting. Overall, this chapter provides evidence on the impact of automation on a country with a strong manufacturing sector and a relatively rigid labour market. When exploring heterogeneous effects across workers and firms, there is a clear distinction between “winners” and “losers”, with less skilled workers facing bigger losses from technology adoption.

Acknowledgements

I am deeply indebted to my supervisors Kai Liu and Hamish Low. They have both provided guidance, support and encouragement throughout my PhD journey and pushed my research forward by generously lending their time and intellect. They were always available to discuss my ideas, to correct my mistakes, and to help me improve my research. I have learnt so much from them and I could have not asked for better mentors.

I am thankful to Alessandra Casarico. She is way more than the co-author of the first chapter of this dissertation. She first gave me the chance to confront myself with research and, as already stated in the acknowledgements of my MSc dissertation a while ago, having attended her Public Economics course is still one of the best choices of my life.

I am indebted to Long Hong, who co-authored the second chapter. His work ethics, inputs and enthusiasm have been a great inspiration since when we were classmates.

My mentor at the Bank of Italy, Francesco D'Amuri, and the wider Family and Labour Market Division members were excellent colleagues to work with. I thoroughly enjoyed my time in Rome and my academic and professional development has flourished on account of their generous input.

Conversations with Toke Aidt, Noriko Amano-Patino, Maarten De Ridder, Dexter Ding, Elisa Faraglia, Eric French, Lukas Freund, Timo Haber, Zeina Hasna, Dario Palumbo, Christopher Rauh, Davide Rigo, Gabriella Santangelo, Alexandru Savu, Weilong Zhang, were of great benefit. The wider group of empirical microeconomics and PhD students were a fantastic source of feedback. The support staff at the Faculty of Economics, in particular Silvana Dean and Louise Cross, were indispensable for administrative support.

I gratefully acknowledge the financial support of ESRC, Cambridge Trust and Keynes Fund, which considerably helped meet the financial needs of a four year long journey.

I would not be in this position today without the support of my family. My mum will never read these words, but I could see her in front of me with her beautiful smile, her dark eyes full of joy and love. She would not say much, but I could be able to sense her happiness. My dad, my brave and kind dad, who has always been there for me, managed to become an even better dad in the past two years, when everything became tougher for us. His example and his life, started from a much less privileged point than mine, have been and will be a never ending source of inspiration. To him and to the loving memory of my mum I dedicate this thesis.

Table of contents

List of figures

List of tables

Chapter 1

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

Abstract This chapter studies the impact of firm pay policy on gender inequality in the labour market. Building on Card et al. (2016), we investigate the contribution of firms to the gender gap in earnings at different deciles of the earnings distribution, by age and cohort, and over time. Using a linked employer-employee dataset for Italy, covering the universe of workers in the private sector, we show that the gap in firm pay policy explains on average 30 per cent of the gender pay gap in the period 1995-2015. When we decompose differences in firm pay policy into sorting and bargaining, we find that sorting of women in low pay firms dominates on average and at the bottom of the distribution, whereas bargaining prevails at the top and has increased in importance over time. We explore gendered mobility patterns towards firms with more generous pay policy as a driver of sorting and exploit exogenous variation in the gender composition of board of directors to study the impact of firm environment on gender differences in bargaining power. We find that women are less likely to move towards more generous firms, especially in the event of firm closures, and that exogenous changes in the gender balance in leadership positions reduce the gender gap in bargaining power, indicating that the latter is partly malleable to institutional changes.[1]

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

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

[1] This chapter is based on joint work with Alessandra Casarico, Bocconi University.

1.1 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 per cent in 2016 against a value above 30 per cent in 1975, but only 1.7 percentage points below its value in 2005,[2] with large cross-country differences and with smaller reductions at the top of the distribution. A large literature documents the extent of gender wage gaps and their evolution over time,[3] 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 (Azmat and Petrongolo, 2014; Bertrand, 2011; Flinn et al., 2018), and on the important role of child penalties (Cortés and Pan, 2020; Kleven et al., 2019). Clearly, gendered outcomes depend not only on the characteristics and behaviour of workers, but also on those of the firms which employ them. Firm-related gender wage differences can show up through labour market segmentation of women into firms with lower pay rates (Bayard et al., 2003; Card et al., 2016; Cardoso et al., 2016; Groshen, 1991; Ludsteck, 2014; Morchio and Moser, 2019).[4] In addition, women may show lower bargaining power compared to men working at the same firm: women may negotiate less aggressively (Babcock et al., 2006; Bowles et al., 2007, 2005; 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 (Blau, 2012; Petersen and Morgan, 1995). From the policy side, 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.

The goal of this paper is to complement this literature by documenting the role of firm pay policy in hindering or reinforcing gender inequality in the labour market. Building on Card et al. (2016), we study the contribution of firms to the gender gap in earnings at different deciles of the earnings distribution, by age and cohort, and over time. After accounting for individual permanent productivity, we decompose gender differences in firm fixed effects and distinguish between gender differences in sorting across firms and

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

[3] 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).

[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).

in bargaining within firms. The latter captures both allocation of men and women across different hyerarchical levels within firms and gender differences in the ability to bargain over wages. The statistical decomposition reveals that sorting accounts for two thirds of the gender gap in firm pay policy, while bargaining accounts for the remaining third. To provide more direct evidence on the determinants of sorting and bargaining, we explore gendered mobility patterns towards firms with more generous pay policy as a driver of sorting and 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.

We first show that differences in firm-specific premia[5] account for approximately 30 per cent of the Italian gender pay gap at the mean, with sorting accounting for two thirds of the firm contribution compared to bargaining, and for one fifth of the overall gender gap in earnings. This is in line with the existing evidence on Portugal, France and Germany (Bruns, 2019; Card et al., 2016; Coudin et al., 2018). The dominant role of sorting compared to bargaining is persistent across age and cohorts and it is more evident for older women. The decomposition analysis within deciles of the pay distribution shows that bargaining is the main factor driving the firm contribution at the top: even when women work for high-pay firms, their earnings are lower than those of men because of their worse bargaining power. The role of bargaining has increased over time: when we estimate the firm contribution and its decomposition into sorting and bargaining in four overlapping time intervals between 1995 and 2015, we find that gender differences in firm fixed effects are practically unchanged over time, but the contribution of each component has varied considerably, with gender differences in bargaining sharply gaining importance. We argue that this may reflect the spreading of decentralised wage setting in the Italian labour market, and highlight that it has differentially affected men and women, to the detriment of the latter. When investigating mobility patterns, we find that a gender mobility gap is present, with women displaying a lower likelihood of moving to better paying firms only if mobility is associated with firm closures. We also find that women tend not to move towards firms with high (unexplained) earnings dispersion, indicating that gender differences in risk aversion or attitudes towards competition may be at play. Note also that gender differences in mobility can explain not only sorting, but also the lower bargaining power of women. As to the latter, in the experimental part of the analysis, we find that the reform introducing gender quotas in boards reduces the gender gap in bargaining power, especially in firms with lower female

[5]Throughout the paper we use the terms firm pay policy and firm-specific premia as synonyms of the gender-specific firm fixed effect in the earning equation.

presence in the board before the reform and for stayers, employed as white-collar workers and executives, suggesting that the lower bargaining power of women compared to men is partly institution-driven.

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 measure of the gender mobility gap, and point out the lower probability of women to move towards firms adopting more generous pay policy only when moves are associated with firm closures. 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.[9] The evidence that gender quotas affect the gender gap in bargaining power reveals that some trickle down effects from boards to workers 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 may influence gender gaps in earnings.

The remainder of the paper is organised as follows: section 1.2 describes the dataset and provides evidence on the gender gap in earnings in Italy; section 1.3 explains the methodology to measure the firm contribution to the gender pay gap and its decomposition

[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 policy 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). 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]See 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.

[9]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.

into sorting and bargaining; section 1.4 presents the results on the decomposition on average, across the distribution of earnings and over time; section 1.5 investigates firm-related mobility; section 1.6 discusses the impact of the gender quota law on the relative bargaining power of female employees; section 1.7 concludes.

1.2 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.[10] The data provide information on gross annual earnings,[11] the number of days and weeks worked in a given year, the type of contract (whether full-time or part-time), the 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. In a separate record, INPS provides information on location, industry,[12] and date of opening and closure of all firms in the data. Firms may comprise multiple establishments. We link firms to balance sheet information, coming from the AIDA-Bureau Van Dijk dataset, that collects information for all companies obliged to file their accounts within the Italian Business Register. Specifically, we use information on sales and value added.

1.2.1 Descriptive Statistics

We build a panel dataset that comprises one observation per worker per year. Since workers who hold multiple jobs are recorded in the data more than once within a year, we select the observation corresponding to the main job, that is, the one 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,

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

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

[12]Sectors are classified according to Nace Rev. 2 codes (whose Italian counterpart is Ateco 2007). All sectors are available in the data, except the public sector, agriculture and domestic workers.

Table 1.1 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. The Table reports, in columns (1) and (2), summary statistics for male and female workers in the entire sample over the period 1995-2015; in columns (3) and (4), it reports summary statistics for the sample used in the analysis in section 1.4.2, which comprises firms that belong to the dual connected set, i.e. the intersection of male and female largest connected sets (see sections 1.3.1 and 1.3.2 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 real (2010 = 100) euros. The number of workers per firm is computed as the average of the yearly male and female workforce at each firm.

we keep only workers who have been employed for at least 4 weeks during the year.[13] We further restrict our analysis to workers with age between 19 and 65, and with at least two years of labour market experience.

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

Table 1.1 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

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

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

workers is 9.1 million. Firms are 1.6 million. Mean age is slightly higher for men than for women, and so is the average job tenure.[15] The average real weekly earnings – the measure of pay we focus on – are larger for men, with a 22 per cent gender gap. The average number of male workers in a firm is 8 and of female workers is 5, both reflecting the small average firm size of Italian firms. The share of blue-collar workers is higher for males (64 per cent versus 44 per cent), whereas that of white-collar workers is higher for females (50 per cent versus 28 per cent). The percentage of executives and middle managers is higher for male workers (1.8 per cent and 3.9 per cent) than for female workers (0.4 per cent and 1.9 per cent). The share of apprentices is higher for women. Around 6 per cent 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 during the year; for part-time workers we have the number of full-time equivalent weeks, which make weekly earnings of full-time and part-time workers comparable. The measure of full-time equivalent weeks is provided directly by INPS, which exploits the information on the actual number of hours worked in a month (available to them, but not to researchers) and computes full-time equivalent weeks 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, differences in hours for part-timers are controlled for and the gender gap in weekly earnings is very close to that in hourly wages.[16]

1.2.2 Evidence on the Gender Earnings Gap in Italy

Figure 1.1 reports the evolution of the gender gap in log average real weekly earnings over the period 1995-2015. 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.

We ask how far the gender pay gap is related to firm-specific factors. A first evidence is provided in Table 1.2, 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 male ones 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

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

[16] For full-time workers weekly earnings may be an imperfect measure of hourly wages in the presence of overtime hours, on which the data provides no information. Reassuringly, however, the gender pay gap in weekly earnings of 22% that we find over 1995-2015 is not too different from that reported by Eurostat for the private sector based on hourly wages, which oscillates between 17.5% and 19.9% over 2009-2015 (source: Eurostat, Gender pay gap in unadjusted form by type of ownership of the economic activity).

Fig. 1.1 Gender pay gap over the period 1995-2015

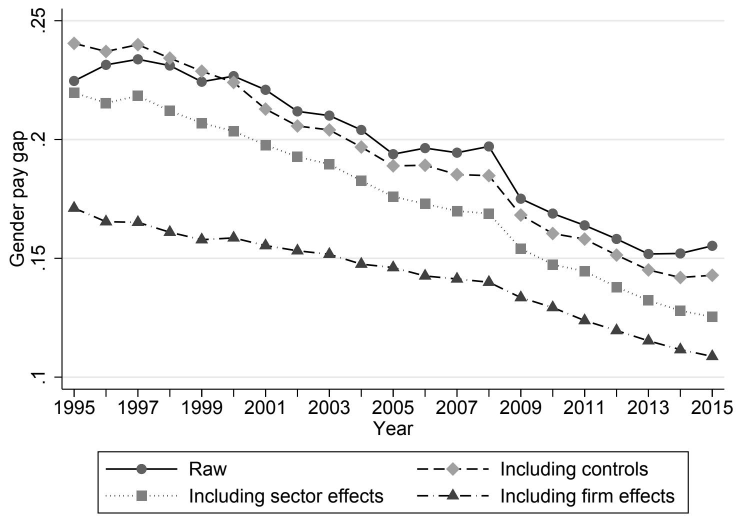

Notes. The Figure plots coefficients of a dummy for male workers from log wage regressions, run for each year in four different specifications: without controls ("Raw"); controlling for observable characteristics of workers, i.e. cubic polynomials in age, experience and tenure, a dummy for full-time 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").

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. Women tend to work for low-pay firms, too, as shown by the specification in Column (4), which includes firm fixed effects (and excludes sector fixed effects, as they would be absorbed by 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 or occupation. 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 1.1, besides the raw gender gap in earnings, we plot the

Table 1.2 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	No
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. The Table reports the coefficients of a dummy variable for male workers from OLS regressions where log weekly earnings are the dependent variable. 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 Rev. 2. Robust standard error, clustered at firm level, in parentheses. Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

coefficients of the male dummy from regressions that control for individual observable characteristics (as in column (2) of Table 1.2), for sector fixed effects (as in column (3) of Table 1.2) and for firm fixed effects, in addition to individual observables (as in column (4) of Table 1.2). Figure 1.1 confirms that firm 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 A.7 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 firm fixed effects (dashed-dotted line).[17] The figure shows the presence of a strong *glass ceiling* effect: 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 earnings inequality for high earners originates within rather than between firms.

[17] Following Canay (2011), we estimate fixed effects quantile regressions in two steps. In the first step, we run an OLS regression, 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.

1.3 Empirical Strategy

We build on Card et al. (2016) and their decomposition method to estimate the impact of firm pay policy, sorting and bargaining along the earnings distribution, by age and cohort, and over time, thus providing novel evidence on the contribution of firms to gender inequality in the labour market. In this section, we describe the details of such decompositions and the regression model used to retrieve the quantities of interest.

1.3.1 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):

$$w_{ijt} = \theta_i + \psi_j^g + X_{it}'\beta^g + \varepsilon_{ijt}, \tag{1.1}$$

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:

$$\psi_j^g = \gamma^g \bar{S}_j, \tag{1.2}$$

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 pay policy, which we allow to vary by gender.[18]

To estimate (1.1), we construct connected sets of firms and workers separately by gender and focus on the largest connected set for female and male workers.[19]

[18]In Appendix A.1, we provide the modelling framework behind equation (1.1).

[19]Abowd et al. (2002) show that identification of equation (1.1) 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.

1.3.2 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.

Hence, we focus on the double connected set of workers and firms. Columns (3) and (4) of Table 1.1 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.

Ideally, given equation (1.1), firm fixed effects should be zero when firms do not share rents with their workers. Thus, we normalise firm fixed effects with respect to the average firm effect in the accommodation and food industry, which is usually identified in the literature as a low-surplus sector (Card et al., 2016; Coudin et al., 2018). The normalisation procedure entails rewriting our estimated firm effects as:

$$\psi_j^g = \widehat{\psi}_j^g - \mathbb{E}\left(\widehat{\psi}_j^g \mid \text{Accommodation and food}\right), \tag{1.3}$$

where ψ_j^g are the normalised firm effects, which are consistent with equation (1.2), $\widehat{\psi}_j^g$ are the estimated firm effects from model (1.1), and the conditioning event means that we are computing the average firm effect in the accommodation and food sector.[20]

1.3.3 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 policy. Following Card et al. (2016), we decompose the difference in firm premia into sorting and bargaining implementing the

[20] We report an alternative normalisation procedure in Appendix A.3, where we empirically identify the set of firms that pay zero rents to their workers. Results do not change and we leave to Appendix A.3 a more thorough discussion of this alternative normalisation procedure.

Oaxaca-Blinder decomposition (Oaxaca, 1973; Blinder, 1973) as follows:

$$\begin{aligned}\mathbb{E}\left[\psi_j^M \mid g=M\right]-\mathbb{E}\left[\psi_j^F \mid g=F\right]=&\mathbb{E}\left[\psi_j^M-\psi_j^F \mid g=M\right]\\&+\mathbb{E}\left[\psi_j^F \mid g=M\right]-\mathbb{E}\left[\psi_j^F \mid g=F\right]\end{aligned}\tag{1.4}$$

$$\begin{aligned}=&\mathbb{E}\left[\psi_j^M-\psi_j^F \mid g=F\right]\\&+\mathbb{E}\left[\psi_j^M \mid g=M\right]-\mathbb{E}\left[\psi_j^M \mid g=F\right].\end{aligned}\tag{1.5}$$

The left hand side of equation (1.4) 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.[21]

This difference can be decomposed in sorting and bargaining in two ways. In equation (1.4), the first term on the right hand side, $\mathbb{E}\left[\psi_j^M-\psi_j^F \mid g=M\right]$, represents the difference in firm premia between men and women, averaged across men. That is, it detects differences in firm premia, fixing the distribution of male jobs. This is a measure of the bargaining channel. It tells by how much the gender pay gap would change if women were given the same firm effects as men, weighted by the male distribution of jobs. The second block, $\mathbb{E}\left[\psi_j^F \mid g=M\right]-\mathbb{E}\left[\psi_j^F \mid g=F\right]$, represents the difference between the average female firm premia evaluated across men and the average female firm premia across women. This difference tells by how much the gender pay gap would change if women were employed in the same firms as men, weighted by the female firm effect.

Similarly, equation (1.5) splits the firm contribution into bargaining, evaluated using the female rather than the male distribution, and sorting, evaluated using male rather than female premia.[22]

We often choose to report the results as averages of sorting and bargaining computed from equations (1.4) and (1.5), i.e.:

$$\begin{aligned}\text{Sorting}&=\frac{1}{2}\sum_{x\in\{F,M\}}\left\{\mathbb{E}\left[\psi_j^x \mid g=M\right]-\mathbb{E}\left[\psi_j^x \mid g=F\right]\right\},\\\text{Bargaining}&=\frac{1}{2}\sum_{x\in\{F,M\}}\mathbb{E}\left[\psi_j^M-\psi_j^F \mid g=x\right].\end{aligned}\tag{1.6}$$

[21] 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.

[22] The first block of equation (1.5), $\mathbb{E}\left[\psi_j^M-\psi_j^F \mid g=F\right]$, 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 (1.5), $\mathbb{E}\left[\psi_j^M \mid g=M\right]-\mathbb{E}\left[\psi_j^M \mid g=F\right]$, 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.

Decomposition across the earnings distribution We know that lower and higher quantiles show a wider gender pay gap (see Figure A.7). Hence, we investigate the impact of the firm contribution on the gender pay gap at various quantiles of the distribution of earnings. Specifically, we select groups in both the male and female samples corresponding to different deciles of the male and female earnings distribution over the whole sample period 1995-2015. In each decile group we drop single-gender firms, in order to avoid contamination of the sorting effect from compositional effects related to the distribution of firms in each group. We then compute, for each gender-specific decile group, the mean male and female firm effects and perform the Oaxaca-Blinder decomposition. In particular, for each gender-specific decile group d_k^g, $k = 1, ..., 10$, we compute:

$$\begin{aligned}
&\mathbb{E}\left[\psi_j^M \mid g = M, i \in d_k^M\right] - \mathbb{E}\left[\psi_j^F \mid g = F, i \in d_k^F\right] \\
=&\mathbb{E}\left[\psi_j^M - \psi_j^F \mid g = M, i \in d_k^M\right] + \mathbb{E}\left[\psi_j^F \mid g = M, i \in d_k^M\right] - \mathbb{E}\left[\psi_j^F \mid g = F, i \in d_k^F\right] \qquad (1.7)\\
=&\mathbb{E}\left[\psi_j^M - \psi_j^F \mid g = F, i \in d_k^F\right] + \mathbb{E}\left[\psi_j^M \mid g = M, i \in d_k^M\right] - \mathbb{E}\left[\psi_j^M \mid g = F, i \in d_k^F\right]. \qquad (1.8)
\end{aligned}$$

In both equations (1.7) and (1.8), 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 (1.7) and (1.8), akin to what we do in (1.6).

1.4 Results

1.4.1 Estimation of Two-way Models

We estimate (1.1) separately for the largest connected groups of female and male workers. We include as controls cubic polynomials in age,[23] 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 1.3 reports sample sizes of the largest connected sets in both the female and male samples. We retain 99.1 per cent and 97.5 per cent of the total person-year observations in the male and female samples, respectively. Men are 98.5 per cent and women are 96.4 per cent of those in the original data. Coverage of firms is 90 per cent and 84.6 per cent in the male and female samples, respectively, compared to the original population.

Panel B of Table 1.3 reports statistics about the fit of the model in equation (1.1) to our data for both samples of female and male workers and it shows that the adjusted R-squared is 0.87 in the male sample and 0.74 in the female one and all the parameters are jointly

[23] 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).

significant.[24] 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. (2019).[25]

Finally, it is important to stress that the validity of the two-way fixed effects model in equation (1.1) relies upon the assumption of conditional random mobility of workers. We test this assumption in Appendix A.2. Overall, we conclude that it holds for both the female and the male sample.

1.4.2 Firm Contribution to the Gender Gap in Earnings and Decomposition into Sorting and Bargaining

To relate our results to those in the existing literature, we start by focusing on the average decomposition in the overall sample. We then move to the age and cohort analysis and to the decomposition across the earnings distribution. Finally, we look at the evolution of the firm contribution over time.

Average Decomposition

Overall sample Focusing on the double connected set of workers and firms, we normalise firm effects as detailed in section 1.3.2 and decompose the difference in firm pay premia as in equations (1.4) and (1.5). Results are in Table 1.4, which 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 per cent 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 (1.4) or (1.5). In both scenarios, sorting accounts for more than 20 per cent of the overall gender pay gap, while bargaining accounts for a smaller share (between 7.6 per cent and 9.8 per cent). This result is similar to the one found by Card et al. (2016) for Portugal, Jewell et al. (2020) for UK and by Coudin et al. (2018) for France.

[24] 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.

[25] A well known problem with the estimation of AKM models is that the sampling errors of worker and firm effects are negatively correlated, implying a downward bias in the covariance between them. This bias is higher the lower the number of movers per firm, as highlighted in Andrews et al. (2008). Recent contributions have proposed different estimation methods for retrieving worker and firm heterogeneity in wage equations (Bonhomme et al., 2019; Kline et al., 2019). It is beyond the scope of this paper to quantify the bias in the estimation of firm and worker effects. Moreover, our decomposition approach focuses on changes in between-group averages involving millions of observations for which, as highlighted in Card et al. (2013), sampling error biases become irrelevant.

Table 1.3 Summary statistics for largest connected sets and AKM estimation

	(1) Male	(2) Female
Panel A: Largest connected sets		
Number of person-year observations	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
Standard deviation earnings	0.486	0.415
Standard deviation worker effects	0.661	0.568
Standard deviation firm effects	0.209	0.195
Standard deviation xb	0.709	0.564
Standard deviation 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 (1.1), separately for men and women.

Thus, sorting is the main factor behind the different premia men and women receive on average.

Decompositions by occupations and by sectors are reported in Appendix A.4.

One caveat with our findings is that we are unable to properly account for hours worked in our data. The literature finds that differences in hours worked and in the willingness to work *long* hours account for a sizeable fraction of the gender wage gap (Bertrand et al., 2010; Goldin, 2014). Our data does not record the number of hours worked by an employee in a year, but luckily it provides an adjusted number of weeks that accounts for differences in hours worked between employees with part-time and full-time contracts, allowing comparisons of weekly wages between them (see section 1.2.1 for further details). Therefore, we are able to keep both full-time and part-time workers in our sample. However, it could still be the case that firm wage premia and part-time status are correlated if, for example, firms that offer high pay policy also allow for less work-time flexibility or require long hours of work. In this case, we would be attributing to the sorting channel part of the differences in wages between gender that are due to different preferences for flexibility. Without further details on hours worked, we cannot really test for how much this potential bias could affect our

Table 1.4 Gender pay gap, firm effects, sorting and bargaining

	(1) Log points	(2) Percent of the gender pay gap
Gender pay gap	0.213	
Male firm effect across males	0.113	
Female firm effect across females	0.049	
Firm effects gap	0.065	30.4%
Decomposition:		
Sorting		
Using male coefficients	0.049	22.8%
Using female coefficients	0.044	20.6%
Bargaining		
Using male distribution	0.021	9.8%
Using female distribution	0.016	7.6%
Observations	183,062,102	

Notes. The Table reports results of the Oaxaca-Blinder decomposition of equations (1.4) and (1.5). Firm effects are normalised with respect to the average gender-specific firm effects in the food and accommodation sector. Column (1) shows results for all workers. Column (2) shows the ratio between the quantities reported in column (1) and the gender pay gap.

results. We have therefore attempted to perform the decomposition of the gap in firm pay policy in sorting and bargaining on the sample of full-time workers. Such decomposition suggests that, indeed, the sorting channel becomes slightly less important, as it amounts to 18.1-18.6 per cent of the gender wage gap, whereas differences in bargaining power account for 10.9-11.4 per cent. However, there might still be correlation between hours and firm wage premia also restricting to full-time workers only, as firms may require long hours of work in exchange of higher compensation. The evidence provided in Card et al. (2016) seems to suggest that this correlation should not be concerning, though. In their paper, they provide correlations between the estimated firm wage premia and hours worked (both regular and overtime) and show that, for both men and women, hours are not a significant predictor of firm pay policy. Given that the Portuguese labour market, on which their analysis is based, shares many features with the Italian labour market – especially in terms of the wage setting mechanism,[26] we are confident that the same results would apply to Italy as well, although we cannot fully rule out possible confounding effects in the Italian context that are not present in the Portuguese one.

[26]Similarly to Italy, in Portugal wage setting is based on sector-level agreements between employers' organizations and unions, and top-up agreements at the firm-level which determine flexibility packages specific to each workplace

Age and cohorts Figure 1.2 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 and round it to the next integer. 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 the firm contribution over the life cycle, highlighted also by Card et al. (2016) and Bruns (2019), is not only an age effect, as they argue, but it is the outcome of both a cohort effect and an age effect *within* cohort.

Decomposition Across the Earnings Distribution

As a first evidence on the magnitude of firm effects along the earnings distribution, we plot in Figure 1.3 the within-decile mean male firm effect across the male distribution and the within-decile mean female firm effect across the female distribution over the period 1995-2015. 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 in the first decile – most likely for the presence of sector-specific minimum wages –, but it starts to widen in the middle part of the distribution. At higher deciles the gap closes, especially in the last one. The closing of the gap at the top of the pay distribution can be due to an increased presence of high-pay female workers in high-pay firms (thus, a better sorting) or to a rise in bargaining power within firm of women relative to men.

We investigate which effect prevails by looking at Figure 1.4, that shows the gender difference in firm effects and its decomposition in sorting and bargaining, as a share of the gender pay gap at each decile. The firm contribution is fairly stable in the central part of the earnings distribution, and smaller at the lowest and highest deciles. As to the determinants of this contribution, sorting is more relevant at the bottom and middle of the distribution. Its importance declines as we move along the distribution. On the contrary, bargaining is the most relevant factor at the 10th decile.

Fig. 1.2 Gender pay gap, firm effects gap, sorting and bargaining by age and cohort

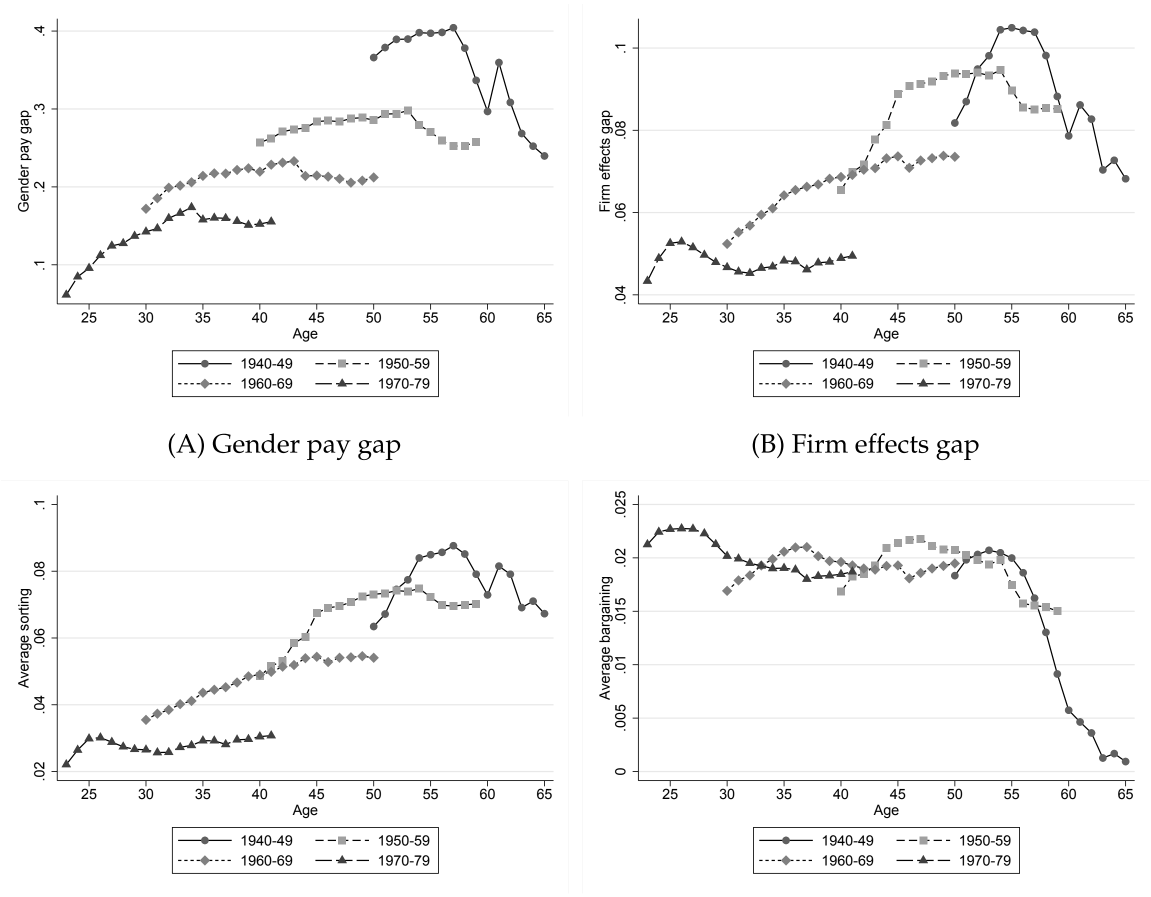

(A) Gender pay gap

(B) Firm effects gap

(C) Sorting

(D) Bargaining

Notes. The Figure plots the gender pay gap (panel A), the gender gap in firm effects (panel B), sorting (panel C) and bargaining (panel D), averaged by age and cohort. The horizontal axis reports age by cohort cells, defined as the mean age for each cohort across the years 1995-2015. Sorting and bargaining are average values, as in equation (1.6).

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.

Evolution of Firm Contribution, Sorting and Bargaining 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 (Alvarez et al., 2018; Barth et al., 2016; Card et al., 2013; Engbom and Moser, 2020; Song et al., 2018) and could impact differently men and women. 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

Fig. 1.3 Firm effects along the earnings distribution

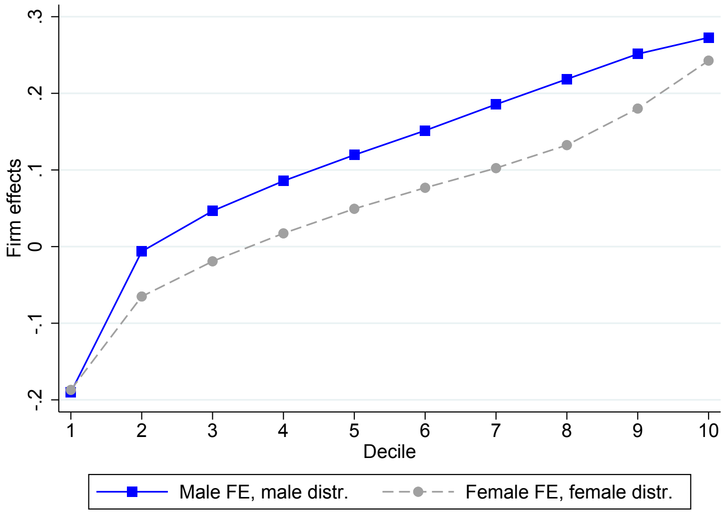

Notes. The Figure shows the average male firm effects across decile bins of the male distribution of earnings and the average female firm effects across decile bins of the female distribution of earnings over the period 1995-2015. We keep only firms that employ at least a man and a woman in each decile group.

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 sub-intervals between 1995 and 2015. Specifically, we run separate AKM regressions in four overlapping intervals of six years each: 1995-2000, 2000-2005, 2005-2010, 2010-2015.[27] 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 1.5, where we plot the average gender pay gap, the firm effect gap, sorting and bargaining in each of the four sub-intervals. As explained in section 1.3.3, we present results averaging sorting and bargaining as in equation (1.6). 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.[28] This evidence, which also includes the aftermath of the Great Recession, differs from that on West Germany in Bruns

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

[28]Table A.3 in the Appendix reports the values used to produce Figure 1.5. It shows that both male and female firm effects increase especially after 2005, but they grow at the same pace, leaving the difference unaltered.

Fig. 1.4 Impact of firm components on the gender pay gap along the earnings distribution

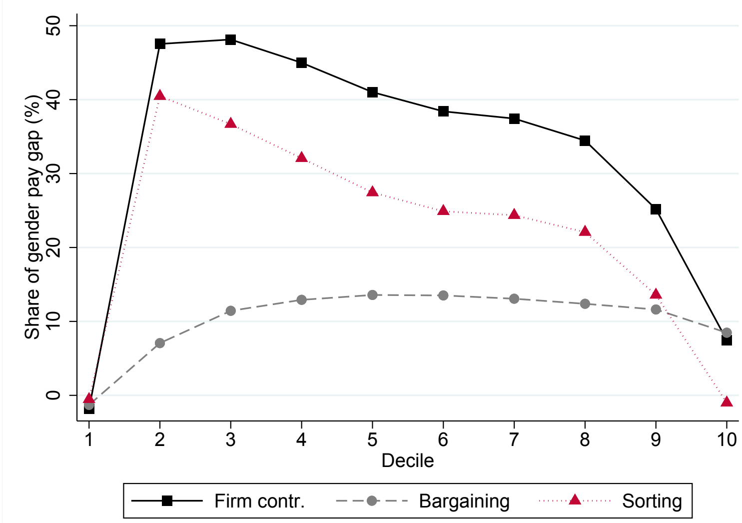

Notes. The Figure shows differences between male and female firm effects and their decompositions into sorting and bargaining across decile bins of the distribution of earnings over the period 1995-2015. We keep only firms that employ at least a man and a woman in each decile group. Sorting and bargaining are average values, as in equation (1.6).

(2019): in Germany, the gender gap in firm effects has increased, rather than having been constant as in Italy, providing an explanation for the stall in the decline of the gender wage gap.

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 per cent), whereas very little is due to within firm differences in firm pay policy. During the period 2010-2015, the two channels have approximately equal weights in explaining the differences between male and female firm effects.[29] Women tend to be employed in "better" firms in more recent years, i.e. in firms with more generous pay policy towards all employees. However, the overall gender gap in firm policy 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 differences in 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. A reform in 1993 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

[29] 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 A.3.

Fig. 1.5 Evolution of gender pay gap, firm effects gap, sorting and bargaining over time.

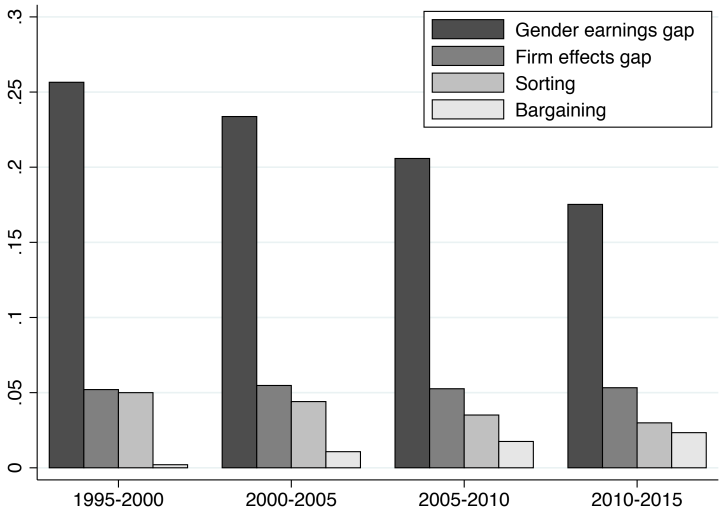

Notes. The Figure shows the evolution of the gender pay gap, firm effects gap, sorting and bargaining over time. We estimate firm effects in each of the four overlapping time intervals. Sorting and bargaining are average values, as in equation (1.6).

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 per cent in 1995 against a value of 50.6 per cent in 2015.[30] 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 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 A.8 shows sorting and bargaining as a percentage of the gap in firm effects in the first and last subinterval. Even though there

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

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 per cent of the gap in firm effects for the 1940 cohort, against a value around 80 per cent in 2010-2015 (panel B). At the same time, bargaining importance has increased for this cohort from values around 4 per cent (panel C) to 20 per cent (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.

1.5 Firm-to-firm Mobility and Sorting

In this section, we investigate gendered mobility patterns in order to gain insights on the sorting channel. We estimate 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 measure of mobility, which takes into account the features of the origin/destination firms and we refer to it as gender mobility gap. In estimating it, we condition on overall mobility.[31] The estimation of gender mobility gaps may seem in contrast with the random mobility assumption required for the identification of firm effects in the AKM model, discussed in Appendix A.2. 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. 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 A.2 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). However, as long as mobility is related to non-wage components, to wage components that do not change over time and are thus absorbed by the individual fixed effect, or to time-varying wage components observable to the researcher, it can be correlated with workers' characteristics.

[31] 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).

[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 per cent of the variance of earnings in the United States.

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

$$\Pr\left\{\mathbf{1}\left[Q^g_{f_1} > Q^g_{f_0}\right]\right\} = \Phi(\alpha + \gamma F_i + \delta Z_{it} + \delta_s + \lambda_t) \tag{1.9}$$

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.

Figure A.9 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 casual employment and fixed-term 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 the two following years. After these restrictions, we are left with a set of 5.2 million job moves. Workers can move more than once over their work career. Overall, 68 per cent of moves in our sample refer to workers who changed job once, 28 per cent twice. Only 4 per cent of moves refer to workers who move three times or more (at most five) between 1995 and 2015. In this restricted sample – Figure A.9 (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 casual 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.

Given that, by definition, workers employed in firms belonging to the top quartile of the distribution cannot move to a firm belonging to a higher quartile, in the main analysis we drop from the sample all moves originating from the top quartile. We include them, instead, in our robustness analysis, in which we estimate whether there is a gender difference in the probability of moving to a firm belonging to the same or to a higher quartile. With this further restriction, we are left with a sample of 3.8 million job moves.

Results on the gender mobility gap Table 1.5 shows average marginal effects from the estimation of equation (1.9). The first column shows the results for all types of job moves.

[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.

Female workers are 1.7 percentage points less likely to move to a better firm within sectors.[34] To characterise gender mobility gaps, we distinguish mobility related to firms' closures from mobility related to 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.[35] These are constrained job moves. All the other moves are classified as "individual". This does not necessarily capture a voluntary choice of the worker, since they can comprise also moves related to, say, occupations disappearing due to technological change or job downgrading following childbirth. In our sample, 2.2 million moves are classified as "individual" and the remaining 1.6 million are "firm" moves. Table 1.5 shows that the gender mobility gap is higher for "firm" moves (column 2) than for "individual" moves (column 3): while the former difference is 3.4 percentage points, the latter is around zero and not significant. When women are constrained to move by their firm closure, they are less likely than men to end up in a firm with a more generous pay policy within a given sector. This may be explained by the network of female workers being weaker and women having worse outside options when forced to leave the firm. In each specification we include as additional covariates a dummy for changing province, occupation and type of contract (specifically from part-time to full-time). Each of these covariates raises the probability of moving to a "better" firm. The effect of switching occupation is the strongest across all moves. We also add age at the moment of the job move, which has a negative impact on the probability of upward mobility, except for firm moves.

The exclusion of movers from the top quartile does not affect our results. Table A.4 in the Appendix reports the estimates for an alternative model that defines the dependent variable as mobility towards a firm in the same or a higher firm effect quartile, therefore including workers moving from firms in the top quartile of the firm effects distribution. Results are quantitatively similar, with women being 1.8 percentage points less likely to move to same/higher quartile firms.

To further explore heterogeneity in gender mobility gaps, we consider their patterns by age. Figure A.10 in the Appendix shows the probabilities for male and female workers of moving to higher-quartile firms by age groups. Women 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 higher for

[34] We also estimate the gender gap in the probability of moving down in the firm fixed effect distribution. In this case, we drop from the sample all moves originating from the bottom quartile of the firm fixed effect distribution, since workers in this quartile cannot worsen their position by definition. In this case, we find that women are more likely to move down in the firm fixed effect distribution. Results are available upon request.

[35] The use of the date of closure may only be a proxy of the effective moment in which firms cease their activity. On the one hand, firms' closures may represent changes in the structure of firms (i.e. mergers or acquisitions). On the other hand, the date of closure may be only a formal record if the firm has effectively ceased its activity in the past and used social benefits to keep employment relationships. To take this into account, we consider as "firm moves" also those happening one year before the firm closure. Note though that the sign and magnitude of the estimates presented reassure against measurement error in the firm closure date.

Table 1.5 Probit model for job moves to a firm in a higher fixed effect quartile

	(1) All	(2) Firm	(3) Individual
Female	-0.017***	-0.034***	-0.001
	(0.006)	(0.009)	(0.007)
Age	-0.001***	0.001**	-0.002***
	(0.000)	(0.000)	(0.000)
Change province	0.028***	0.023***	0.021***
	(0.005)	(0.007)	(0.005)
Change occupation	0.036***	0.039***	0.023***
	(0.004)	(0.006)	(0.005)
Change to full-time	0.018**	-0.004	0.017
	(0.008)	(0.006)	(0.011)
Observations	3,778,512	1,571,607	2,206,905
Sector FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Baseline Probability	0.385	0.374	0.392

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 happening because of firm closure. Column (3) shows results for moves not determined by firm closure. Robust standard errors, clustered at the firm level, in parentheses. Significance levels: * $p<0.1$, ** $p<0.05$, *** $p<0.01$

individual moves, rather than for firm moves when the worker is young, but the decline with age in such probability is faster for individual moves. Moreover, the gender mobility gap tends to be lower or insignificant for older workers. The gap is 4 percentage points for workers aged 19-25 years old, and it becomes approximately zero for workers aged 56-65.

The likelihood of moving to a higher-quartile firm might be influenced by the unobserved ability of the worker. In Figure 1.6 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.[36] 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 across all types of move and workers, but it is only significant for women with low individual fixed effects, especially when moves are associated with firm closure.

Firms offering high pay policy to female workers may be geographically concentrated in some provinces or cities, whereas those offering high pay policy to male workers may

[36] The individual fixed effect captures the portable component of earnings between jobs, which is influenced – among other factors – by individual ability.

Fig. 1.6 Gender-specific probabilities of moving to a higher-quartile firm by worker effect and type of move.

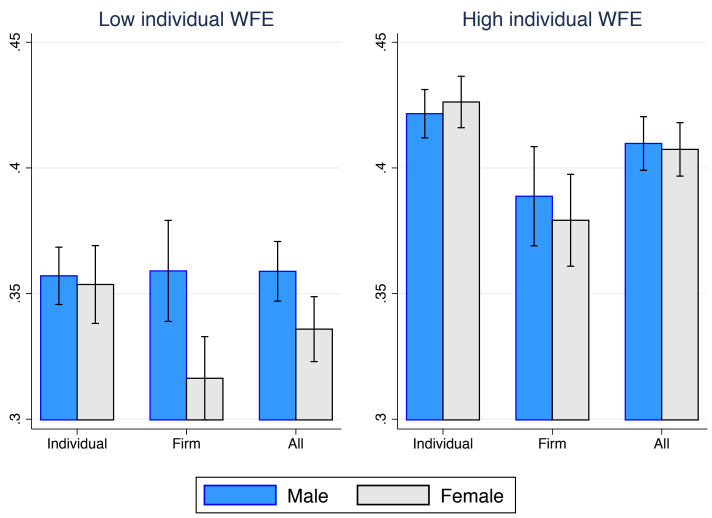

Notes. The Figure reports the probabilities for men and women of moving to a firm in a higher gender-specific firm effect quartile for different types of moves and workers, conditional on the controls included in equation (1.9). Low (high) individual fixed effect workers are defined as those having a worker fixed effect – WFE – below (above) the median of the worker fixed effect distribution. Vertical lines are 95 per cent confidence intervals.

be more geographically dispersed, providing the opportunity to improve earnings to a potentially wider set of them. We graphically inspect the geographical distribution of average male and female firm effects. Figure A.11 in the Appendix plots the map of average male and female firm effects across Italian provinces. 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 of those effects for men and women across provinces is fairly similar, 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 formally test whether the gender mobility gap is a within or between province phenomenon: Figure A.12 in the Appendix plots the marginal effects from probit models estimated as in equation (1.9) separately for each province (thus, excluding the dummy for change of province). Different colours indicate the sign and significance (at 95 per cent confidence level) of coefficients. They are negative for 95 of the 110 Italian provinces, and 54 of them are significantly different from 0 at a 95 per cent confidence level. Hence, also within province women tend to move to employers that offer less generous pay policy relative to men. Note that this evidence could be consistent with women having a higher cost of commuting relative to men (even within the smaller province scale). Interestingly, the size of the coefficient of the female dummy is smaller in absolute terms in Southern provinces. This

Fig. 1.7 Gender-specific probabilities of moving to a higher-quartile firm by firm's variance of residual earnings.

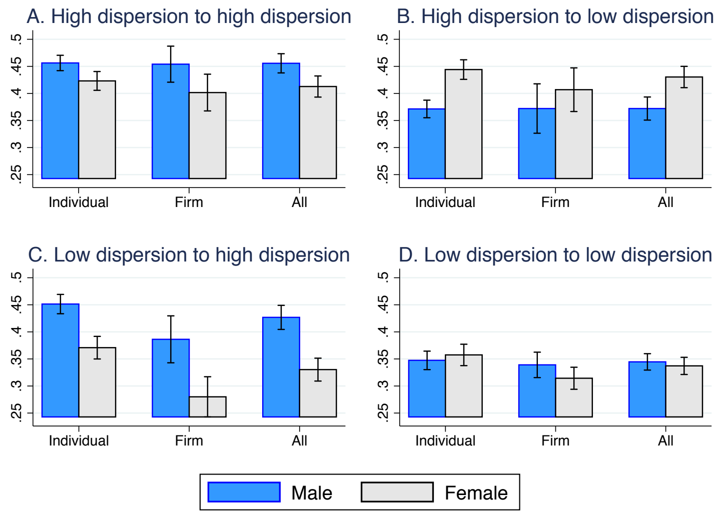

Notes. The Figure reports the probabilities for men and women of moving to a firm in a higher gender-specific firm effect quartile for different types of moves, conditional on the controls included in equation (1.9). We define high (low) dispersion firms those having a standard deviation of residual 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 from high dispersion firms to high dispersion firms. Panel B displays the same probability for movements from high to low dispersion firms. Panel C displays the same probability for movements from low to high dispersion firms. Panel D displays the same probability for movements from low to low dispersion firms. Vertical lines are 95 per cent confidence intervals.

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 make.

Last, we focus on the characteristics of firms and distinguish between those with high and low earnings dispersion. Women may move less frequently to high-pay firms because these firms display higher earnings dispersion, for instance because they use incentive pay more,[37] 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,[38] they may be less willing than men to move to higher quartile firms if these firms have a higher dispersion of earnings. In Figure 1.7 we divide our sample of moves in four groups, defined by the level of earnings dispersion of the origin-destination firm. We define high earnings dispersion firms as follows. We first compute residuals from regressions of weekly earnings on a full set of sector, occupation and full-time dummies. We then compute the firms' standard deviation of

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

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

residual earnings across all workers and periods. High earnings dispersion firms are 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.[39] We finally estimate the probit model (1.9) for each of the four groups of workers that move between firms with different levels of earnings dispersion.

We find evidence that the gender mobility gap widens when movements happen from low to high-dispersion firms (panel C), whereas it is reversed when moves are from high- to low-dispersion firms (panel B) and individual. In other words, women tend to move 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, 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 may be 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. It may also be consistent with firms characterised by a high dispersion of earnings being less willing to hire women.

Overall, this evidence highlights that sorting comes from the lower probability of women to move to better firms. We have shown that this gender mobility gap depends on age and workers' unobserved ability, it persists within provinces and it disappears when the destination firm is characterised by low earnings dispersion. The presence of dependents could be a further factor influencing gendered mobility patterns and sorting, since the time cost of child-rearing disproportionately affects women (Kleven et al., 2019), reducing their possibility to search effectively for a better job, and to end up in a firm with more generous pay policy. 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.[40] 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.

[39] 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 occupational structures.

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

1.6 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 only on pay rises for a given job, but also for promotions within firms, since differences in bargaining power capture both. To what extent can changes in the firm environment influence the gender gap in 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 a 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 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 limits the advantage of male workers. In addition, female workers may be more inclined to ask for increases in pay or for promotions, whereas male workers may have a weaker network to rely on, when the top of the corporate hierarchy is more gender balanced. According to Azmat and Petrongolo (2014), "experimental settings suggests that women and men may differ in traits that are potentially related to labour market success" (p. 37). The ability to negotiate on pay is among them. However, "the causes – nature or nurture – and the economic consequences of such differences are not entirely understood", although an answer to these issues is important, given the clear policy implications. Our analysis contributes to this debate by exploiting a natural experiment, which brings about a change in gender balance at the top of the firm governance. In particular, we focus on a law which prescribes gender quotas in boards of listed firms in Italy. 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 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.[41]

[41] 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. (2019), Gagliarducci and Paserman (2015).

1.6.1 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, ensuring "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.[42] 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.[43] The policy had a clear impact on the share of women in the boards of listed companies, as Figure 1.8 shows. Until 2011, the share of women in the board of directors was 7.4 per cent, 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 per cent, and it kept rising until 33.3 per cent in 2017.

1.6.2 Empirical Analysis

In our modelling framework, outlined in Appendix A.1, we show that the firm fixed effect can be rationalised in a wage equation as a rent-sharing coefficient, i.e. as the share of surplus that a firm pays to its employees. As a consequence, the firm fixed effect can be expressed as in equation (1.2), 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 A.13 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 A.13 is an estimate of the relative bargaining power of women, γ^F/γ^M, and equals 0.85, meaning

[42] Gender quotas have been extended for six additional board renewals by the 2020 Budget Law (Law 160/2019). Moreover, the requirement for the presence of the least represented gender was raised to 40 per cent of the board members, starting from the first board renewal in 2020.

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

Fig. 1.8 Share of women in the board of directors of listed companies between 2008 and 2017.

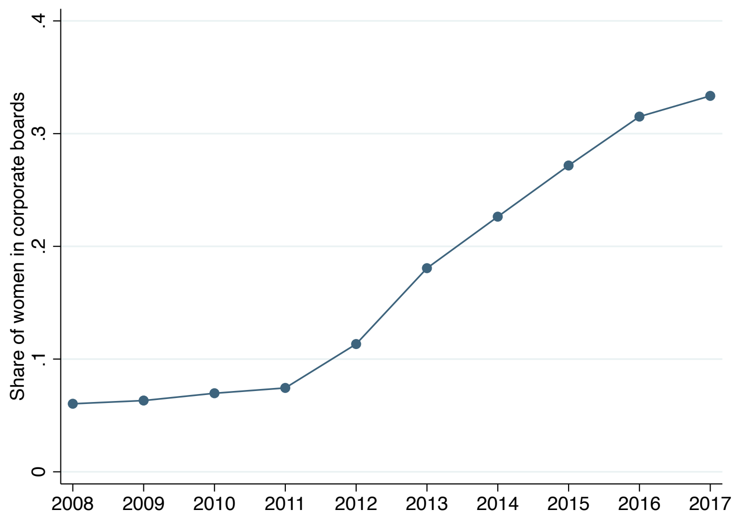

that firms share a lower fraction of increases in value added with female employees relative to males.

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,[44] and we compare the outcome of listed and non-listed companies. In particular, we identify the 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 pre-reform period 2008-2011: log weekly earnings, female log weekly earnings, value added per worker, sales per worker, 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, sector dummies and region dummies.[45] In addition, 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.[46]

To support the validity of our research design, we test the balance of covariates between listed and non-listed firms, before and after the matching, by comparing means of covariates

[44] 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.

[45] The literature on the effects of gender quotas highlights the importance of selecting an appropriate control group for the firms targeted by the reform (Bertrand et al., 2019; Comi et al., 2020; Ferrari et al., 2018; Maida and Weber, 2019).

[46] 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 166 have no missing balance-sheet information, and thus constitute our sample of treated firms.

in the treated and control groups, averaged over the period 2008-2011. Table A.5 shows the comparison. Columns (1) to (3) report means in the unmatched and matched control groups and in the treated group. Columns (4) and (5) report differences in means and the p-value of the significance of the difference in the unmatched sample. Columns (6) and (7) report the same quantities for the matched sample.[47] The table shows that there are no significant differences between treated and control firms in the pre-reform period after matching.

On the matched sample of firms, we estimate, separately by gender, worker-level regressions of the form:

$$w_{ijt} = \kappa + \gamma^g_\Delta Treat_j \times Post_t \times \overline{S}^{pre}_j + f(Treat_j, Post_t, \overline{S}^{pre}_j) + \delta^g X_{it} + \eta_t + \theta_i + \varepsilon_{ijt} \qquad (1.10)$$

where w_{ijt} are log real weekly earnings of worker i in firm j at time t; κ is a constant; $Treat_j$ 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; $\overline{S}^{pre}_j$ is the surplus of firm j, as measured by the log average value added per worker at firm j over the period 2008-2011 – the source is AIDA-Bureau Van Dijk;[48] $f(Treat_j, Post_t, \overline{S}^{pre}_j)$ contains the levels and all the two-way interactions of $Treat_j$, $Post_t$ and $\overline{S}^{pre}_j$. X_{it} are observable characteristics of workers: cubic polynomials in age[49] and experience, and occupation dummies. η_t and θ_i are year effects and worker fixed effects, respectively. The parameter of our interest is γ^g_Δ, which measures the change after the reform in bargaining power by gender in treated versus control firms. Note that we depart from the model with two fixed effects in section 1.3.1 and estimate, instead, a "reduced-form" model in which we regress log male and female weekly earnings on log average value added per worker – controlling for worker time invariant and time-varying characteristics – to recover a rent-sharing coefficient and, with that, measure the bargaining power of workers.

We also estimate models for all workers, without distinguishing by gender, and including interactions of $Treat_j$, $Post_t$ and $\overline{S}^{pre}_j$ with a female dummy, in order to measure the differential change in bargaining power of female workers relative to male workers in treated versus control firms after the reform. Finally, we report results of an event study specification, in which we replace the single indicator for the after reform period $Post_t$ with year dummies, and plot the coefficients of the triple interaction with $Treat_j$ and $\overline{S}^{pre}_j$ by gender. In this way, we ensure that pre-trends are absent and we measure the dynamic impact of the reform over time.

[47] The p-values are computed from univariate regressions of each covariate on a treatment indicator dummy and are robust to heteroskedasticity.

[48] 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).

[49] As in section 1.4, we normalise the age profile to be flat at 40 and exclude the linear term in age from the regression.

1.6.3 Results

Table 1.6, panel A, columns (1) to (3), reports results from the estimation of equation (1.10). Column (1) shows results for men, column (2) for women and column (3) reports the differential impact of the reform on women relative to men in treated versus control firms. The coefficient in column (1) indicates that a 10 per cent increase in value added per worker entails a 0.3 per cent reduction in male earnings in treated firms compared to control firms. At the same time, given that there is no change in female earnings, as shown in column (2), the relative bargaining power of female workers increases: in response to a 10 per cent higher value added per worker, female earnings increase by 0.3 per cent relative to male earnings in treated versus control firms after the reform. Overall, the estimates suggest that gender quotas increased the relative bargaining power of female workers by reducing male bargaining power.

Table 1.6, panels B and C, report results for the samples of new hires and stayers.[50] We show that there are no significant changes in bargaining power for newly hired employees, whereas the effect on stayers is significant and of similar magnitude to that estimated on the full sample of workers. This is in line with gender quotas changing rent sharing for workers already employed in the firm. It is also important to stress that these changes do not happen because the skill composition of the workforce is changing over time, being the estimates conditional on worker fixed effects.

The results are confirmed by the event study specification, shown in Figure 1.9. In particular, we can exclude the presence of different pre-trends for men and women: both sets of coefficients evolve parallelly before the reform and are statistically indistinguishable from 0. Men experience a significant drop in their bargaining power right after the reform in the full sample (panel A) and in the stayers' sample (panel C), whereas women do not experience any change. There are no statistically significant changes for new hires (panel B) for both genders.

The reduction in male bargaining power may come from male employees having a weaker network to rely on after the reform, and therefore lower power in wage and career negotiations. This would be consistent with detecting significant effects for stayers, but not for new hires. Alternatively, new female board members may influence wage setting by adopting as a fair reference point for all workers the average rent-sharing of female workers, which is lower than that of male ones. This result is in line with Bennedsen et al. (2019), that examine the effect of pay transparency laws on within-firm gender inequality and find that the gender pay gap decreases in firms required to disclose pay statistics, mainly through a reduction in male wage growth. Changes in bargaining power can be one mechanism behind this finding.

[50]New hires are identified in the INPS data via the hiring date. Stayers are workers employed by the same firm for the period 2008-2017.

Table 1.6 Impact of gender quotas on bargaining power of workers

	(1) Men	(2) Women	(3) Interaction
Panel A: All workers			
Change in bargaining power	-0.032**	-0.002	0.031***
	(0.013)	(0.009)	(0.008)
Observations	2,413,309	1,356,825	3,770,134
R-squared	0.954	0.928	0.946
Panel B: New hires			
Change in bargaining power	-0.028	0.016	0.047
	(0.046)	(0.036)	(0.045)
Observations	142,392	87,693	230,085
R-squared	0.990	0.988	0.989
Panel C: Stayers			
Change in bargaining power	-0.039***	-0.005	0.035***
	(0.013)	(0.011)	(0.010)
Observations	1,241,290	597,450	1,838,740
R-squared	0.948	0.918	0.939

Notes. The Table shows coefficients from the estimation of equation (1.10) in columns (1) and (2) and from a model that interacts $Treat_j$, $Post_t$ and $\overline{S}_j^{pre}$ with a female dummy in column (3), where the coefficient shown is the one on the quadruple interaction. The dependent variables are log weekly earnings of all workers (panel A), new hires (panel B), and stayers (panel C). New hires are identified in the INPS data via the hiring date. Stayers are workers staying in the firm for the period 2008-2017. All regressions also control for the levels and interactions of $Treat_j$, $Post_t$ and $\overline{S}_j^{pre}$, cubic polynomials in age and experience, occupation dummies, year fixed effects and worker fixed effects. Robust standard errors, clustered at the firm level, in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.10.

Heterogeneous effects We perform two sets of heterogeneity analyses. First, we investigate the presence of differential effects according to how large the change was in the gender composition of the board that the reform imposed on the firm. Namely, we use the average share of female board members in the period 2008-2011 and we include firms in a high intensity group when the pre-reform share of women in board of directors is below 10 per cent, and in a low intensity group when the share is above 10 per cent.[51] We then estimate equation (1.10) separately for workers employed in each of these two groups of firms. The control group contains the full set of matched non-listed firms. Results are reported in Table 1.7, and show that the effect on the relative bargaining power of female versus male workers is positive for workers employed in both groups of firms, but statistically significant only

[51] We have information on the share of women in the pre-reform period for 157, out of 166 firms. 109 firms are classified as high intensity, and 48 as low.

Fig. 1.9 Event-study of the impact of gender quotas on male and female bargaining power

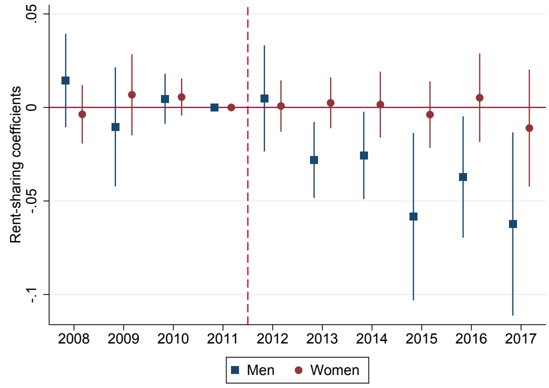

(A) All workers

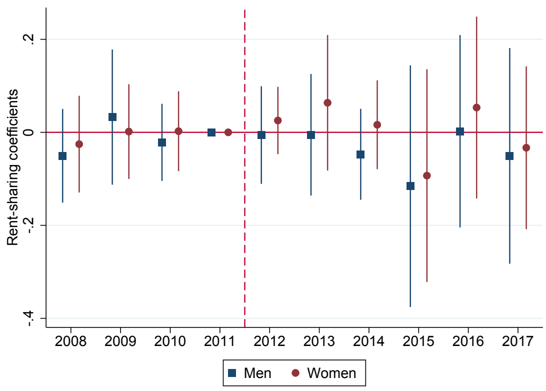

(B) New hires

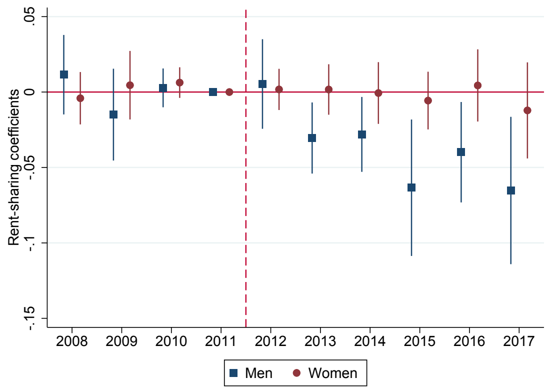

(C) Stayers

Notes. The Figure plots rent-sharing coefficients from the event-study specification explained in section 1.6.2. Vertical lines are 95 per cent confidence intervals.

for those employed by high intensity firms (panel B), and driven by a reduction in male bargaining power, as in the main analysis. Hence, the reform had a greater impact on the relative bargaining power of female workers in those firms, that in the pre-reform period had a worse gender balance in their board of directors and therefore experienced a stronger adjustment.

We then investigate whether the impact of the reform depends on the workers' occupation.[52] Table 1.8 reports the results and shows that the positive and significant effect on the relative bargaining power of female workers is concentrated among white-collar and executives (panels B and C), whereas the effect is negative and not significant for blue-collar workers. Thus, the reform affects workers in highly paid occupations, for which differences in bargaining power represent a larger share of the firm contribution to the gender pay gap,

[52] When estimating equation (1.10) separately by occupation, we omit occupation dummies from the covariates included in X_{it}.

Table 1.7 Heterogeneity in the effect of gender quotas on bargaining power by treatment intensity

	(1) Men	(2) Women	(3) Interaction
Panel A: Low intensity			
Change in bargaining power	0.002	0.017	0.015
	(0.017)	(0.018)	(0.013)
Observations	874,834	418,356	1,293,190
R-squared	0.943	0.940	0.942
Panel B: High intensity			
Change in bargaining power	-0.028**	-0.000	0.030***
	(0.012)	(0.009)	(0.008)
Observations	2,260,717	1,267,530	3,528,247
R-squared	0.954	0.928	0.947

Notes. The Table shows the estimates of the impact of gender quotas on bargaining power by intensity of treatment at the firm level, measured by the share of female board members in the firm prior to the introduction of gender quotas: *Low intensity* are firms with more than a 10 per cent share (panel A), *High intensity* are firms with less than a 10 per cent share (panel B). The control group contains the full set of matched non-listed firms. Robust standard errors, clustered at the firm level, in parentheses. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

as discussed in section 1.4.2. The result suggests that gender quotas could be an effective policy tool to tackle differences in bargaining power between female and male workers in the top echelons of the firm hierarchy.

Overall, the evidence that the reform affects workers' bargaining power, especially in high-intensity firms, is consistent with the gender gap in bargaining power being partly institution-driven.

1.7 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 per cent of the average gender pay gap, with sorting explaining roughly 20-22 per cent of the gender pay gap and bargaining playing the dominant role at the top of the earnings distribution.

When we study gendered mobility patterns, we find that a gender mobility gap is present, with women displaying a lower likelihood of moving to better paying firms, compared to men with similar characteristics, especially when they have low individual fixed effect,

Table 1.8 Heterogeneity in the effect of gender quotas on bargaining power by occupation

	(1) Men	(2) Women	(3) Interaction
Panel A: Blue-collar			
Change in bargaining power	0.012	-0.005	-0.022
	(0.017)	(0.017)	(0.022)
Observations	596,761	241,922	838,683
R-squared	0.820	0.843	0.856
Panel B: White-collar			
Change in bargaining power	-0.034**	-0.003	0.035***
	(0.014)	(0.010)	(0.009)
Observations	1,192,916	851,396	2,044,312
R-squared	0.898	0.836	0.873
Panel C: Executives			
Change in bargaining power	-0.013	0.006	0.021*
	(0.019)	(0.020)	(0.012)
Observations	623,632	263,507	887,139
R-squared	0.925	0.855	0.913

Notes. The Table shows the estimates of the impact of gender quotas on bargaining power by occupation. Panel A reports results for blue-collar workers, panel B for white-collar workers and panel C for middle managers and executives. Robust standard errors, clustered at the firm level, in parentheses. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

and only if the move is triggered by firm closure and if the destination firm displays high earnings dispersion. 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 power, 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 instance, gender differences in peer effects.

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.

Chapter 2

The Peer Effect on Future Wages in the Workplace

Abstract We study a critical driver of wage growth: peers. Using linked employer-employee data for Italy, we explore the peer effect on wage growth in two directions. First, using a novel estimation method and accounting for the endogenous sorting of workers into peer groups and firms, we estimate the impact of the average peer quality on future wages. We find that a 10 percent rise in peer quality increases one's wage in the next year by 1.8 percent. The effect decreases gradually over time and becomes about 0.7 percent after five years. Second, we delve deeper into the channels that identify the peer effect and, using an event-study specification around mobility episodes, we study how the entry and leave of high-quality and low-quality workers affect wages of movers and coworkers. We find that hiring a high-quality worker is an important driver of wage growth, as well as separating from a low-quality worker. Movers experience an immediate gain when moving into high-quality peers. Knowledge spillover and peer pressure play an important role in explaining the mechanisms behind our findings.[1]

Keywords: Peer effects, wage differentials, wage growth, linked employer-employee data

***JEL* codes**: J24, J31, J41, L14, M52

[1]This chapter is based on joint work with Long Hong, University of Wisconsin-Madison.

2.1 Introduction

The literature has well documented that wages typically increase over the life cycle. However, significant heterogeneity exists in wage growth among workers. In line with the canonical models in Becker (1964) and Ben-Porath (1967), wage growth reflects workers' accumulation of knowledge and skills on-the-job. As interaction is essential in the workplace, it is natural that on-the-job learning is primarily the result of interaction with coworkers. Despite the importance, little is known on the link between coworkers and wage *growth* from both empirical and theoretical perspectives. If two workers have the same ability to learn, does the worker with better coworkers experience a faster wage growth over time? How persistent are such peer effects? How does the move of a high-quality worker in a firm contribute to wage growth of his or her new coworkers and how are past ones affected? This paper answers these questions.

There is a growing literature that investigates the relationship between coworkers and wages.[2] Earlier empirical evidence mainly focuses on the effect of coworkers on the contemporaneous wage level in a specific workplace (e.g., Mas and Moretti, 2009) or based on laboratory experiments (e.g., Falk and Ichino, 2006). For example, Mas and Moretti (2009) provide persuasive evidence in a supermarket chain that a cashier's productivity increases when they work alongside more productive coworkers. Nevertheless, it is unclear to what extent these findings, based on a specific firm or laboratory experiment, apply to more general settings and labour markets. Increased access to administrative data allowed researchers to investigate peer effects in one or more local labour markets. For example, Cornelissen et al. (2017) use German employer-employee matched administrative data to study the overall impact of coworkers on contemporaneous wage levels in the Munich local labour market. They find, surprisingly, there is only a small positive effect. However, if knowledge transmission takes time to be reflected in wages, it would be appropriate to consider dynamic effects. That is, the impact of coworkers could materialize on future wage *growth* rather than the current wage level. Despite the potential relevance, only a handful of papers have examined the link between coworkers and wage growth. Two related and complementary papers, Jarosch et al. (2021) and Herkenhoff et al. (2018), both find a substantial knowledge spillover from coworkers that facilitate wage growth, using data from Germany and the United States, respectively.

We aim to explore the causal effect of coworkers on wages along several directions. To do so, we use a matched employer-employee administrative dataset – the Veneto Worker History panel –, which covers the universe of private-sector workers and firms in Veneto, one of the largest Italian regions, from 1975 to 2001. The availability of the full employment information for all firms in the dataset is crucial for our purposes as it allows us to track workers and coworkers over a long period of time. We proceed in this way. First, we explore

[2] Throughout the paper we use the terms peer and coworker interchangeably.

the overall impact of peer quality on one's future wages. By incorporating a coworker component in the canonical AKM model (Abowd et al., 1999), our econometric strategy helps circumvent the common reflection problem (Manski, 1993) and account for workers' endogenous sorting into peer groups and firms. We adopt the novel estimation strategy developed by Hong and Mikkel (2020) to overcome the estimation challenge induced by the high-dimensional fixed effects. Specifically, we exploit two sources of variation to identify the peer effect: changes in peer quality for workers who switch peer groups, and changes in peer quality for workers who remain with their peer group as other workers join or leave the peer group.

Our analysis reveals that peer effects are an important and persistent driver of wage growth. Our baseline results show that a 10 percent increase in coworker quality, as measured by the average AKM worker effect of one's peers, implies an increase of contemporaneous weekly wages by 2.6 percent and next year's wages by 1.8 percent. The effect decays over time, but we find that it is present and significant also in future years, as we find that better coworkers in the past increases one worker's wage by 0.7 percent after five years. The effect is stronger for workers moving to a new job, indicating that part of the peer effect materializes through complementaries between knowledge acquired from different peer groups over time or through better outside options when the worker moves into a new firm. Moreover, peer effects are larger for younger and junior workers (with lower tenure), as learning from peers is probably more important for them, but we do not find differences between workers employed by firms of different size.

In the second part of the paper, we delve deeper into the mechanisms that identify the peer effect and study mobility of workers across firms and how it affects coworkers' wages in the origin and destination firms. Specifically, we exploit an event study analysis around the mobility of workers and analyse how the *entry* of a high-quality worker, who could potentially transmit knowledge to her peers, changes the trajectory of peer average wages in the destination firm. As the choice to hire a worker is not random, by simply comparing outcomes in firms that hire to those that do not hire would likely bias our estimates upward, as hiring choices would be correlated with firm performance and therefore average wages. We overcome this problem by selecting a sample of firms that hire a worker in a given year and analysing coworkers' wages in firms hiring a high- or low-quality worker relative to firms hiring average-quality workers: we define high-, average- and low-quality workers by comparing the estimated AKM worker effect of the mover, from the first part of our analysis, to that of coworkers in the firm. Besides focusing on firms hiring or separating from a worker only, we perform ex-ante propensity score matching before mobility and assess the absence of observable differences between treated and control firms. On top of that, we use data for a subset of firms in our data to show that mobility decisions in the matched sample are not correlated with leads and lags of sales and value added per worker, reassuring on the validity of our empirical strategy. These analyses reveal that hiring a high-quality worker is associated with an increase of peer wages by 3.1 percent relative to firms hiring average-quality workers three years after mobility. The effect is stronger for coworkers belonging to the same occupation of the new hire, as one would expect that more interaction occurs in the workplace among them. The effect is not statistically significant, instead, for firms hiring low-quality workers.

We also examine how wages of coworkers in the origin firm are affected by the departure of high- and low-quality workers. We find opposite results to those outlined above, as the departure of a high-quality worker leads to a drop in wages of about 1.3 percent for coworkers in the origin firm relative to firms separating from an average-quality worker. On the contrary, we find that the loss of low-quality workers benefits coworkers in the firm, whose wages increase by 2.4 percent on average in the three years after the move.

Finally, we shift the focus to movers and examine how wages of workers who move into different peer groups evolve over time, by comparing the wage trajectories of workers moving into high- and low-quality peer groups compared to that of matched workers moving into average-quality peer groups. We find that workers gain from moving into high-quality peer groups. On average their weekly wage increases by 3.9 percent relative to workers moving into average peer groups. We find instead a null effect for workers moving into a low-quality peer group. Taken together, these findings highlight the importance and persistence of coworkers – and high-quality ones, especially – in shaping wage growth.

Our paper makes four contributions to the literature. First, we improve on the existing evidence on peer effects by studying how they evolve over time, allowing for dynamic effects that may not be captured by research focusing only on contemporaneous effects, either in a specific workplace (e.g., Brune et al., 2020; Mas and Moretti, 2009; Papay et al., 2020; Sandvik et al., 2020) or in a more general labour market (e.g., Battisti, 2017; Cornelissen et al., 2017; Lengermann, 2002). In particular, the closest papers to ours in this strand of the literature are Cornelissen et al. (2017) and Battisti (2017). The former studies the peer effect in the Munich local labour market and the latter in Veneto, using the same data that we use. They both use an empirical model similar to ours and find a positive peer effect on wages, but the magnitude is larger in Italy than in Germany. We innovate with respect to these papers, by adopting a different estimation method, as detailed below, and by providing evidence on how peers influence not only contemporaneous but also future wages. We further depart from these papers by adding evidence on which mechanism plays a more prominent role in determining the peer effect, by separating job stayers who change their peer group as workers join or leave the firm and job switchers who change their peer group by moving to a new firm. This brings to our second contribution to the literature, as we exploit the richness of our data to explore the impact of mobility of workers across firms on wage growth of movers and peers. Researchers in the past proved the importance of hiring "good" workers for firm performance (see, e.g., Serafinelli, 2019, who examines the effects of hiring high-quality workers on firms' productivity in an event study framework). However, little is known on the trickle-down effects on coworkers. We fill this gap in the literature by providing comprehensive evidence on the effect of hiring high-quality workers or moving into high-quality peers on future wages of movers and peers. Third, we contribute to a growing literature that studies the firm environment as a driver of wage growth (Arellano-Bover and Saltiel, 2021; Gregory, 2019; Herkenhoff et al., 2018; Jarosch et al., 2021; Nix, 2020). We pin down one mechanism, the influence of peers, that contributes to the understanding of how firms impact workers' wages. Finally, we apply a new method for estimating coworker effects in an AKM wage regression that controls for non-random sorting across firms and occupations.[3] Specifically, we apply the method described in Hong and Mikkel (2020), which derives a moment condition that allows the joint estimation of worker and peer effects. This method improves on the iterative procedure pioneered in Arcidiacono et al. (2012), whose accuracy and computational time for convergence is highly dependent on the chosen tolerance level. The algorithm we use imposes minimal assumptions to ensure consistency and allows a faster computation time.

Our findings are consistent with learning from coworkers being one important driver of wage growth. In particular, Cornelissen et al. (2017) highlight how peer effects may show up through peer pressure or knowledge spillover. The effect we find on future wages – especially for movers – cannot be determined by pressure from past peers, but rather reflects knowledge acquired from peers and accumulated in the workers' human capital. Our analyses cannot exclude, however, other mechanisms, such as better networks to rely on, which may increase workers' outside options and, therefore, their wages.

The rest of the paper is organized as follows. Section 2.2 describes the data and provides descriptive statistics. Section 2.3 provides results about peer effects in the AKM framework. Section 2.4 describes the event study analysis around mobility episodes and presents the results. Finally, section 2.5 concludes.

2.2 Data and Descriptive Statistics

Data We use social security administrative data that contains the entire working population and private firms in the region of Veneto in Northern Italy[4] – the Veneto Worker History (VWH) dataset – from 1975 to 2001. We can observe every coworker of each worker over their working life. The database contains three types of administrative datasets: (1) a worker-level demographic register, (2) a firm-level record, and (3) an annual firm-worker social security contribution register. A brief description of each follows.

[3] We build on AKM as it allows to avoid Manski's reflection problem (Manski, 1993): we use a pre-determined long-term worker characteristic (the AKM worker effect) to measure how peers impact wages and not contemporaneous endogenous productivity measures (such as, average wages).

[4] Veneto is the fifth largest region in terms of population and the third most prosperous region in terms of GDP in Italy.

1. The worker register tracks over 3 million workers from 1975 to 2001. It records the entire working history of a worker in the private sector, as long as he/she worked one day in Veneto.[5] It contains basic demographic information, including birth year and place, gender, nationality.

2. The firm register contains all private firms that employ each worker in the worker register.[6] It includes a firm's detailed information such as national tax code, address, start and closure dates, industry.[7] This register also includes information on firms outside of Veneto if the worker has been employed in one such firm.[8]

3. The last register links the firm and worker registers. A private firm has to report the payment to its workers and the corresponding labour contract to the National Institute of Social Security (INPS) so that the authority could calculate each worker's social security contribution. Therefore, the register contains accurate information on annual earnings[9] (with no top-coding), weeks worked, occupation (white-collar, blue-collar, manager, apprentice), type of contract (fixed-term or open-ended) and type of working schedule (full-time or part-time). Annual earnings have been inflation-adjusted to the price level of the year 2003.

Sample selection We use all workers and firms within the Veneto region only. In other words, we use the working population data within the Veneto labour markets. We only use the period from 1982 to 2001 because the information on working weeks before 1982 is not accurate (Battisti, 2017). Besides, we have a few minimal restrictions, mainly following the standard practice in the literature. First, we keep only a worker's primary job if he or she works in multiple positions,[10] and we restrict the working ages from 16 to 65. Also, we exclude part-time jobs and apprentices because their wages cannot be compared to regular full-time employment (as we have no information on working hours). Since we are interested in coworkers, we drop single-worker firms. Following the practice of Cornelissen

[5] More precisely, we observe the working history before and after he/she worked in Veneto as long as it is within Italy. Besides, around 80 percent of the workers have never worked outside Veneto.

[6] There are two important related points. First, the public sector is not included in this database. Second, the firm is not at the establishment level. It might be ideal to use establishment-level data for our analysis, but using firm-level data would not make a difference for two reasons. First, most firms, especially in our sample period where the franchise is not typical, are single-establishment firms. Second, the firm size is typically small, with a median size of six workers, and firms with 200 or fewer employees take up around 90 percent of the observations.

[7] We also use the identified national tax code to link the data to balance sheet information from Bureau Van Dijk-AIDA, a firm-level dataset that covers a fraction of the firms in Veneto after 1996.

[8] In the main analysis, we will keep firms in Veneto only, as for firms outside it we do not have complete information on coworkers.

[9] The annual earnings equal full net earnings, plus all kinds of pecuniary compensation, grossed up with labour income taxes and social security contributions on the employee.

[10] Specifically, if a worker has two or more employment contracts in a year, we keep the job with the highest annual earnings or the highest number of weeks worked. We break the very few ties (less than 1 percent of the data) by choosing randomly the primary job.

et al. (2017) and Caldwell and Harmon (2019), we also restrict the firm size to be smaller than 5000.[11]

Peer group definition We define the peer group as all the workers employed in the same firm with the same occupation in a given year, where the occupation is given by broad professional levels (blue-collar, white-collar and executive).[12]

Descriptive statistics Table 2.1 presents descriptive statistics of the sample used in the analysis. We have 17.7 million person-year observations, 2.5 million workers and 169 thousand firms. Full-time workers earn annually on average 33.3 thousand euros (in 2003 prices) and the mean weekly wage is 744 Euros. The average number of weeks worked is 42 (and the median is 52, indicating that the median worker is a full-time and full-year one). Average age is 34.5 and the tenure is 2.5 years. As expected, firms are small, reflecting the structure of the Italian labour market, with a mean firm size of 17 employees and a median of 6. Similarly, the peer group size – that is, workers in the same occupation and firm – is on average 12 and 4 at the median. The mean number of movers per firm is 4, indicating that on average, annually, 4 workers move to other firms or to non-employment. Overall, 61% of workers change job at least once throughout the whole period of analysis. The share of women is 36%, reflecting the relative low female labour force participation. The majority of worker is employed in blue-collar occupations (70%) and are on open-ended contracts (3%). More than half of the workers are employed in manufacturing (53%).

Motivating evidence There exists considerable heterogeneity in wage profiles for workers employed in peer groups of different quality. To see this point, we run a canonical two-way fixed effects AKM regression (Abowd et al., 1999), i.e.

$$w_{i,t} = \alpha_i + \psi_j + \mathbf{x}'_{it}\gamma + \epsilon_{it}, \tag{2.1}$$

where $w_{i,t}$ are log weekly wages of individual i at time t, α_i are worker fixed effects, ψ_j are firm fixed effects, $\mathbf{x}_{it}$ contains a cubic polynomial in age and tenure, a dummy for women and a dummy for part-time workers. ϵ_{it} is an error term. Using the estimates of α_i from equation (2.1), we compute for each worker the leave-one out average peer quality as the

[11] Lastly, due to the identification requirement in the AKM analysis below, we need to restrict the sample to the largest connected set (Abowd et al., 1999), which takes up around 97 percent of the sample.

[12] Compared to Cornelissen et al. (2017), who use a similar definition to ours, we have a lower detail of occupational categories, as we do not have detailed occupation codes. However, given the small average size of firms in Veneto – especially in the period of time we focus on – we end up with peer groups that are comparable in size to those reported in Cornelissen et al. (2017): the average peer group size in our data is 12, whereas in their paper is 9.3.

Table 2.1 Summary statistics

	(1) Mean	(2) S.D.	(3) Median
Annual earnings	33350.06	40250.33	31730
Weekly wage	744.38	1652.81	652
Weeks worked	42.41	15.24	52
Age	34.54	10.69	32
Tenure	2.45	2.58	2
Firm size	17	75	6
Movers per firm	4	26	1
Peer group size	12	54	4
Mover	0.61	0.49	
Woman	0.36	0.48	
Blue-collar	0.70	0.46	
White-collar	0.29	0.46	
Executive	0.01	0.10	
Open-ended contract	0.97	0.16	
Temporary contract	0.03	0.16	
Manufacturing	0.53	0.50	
Services	0.03	0.18	
Construction	0.06	0.23	
Domestic work	0.08	0.28	
Finance	0.04	0.18	
Health	0.03	0.16	
Transports	0.04	0.18	
Wholesale	0.10	0.29	
Accomodation	0.02	0.13	
Other	0.03	0.16	
Person-year observations		17,723,260	
Number of workers		2,531,411	
Number of firms		168,613	

Notes. The table reports means, standard deviations and medians of each variable in columns (1) to (3), based on the largest connected set of workers and firms from the Veneto Worker History Panel. See text for details about data and sample restrictions.

average worker effect of his or her coworkers in a given firm and year.[13] We then show descriptively how wage growth varies for workers joining a firm with better or worse peers,

[13]In other terms, we compute for each worker i the following quantity:

$$\bar{\alpha}_{-i,t} = \frac{1}{|N_{-i,j,t}|} \sum_{k \in N_{-i,j,t}} \alpha_k$$

where $|N_{-i,t}|$ is the number of coworkers of i in firm j.

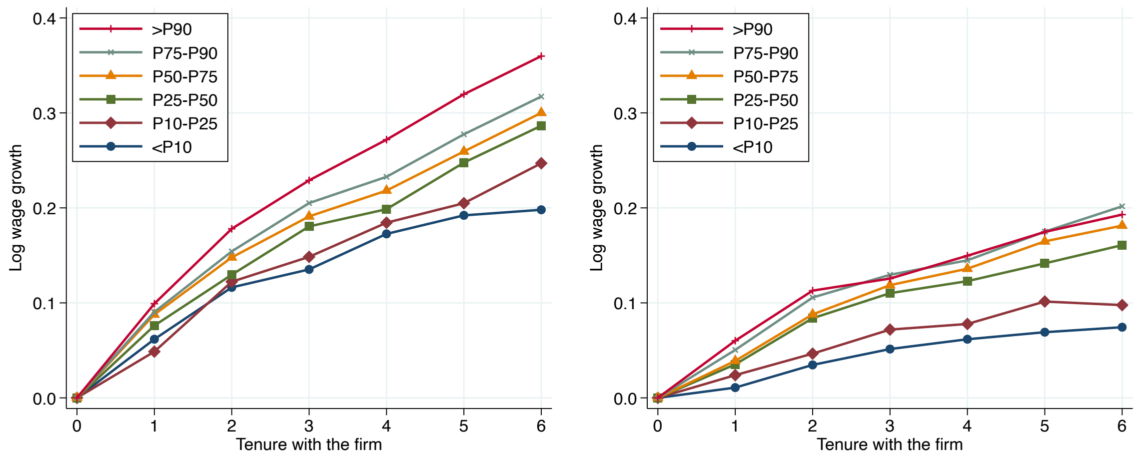

(A) White-collar, hired at age 25 (B) Blue-collar, hired at age 25

Fig. 2.1 Log wage growth for workers entering the labour market in different quantiles of the peer quality distribution, by years of tenure and occupation

Notes. The figure reports the difference between average log weekly wages in tenure years 1 to 6 and baseline log weekly wage in tenure year 0 for workers entering the labour market in different quantiles of peer quality, defined as the average firm-level leave-one out worker effect estimated from equation (2.1). The sample includes workers hired at age 25. Panels (A) and (B) show results for white- and blue-collar workers, respectively.

by grouping workers into different quantiles of the peer quality distribution. The results are reported in Figure 2.1 for white-collar and blue-collar workers in panel (A) and (B), respectively. The figure depicts the growth in log weekly wages (i.e. the difference relative to the entry log wage) for workers entering the firm at age 25 by tenure with the firm and by peer quality, grouped in six discrete groups based on percentiles of its distribution: below 10, 10-25, 25-50, 50-75, 75-90 and above 90. The figure shows that there is wide heterogeneity in wage growth for workers joining a firm in different parts of the peer quality distribution, for both blue- and white-collar employees. Six years after joining the firm, white-collar workers in the top decile experience 0.16 ($= 0.36 - 0.20$) log points larger wage growth than a worker in the bottom decile of peer quality, whereas for blue-collar workers the additional wage growth equals 0.12 ($= 0.19 - 0.07$) log points. The difference in wage growth between workers with better and worse peers signals the contribution of the workplace environment as a driver of wage growth. However, we can say little about the direct effect of peers on wage growth as this descriptive analysis does not rule out sorting on productivity between workers and firms. This evidence likely suggests that part of the differential in wage growth is attributable to peers and part to firms. In the next section, we adopt a more formal empirical strategy to separately measure the contribution of both to wages.

2.3 An AKM Approach to Estimate Peer Effects

In this section, we explore the overall effect of coworkers on future wages. In particular, we build our empirical strategy on the canonical AKM model (Abowd et al., 1999), by incorporating the average peer quality and additional fixed effects to better deal with the sorting of workers across firms and occupations. Besides, we briefly discuss how we adopt the novel method developed by Hong and Mikkel (2020) for estimation. We motivate our empirical analysis by developing a simple two-period principal-agent model, based on Cornelissen et al. (2017) and reported in Appendix B.1.

2.3.1 Empirical Strategy

In our regression specification, we follow Cornelissen et al. (2017), who build on Abowd et al. (1999), as expressed in equation (2.2):

$$w_{i,t+h} = \alpha_i + \beta\bar{\alpha}_{-i,t} + \mathbf{x}'_{it}\gamma + \psi_{jt} + \eta_{ot} + \theta_{oj} + \varepsilon_{it}, \tag{2.2}$$

where

$$\bar{\alpha}_{-i,t} = \frac{1}{|M_{-it}|}\sum_{k\in M_{-it}} \alpha_k, \text{ and } M_{-it} = \{k : oj(k,t) = oj(i,t),\, k \neq i\}.$$

In equation 2.2, $w_{i,t+h}$ is the log weekly wage at time $t+h$, where $h \geq 0$. α_i is the worker fixed effect, which measures the portable component of earnings and is a proxy for quality or innate ability of a worker. $\alpha_{-i,t}$ is the average coworker's quality at time t.[14] $\mathbf{x}_{i,t}$ is a set of individual time-varying characteristics, including age, age squared, tenure, tenure squared, and a dummy on whether tenure is larger than ten years. ψ_{jt}, η_{ot}, θ_{oj} are firm-year, occupation-year, firm-occupation fixed effects. β is our parameter of interest. It describes how contemporary coworker quality could change future wages.

2.3.2 Identification Challenges

As discussed in Arcidiacono et al. (2012), we face three challenges in the identification of peer effects: (*i*) the reflection problem; (*ii*) the non-random sorting of workers across peer groups; (*iii*) the presence of unobserved correlated shocks. The reflection problem was first introduced by Manski (1993) who referred to it when discussing the problem of identifying the peer effect from contemporaneous peer effort or productivity (e.g., wages). For example, in a firm, effort of peers influences a worker's effort, who in turn affects his or her peers'. In the presence of such "reflection" it is difficult to identify the peer effect. As suggested by Cornelissen et al. (2017), using long-term predetermined characteristics of

[14] $|.|$ defines the modulus of the coworker vector M_{-it}; hence, it represents the number of coworkers.

peers solves the reflection problem as it avoids contemporaneous productivity measures interacting with each other. For this reason, we measure peer quality with the leave one out average AKM worker fixed effect, which we interpret as a proxy of peers' long-term productivity. We address the endogenous sorting of workers across peer groups and the presence of unobserved shocks by controlling for a rich set of fixed effects. Peer quality may be correlated with worker's wages if high-quality workers sort into high-quality peers. We therefore control for worker fixed effects α_i in equation (2.2), so to estimate the impact of within-individual changes in peer quality on wages. Moreover, peer quality can be correlated with worker's wages in the presence of sorting between high-quality workers into high-quality firms or occupations. For this reason, we include firm-time fixed effects ψ_{jt} that control for firm-level shocks, occupation-time fixed effects η_{ot} that control for different time trends in occupation-specific pay and occupation-firm fixed effects θ_{oj} that control for the possibility that firms pay higher wages to specific occupations.[15]

Even in the presence of the rich set of fixed effects discussed above, the estimate of β can still be biased if there exist unobserved background characteristics that vary at the occupation-firm-time level (i.e. at the peer group-time level) that are correlated with changes in peer quality observed between consecutive periods. To see this point, observe that there are two sources of variations for the identification of β. For job switchers, peer quality changes when they move to another firm. For job stayers, peer quality changes when other workers join or leave the peer group.[16] Both these variations entail changes in the peer group that allow identification of β.

We can show the potential endogeneity of peer quality to unobserved time-varying shocks as follows. First, we follow Cornelissen et al. (2017) and focus on job stayers only. We denote with $\widetilde{\alpha}_{-i,t}$ the residual peer quality after controlling for time-varying observables $\mathbf{x}'_{it}$ and occupation-time effects η_{ot}. Then, first differencing equation (2.2) removes the time-invariant worker (α_i) and firm-occupation fixed effects (θ_{oj}),[17] but does not remove firm-level time-varying effects ($\triangle\psi_{jt}$): $\triangle w_{i,t+h} = \beta\triangle\widetilde{\alpha}_{-i,t} + \triangle\psi_{jt} + \triangle\varepsilon_{it}$. Such effects can be removed by exploiting variation between occupations, which experience within firm different changes in peer quality. Assuming there are only blue- and white-collar occupations (discarding executives, for the moment) and denoting them as o and o', respectively, we can net out firm-level shocks by taking a second difference across occupations. Taking firm-level averages, we have:

[15] One concern with our empirical strategy is that we control for contemporaneous fixed effects even when estimating equation (2.2) at time $t+h, h>0$. As a matter of fact, using contemporaneous fixed effects has little impact on our findings.

[16] Note that our definition of the peer group allows us to take into account also within-firm mobility due to promotion, e.g. from blue-collar to white-collar or from blue- and white-collar to executive, as we define the peer group as workers in the same firm *and* occupation.

[17] The first difference also removes any unobserved time invariant job match effect.

$$\triangle\overline{w}_{j,o,t+h} - \triangle\overline{w}_{j,o',t+h} = \beta\left(\triangle\overline{\overline{\alpha}}_{j,o,t} - \triangle\overline{\overline{\alpha}}_{j,o',t}\right) + \triangle\overline{\varepsilon}_{j,o,t} - \triangle\overline{\varepsilon}_{j,o',t}.$$

Therefore, if $\mathrm{cov}\left(\triangle\overline{\overline{\alpha}}_{j,o,t} - \triangle\overline{\overline{\alpha}}_{j,o',t}, \triangle\overline{\varepsilon}_{j,o,t} - \triangle\overline{\varepsilon}_{j,o',t}\right) = 0$, i.e. if there is no correlation between unobserved time-varying occupation-firm specific shocks and changes in peer quality between consecutive periods, we are able to retrieve consistent estimates of β. This condition basically requires that different occupations experience the same time shock within firm or, in other terms, that different occupations are on parallel trends within firm before a change in peer quality happens (akin to a standard difference-in-differences). Violations of such parallel trend assumptions are yet possible. For example, the firm may decide to invest in automation which complements white-collar workers and substitutes for blue-collar workers. Assuming that skills and occupations are correlated, and therefore white-collar workers are more skilled, this would raise peer quality and firm output (and therefore wages) simultaneously, leading to an upward bias in the estimate of β. The opposite would be true if a firm decides to divest in some occupation-specific technology which would decrease peer quality and firm wages simultaneously biasing downwards the estimate of β.[18]

The discussion so far implicitly assumed that peer quality was an observed quantity, which in fact is not. The following section discusses how we estimate peer quality and how we use such estimate to measure its impact on contemporaneous and future wages.

2.3.3 Estimation of Peer Effects

There are at least two main difficulties to estimate β in Equation 2.2. The first one is that the worker fixed effect needs to be estimated, but at the same time, the average coworker quality is a function of the worker fixed effects. The other difficulty comes from the high dimensionality of the fixed effects, making it hard to solve the system. The literature has traditionally solved these issues by employing the iterative procedure pioneered by Arcidiacono et al. (2012), who estimate peer effects by first estimating Equation 2.2 setting $\beta = 0$ and then using the estimated worker fixed effects to compute β in following iterations. We depart from the literature and employ instead the novel estimation method developed by Hong and Mikkel (2020), which we now discuss in greater detail.

First, we write Equation 2.2 in matrix form:

$$w = X\delta + C\delta\beta + \varepsilon, \tag{2.3}$$

where $w \in \mathbb{R}^n$ is the wage. $X \in \mathbb{R}^{n\times k}$ contains all the fixed effects and time-varying individual characteristics. $C \in \mathbb{R}^{n\times k}$ is a coworker averaging matrix (see Appendix B.2 for details on

[18] One way to deal with time-varying occupation-firm shocks would be to include an occupation-firm-time effect and therefore exploit within peer group variation to estimate β. However, as discussed in Cornelissen et al. (2017) this would limit the identification to job stayers only. For this reason, we prefer our baseline specification, which exploits variation coming from both stayers and movers.

the construction of C). $\delta \in \mathbb{R}^k$ is a nuisance parameter that corresponds to the coefficients of the fixed effects. $\beta \in B$ is our parameter of interest, where B is a compact parameter space. Note that $C\delta$ is equivalent to $\bar{\alpha}_{-i,t}$ in Equation 2.2. In addition, we make the following assumptions:

- exogeneity, $\mathbb{E}[\varepsilon|\boldsymbol{X},\boldsymbol{C}] = \boldsymbol{0}$;
- homoskedasticity, $\mathbb{E}\left[\varepsilon\varepsilon'|\boldsymbol{X},\boldsymbol{C}\right] = \sigma^2\boldsymbol{I}_n$ where $\sigma^2 > 0$ is unknown;
- the designed matrix $X + C\beta$ has full rank k for any $\beta \in B$.

Our parameter of interest β is estimated by solving the objective function Q_n, i.e. it is the solution to the following inner minimization problem:

$$\hat{\beta} = \arg\min_{\beta \in B} Q_n(\beta) = \arg\min_{\beta \in B}\left\{\min_{\delta \in \mathbb{R}^2} \|w - X\delta - C\delta\beta\|^2/n\right\} \tag{2.4}$$

First, Hong and Mikkel (2020) proves that β is the unique minimizer of the population analogue to Q_n under the assumptions we impose, so consistency is ensured. We derive the moment condition for β by taking the first order conditions of β and δ in Equation 2.4 and get:

$$S_n(\beta) = w'MC\left(R'R\right)^{-1}R'w/n = 0, \tag{2.5}$$

where

$$R = X + C\beta,$$
$$M = I_n - R\left(R'R\right)^{-1}R'.$$

We solve numerically equation 2.5, by using sparse matrices and the conjugate gradient method, which allow to speed up the computation.[19]

2.3.4 Results

Estimation Table 2.2 reports summary statistics from the estimation of equation 2.2 for $h = 0$. The standard deviation of log weekly wages is 0.44. As typical in the literature that studies decompositions of the variance of wages, most of the variability in wages is accounted by variability in worker fixed effects (see Card et al., 2018, for a review). The standard deviation of the average peer fixed effect is 0.18 and, more importantly, the correlation between worker fixed effect and average peer fixed effect is 0.55 highlighting a strong degree of sorting among coworkers, documented as well, for example, in Lopes de Melo (2018). The correlation between worker fixed effect and firm-time fixed effect is slightly

[19] In particular, the sparse matrix helps with the standard matrix operation and memory issues, and the conjugate gradient method could vastly improve the matrix inversion speed.

Table 2.2 Standard deviation of wages and fixed effects and correlation between fixed effects

Statistic	Value
Standard deviation log weekly wages	0.436
Standard deviation worker fixed effect	0.269
Standard deviation peer fixed effect	0.178
Standard deviation occupation-time fixed effect	0.065
Standard deviation firm-occupation fixed effect	0.103
Standard deviation firm-time fixed effect	0.137
Standard deviation change of peer fixed effect between t and $t-1$	0.090
Standard deviation change of peer fixed effect between t and $t-1$ for movers	0.173
Standard deviation change of peer fixed effect between t and $t-1$ for stayers	0.066
Correlation worker fixed effect/peer fixed effect	0.551
Correlation worker fixed effect/firm-time fixed effect	-0.041
Correlation worker fixed effect/firm-occupation fixed effect	-0.001

Notes. The table reports summary statistics from the estimation of equation 2.2 for $h = 0$, based on the largest connected set of workers and firms from the Veneto Worker History Panel. See text for details about data and sample restrictions.

negative, one feature documented for the Italian labour market as a whole in Chapter 1 and Veneto specifically in Devicienti et al. (2019). The table also reports the standard deviation of the change in peer effects between consecutive years, as suggested by Cornelissen et al. (2017). The identification of β in equation 2.2 rests on changes in peer quality between subsequent years. Hence, one needs sufficient variation in peer quality to identify β. The standard deviation of the change in the average peer fixed effect equals 0.09. We also distinguish between movers and stayers: for the former, changes in peer quality happen because they move into a new peer group; for the latter, changes in peer quality happen if peers join or leave the current peer group. Not surprisingly there is larger variation in the change in peer quality for movers than for stayers, with standard deviations of 0.17 and 0.07, respectively, as for stayers it may happen that the peer group does not change at all between consecutive years. We corroborate this finding by plotting the density of the change in peer quality for movers and stayers in Figure B.1, which shows the existence of a mass around 0 for stayers (i.e. when the peer group does not change) and more variability for movers. Anyway, even for stayers, the standard deviation of peer quality changes amounts to approximately 38% of the overall variability in peer quality, indicating that we have enough variation in the data to identify the peer effect.

Main estimates of the peer effect Figure 2.2 shows our baseline results. Each dot in the graph represents the estimate β in equation (2.2) using the future wages as the dependent variable in each year ahead (h), where $h \geq 0$.[20] Vertical bars are 95 percent confidence intervals, retrieved from bootstrapped standard errors. The figure shows that the peer effect is large not only for the contemporaneous wage but also for the wages in the following years. A 10 percent increase in peer quality increases the contemporaneous wage by 2.6 percent[21] and next year's wage by 1.8 percent, which is similar to the size of the return to college (Nix, 2020). The effect gradually fades out to around 0.7 percent after five years. It is consistent with other papers that coworkers in the past three years play the most important role in wage growth (e.g., Caldwell and Harmon, 2019; Nix, 2020).

The peer effect can be a result of different factors. As highlighted in Cornelissen et al. (2017) peers may boost productivity and, therefore, wages by a mechanism of peer pressure, according to which a worker increases her own effort in response to increased effort by her coworkers. At the same time, workplace interaction is crucial for human capital accumulation, as workers transmit knowledge among each other which may make them more productive on-the-job. Our findings can be seen as a way of disentangling the two effects. Peer pressure is a mechanism through which the peer effect materializes on contemporaneous wages, but cannot be persistent over time. Hence, the finding of the peer effect on future wages may be more related to learning and knowledge transmission among coworkers – knowledge that is then incorporated in workers' wages through greater productivity, greater rent extraction or better outside options – rather than peer pressure. At the same time, the decaying pattern that we find suggests that the returns to having worked with better peers in the past are each year smaller. This may signal that most of knowledge transmission happens at $h = 0$, i.e. that the peer effect materializes mostly as a contemporaneous effect, whereas in subsequent years $h > 0$ the marginal returns to additional knowledge from peers are decreasing. This likely highlights the importance of team work as a fundamental mechanism behind peer effects, i.e. it is by working close to his/her peers that a worker feels motivated or pressured to exert more effort. Part of such effect percolates to future wages through knowledge acquired on the job, which contributes to human capital accumulation. It is also important to stress that our findings are conditional on tenure with the firm, which effectively capture time-varying accumulation of firm-specific human capital beyond knowledge acquired from peers. Finally, our finding cannot exclude network effects from interaction with better coworkers: working with better peers may deliver better relationships with suppliers, customers or links to other firms which increase workers' outside options.

[20] Different future wages are used as outcomes in separate estimations. In the cases when $h > 0$, workers who do not have wages in year $t + h$ are excluded. Therefore, each estimate, to some extent, is estimated on different samples. Specifically, $N = 17{,}723{,}260; 14{,}480{,}425; 12{,}387{,}027; 10{,}736{,}176; 9{,}319{,}245; 8{,}068{,}687$ for $h = 0, ..., 5$. We have also conducted a robustness check using a sample restricted to workers who have been working continuously for at least five years. The results are very similar, except that the estimates in the first two years are slightly smaller. Besides, the decreasing pattern remains. Results are available upon request.

[21] This results is smaller than what found in Battisti (2017) on the same data. The difference is likely due to the different algorithm used to compute the peer effect and to different sample selection choices.

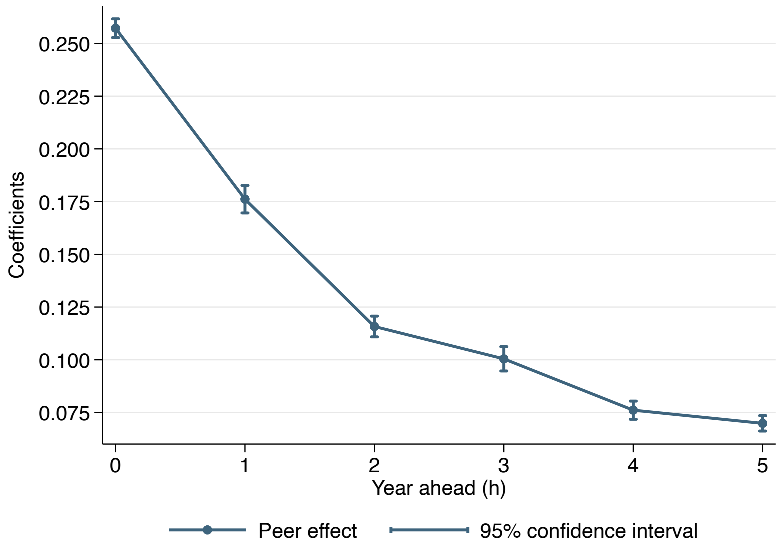

Fig. 2.2 Peer effect on future wages (β)

Notes. $N = 17{,}723{,}260; 14{,}480{,}425; 12{,}387{,}027; 10{,}736{,}176; 9{,}319{,}245; 8{,}068{,}687$ for $h = 0, ..., 5$. The figure reports the estimates of β from equation (2.2) using the future wages as the dependent variable in each year ahead (h), where $h \geq 0$. Vertical lines are 95% confidence intervals, obtained from bootstrapped standard errors

One concern about our finding is the effective ability of workers to bargain over their wages, which requires some degree of wage flexibility. Italy is characterized by a two-tier wage setting scheme, with collective contracts defining wage floors at the sector-level. At the firm-level, workers can bargain, individually or through employees' organizations, top-up agreements that increase their compensation. Moreover, part of the pay that we observe is composed of bonuses and premiums that employees receive on top of the basic compensation established by either the collective or firm-level contract. These wage premia are sizeable: Bartolucci et al. (2018) document that in the Veneto sample, wage premia are about 24 percent above the basic pay at the median. Hence, wage setting in the Italian context should be flexible enough to incorporate peer effects.

Heterogeneity We explore heterogeneous effects across different groups of workers and firms.[22] Figure 2.3A separates movers and stayers, where the former are workers that change employer in at least one year $t+h, h \geq 0$, whereas the latter are workers employed in the same firm throughout the whole time window. The figure shows that peer effects are more important for movers: after five years, the coefficient for them is 0.5 percent, whereas the coefficient for stayers is 0.1 percent.[23] The difference between movers and stayers can reflect differences in the ability to learn from peers for these two types of workers or the

[22] We conduct the heterogeneous analysis using the pre-estimated fixed effects from the baseline regression.
[23] The coefficients are smaller than those in the main analysis because the sample used in this case contains only continuing workers, that is workers observed continuously over the whole sample period.

endogeneity of mobility to learning chances in the incumbent and poaching firm: the latter may offer better learning and, therefore, wage prospects that the incumbent firm cannot offer. At the same time, we have highlighted that changes in peer quality are much more common for movers than for stayers (Figure B.1), hence movers may have better chances of acquiring knowledge as they move into new peer groups.

Figure 2.3B shows that there are no evident differences between firms of different size. Specifically, we divide firms among those with less than 10 employees, between 10 and 250 and more than 250. The evolution of peer effects is not dissimilar in the three groups, indicating that peer effects are an important channel of wage growth irrespective of firm size.[24]

Figure 2.3C shows how the effects differ across different tenure years. Specifically, we assign workers at $h = 0$ to three discrete categories of tenure with the firm – 0-2 years, 3-5 years and 6 or more years – and explore how peer effects change for each group of workers. The results illustrate a clear pattern that peers matter the most for low tenure workers, while the effect decreases as one experiences more years in the same firm. The finding is consistent with a learning process: there is more room for a new hire (a worker with low tenure) to learn in a firm. If the new hire has better peers as she enters the firm, she will more likely learn quicker and thus boost future wages.

A similar pattern arises when we explore the heterogeneous impacts across different age groups, as shown in Figure 2.3D. Like the finding above, the effects are higher for younger workers (below age 30). For older workers, the peer effect is lower in levels. The decay over time is similar across age brackets. Again, the same logic applies here. There is greater space for younger workers to learn as they enter the labour market. A better peer group could help them accumulate human capital faster, thus expedite wage growth.

2.4 Mobility, Workers' Quality and Wage Growth

As mentioned earlier, the identification of β is achieved through the following channels:

1. for job stayers, the peer quality changes when a worker *enters* the peer group or when a worker *leaves* the peer group;

2. for job switchers, the peer quality changes as they move to another firm.

We separately study these channels, following the wage trajectories of workers and coworkers around mobility episodes in our data. Specifically, we study channel 1 by setting up a coworker-level event study around mobility of a worker and analyse the evolution of wages of coworkers of the mover in the destination and origin firms. We distinguish three types of movers: high-quality, low-quality and average-quality, where the quality of the

[24]We find the same results if, instead of looking at firm size, we focus on peer group size.

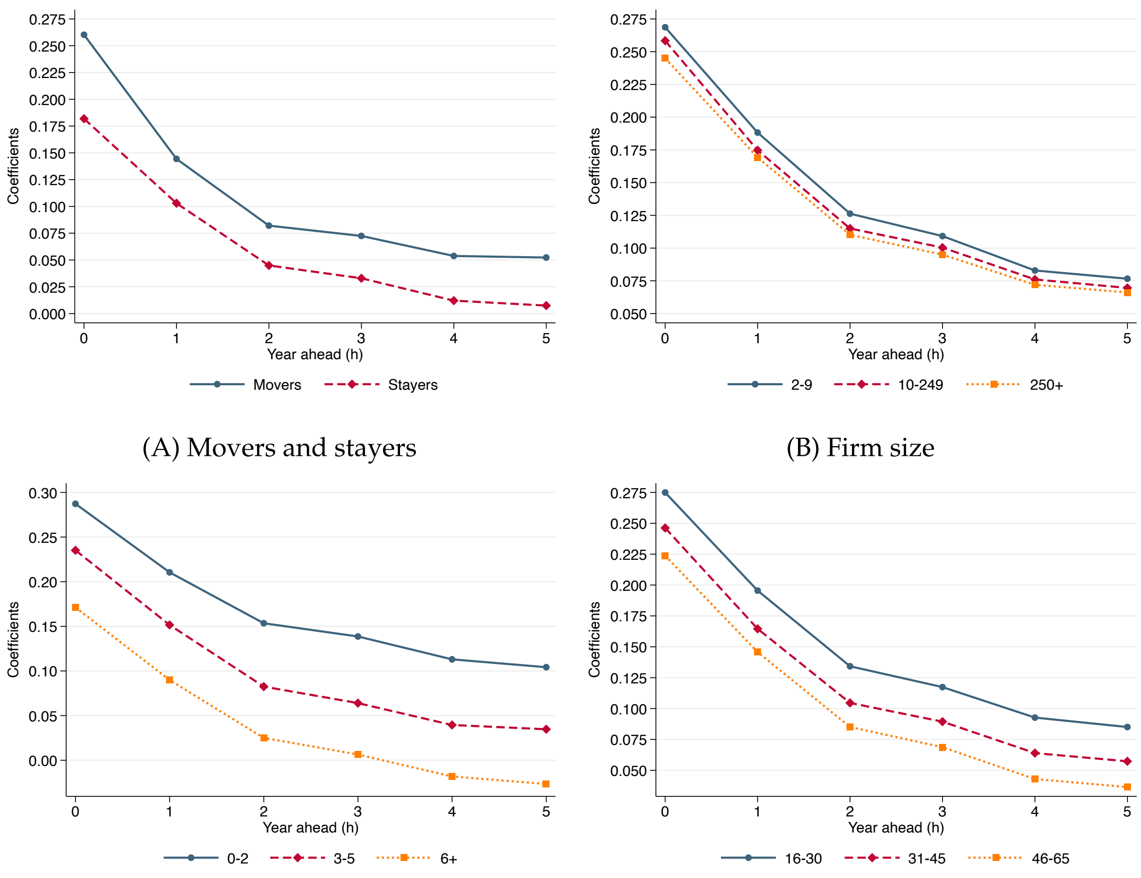

(A) Movers and stayers

(B) Firm size

(C) Tenure brackets

(D) Age brackets

Fig. 2.3 Heterogeneous peer effects across tenure and age brackets

Notes. The figure reports the estimates of β from equation (2.2) using the future wages as the dependent variable in each year ahead (h), where $h \geq 0$, focusing on different groups of workers. Panel (A) shows estimates for movers and stayers ($N = 4{,}587{,}715; 6{,}096{,}654$), panel (B) for workers in firms of different size ($N = 2{,}895{,}186; 10{,}332{,}778; 4{,}495{,}296$ for firms size 2-9, 20-249 and 250+, respectively), panel (C) for different tenure brackets ($N = 8{,}997{,}227; 4{,}174{,}367; 4{,}551{,}666$ for bracket 0-2, 3-5 and 6+, respectively) and panel (D) for different age brackets ($N = 7{,}678{,}846; 6{,}733{,}816; 3{,}310{,}598$ for bracket 16-30, 31-45 and 46-65, respectively).

mover is based on a comparison of her worker fixed effect with the average peer fixed effect of coworkers from equation 2.2. Specifically, we classify firms as hiring a high-quality or low-quality worker if her worker fixed effect is 10 percent higher or lower, respectively, than the average peer fixed effect of coworkers in the destination firm. We classify, instead, firms as hiring an average-quality worker if her worker fixed effect is between -10 and $+10$ percent of the average peer fixed effect in the destination firm. When we focus on coworkers in the origin firm, we use the same classification, based on the comparison between leavers and coworkers in the origin firm.

We study channel 2 in a similar fashion. We follow the wage trajectories of workers moving into peer group of different quality, where, again, we define a peer group as high-, average- or low-quality if the mean peer fixed effect of coworkers is 10 percent higher, between -10 percent and $+10$ percent, or 10 percent lower than the worker fixed effect of the mover.

We now turn to a more formal illustration of our empirical strategy.

2.4.1 Empirical Strategy

Coworker-level Event Study

Empirical design Figure 2.4 illustrates the definition of treatment and control groups in the event-study. Primarily, we define the event as a worker who moves to a firm and stays there for three years. We choose three years because we think that it takes time for the mover to have an effect on her coworkers' wages, especially because knowledge may need some time to be transmitted.[25] We then use as an outcome wages of coworkers of the mover (thus excluding the mover from the sample) in either the destination or origin firm, when we study the impact of the move on the "new" or "old" coworkers, respectively: we define these two analyses the "hire" design and the "leave" design.

We restrict the sample firms to be observed for eight consecutive years, in which there is no worker mobility in the four years before and three years after the event. While we risk losing generalization by such a substantial restriction, we believe it is essential for our analysis. First, as shown in Section 2.3.4, the coworker's effect substantially decreases after three years. Second, we need the pre-event period to examine the pre-event parallel trend assumption from a methodological perspective. Finally, since learning and knowledge spillover take time to be reflected in wages, we need a few years after the event to show the wage trajectory after the new worker enters.

In order to measure the impact of mobility on coworkers' wages we cannot simply compare outcomes in firms that choose to hire a worker to those in firms that choose not to hire any worker, as the worker flows are very likely endogenous to firm performance and, ultimately, workers' compensation. We overcome this issue by selecting only firms that hire a given worker in a specific year and compare firms hiring high-quality or low-quality workers to firms hiring average-quality workers. Similarly, when we look at the impact of mobility on coworkers' wages in the origin firm we select only firms that separate from a worker in a given year. Specifically, in the first case, we define the treatment groups as the coworkers in the firms that hire a high- or low-quality worker in period $t = 0$, where high-quality and low-quality workers are defined as detailed above, i.e. as workers with worker fixed effect that is 10 percent higher or lower than the average peer effect at the firm. We denote the group that hires a high-quality worker as treatment 1 in Figure 2.4 and the the group that hires a low-quality worker as treatment 2. We define the control group as the coworkers in the firms that hire an average-quality worker, whose ability is similar

[25] When the event year has multiple workers entering, we exclude the firm from the analysis.

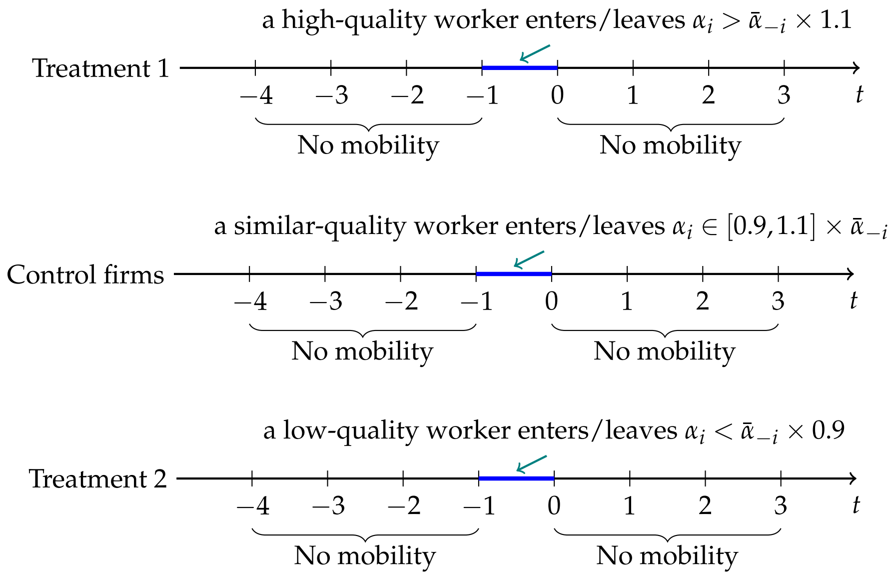

Fig. 2.4 Treatment and control groups in the event-study analysis, hire and leaver design

Notes. The diagram shows the empirical design behind the coworker-level event study estimated in equation (2.6). *Treatment 1* refers to mobility of high-quality workers, whose worker effect is more than 10% larger than the average peer effect in the origin or destination firm. *Treatment 2* refers to mobility of low-quality workers, whose worker effect is more than 10% smaller than the average peer effect in the origin or destination firm. *Control firms* refer to mobility of average-quality workers, whose worker effect is between −10% and +10% of the average peer effect in the origin or destination firm.

(within 10 percent difference) to the workers in the firm. When we analyse the effect on "old" coworkers, we maintain the same definitions, but the reference for defining high-, similar- and low-quality workers are peers in the origin firm.

Propensity score matching A critical issue that prevents us from identifying the effect is that a worker's mobility decision is not random. For example, the decision to hire a high-quality worker might be endogenous to firm performance, which also affects a worker's wage growth.[26] While there is no perfect remedy for this, we construct comparable firms between the treatment and control groups through ex-ante propensity score matching. The implicit assumption is that similar firms have similar hiring strategies, leading to a quasi-random hiring on average, such that the only difference between firms hiring a high- or low-quality worker rather then an average-quality worker is precisely the worker quality.[27]

We estimate the propensity score using a wide range of firm-level variables and some industrial and geographic variables. Specifically, we match on the following set of covariates

[26] Anecdotally, a firm may decide to hire a high-quality worker because he is complementary to some technology the firm decides to invest in. Such an investment could raise the productivity of the firm and eventually compensate all employees with wage raises.

[27] Our analysis will still be biased if the decision to hire is based on firms' unobservables which we cannot control for. We try to minimize the risk of the presence of such bias by including firm fixed effects in the regression analysis.

at time $t = -3$: the AKM firm-time effects and AKM average worker effects estimated from Section 2.3, the average age of employees, the share of female workers, the share of blue-collar workers, firm size, firm age, sales, value-added,[28] industry dummies, and province dummies. We also match average weekly wages at time $t = \{-4, -3, -2\}$. Besides, we use the single nearest neighbor matching without replacement to match the treatment groups 1 and 2 with the control group, separately. In other words, two different matched control groups are respectively comparable to treatment groups 1 and 2. Tables B.1 and B.2 in the Appendix report the mean differences and the p-values of the differences in the covariates we have used for matching for the hire and leaver design, respectively, for both comparisons of high-quality vs average-quality (columns 1-4) and low-quality vs average-quality (columns 5-8). Both tables highlight the presence of significant differences in observables between treated and control firms before matching, which, apart for very few exceptions, disappear in the matched samples. Moreover, we check the common support assumption by plotting the density of the propensity score in Figure B.2 and B.3 in the Appendix for the analysis on the entry of a worker and on the leave of a worker. In both cases, and for both treatments 1 and 2, there is a wide overlap in the propensity score densities.

Overall, when studying the entry of a worker in a firm, our sample consists of 2,164 firms hiring a high-quality worker, matched with the same number of firms from the control group. 1,848 firms have hired a low-quality worker, and they are matched with the same number of firms from the control group. The two matched samples consist of 285,350 and 238,046 person-year observations, respectively. When we focus on the leave of a worker from a firm, we have 2,905 firms where a high-quality worker leaves and 2,046 where a low-quality worker leaves, both matched with the same number of firms in the control group. The two matched samples consist of 390,135 and 234,016 person-year observations, respectively.

Event study On the matched sample of firms, we use the following event study specification to analyse the impact of a high- or low-quality worker's entry or leave on past and current coworkers' wages.

$$w_{-i,j,t} = \eta_t + \psi_j + \sum_{k \neq -1} \beta_k (Treat_j \times \mathbf{1}\{t = k\}) + \epsilon_{-i,j,t}, \tag{2.6}$$

where $w_{-i,j,t}$ is the log weekly wage of coworkers, excluding the mover, in period t and firm j, where firm j is the firm the worker joins when in the hire design and it is the firm the worker leaves in the leaver design. η_t and ψ_j are year and firm fixed effects, respectively. $\epsilon_{-i,j,t}$ is an error term. $Treat_j$ is a dummy variable for treated firms. The coefficients of

[28] Sales and valued-added variables are merged from the external balance-sheet firm-level database, AIDA. However, there is a good portion of firms that are not covered in AIDA. To utilize the information from balance-sheet data, we impute the missing value and create a dummy to indicate the missing observations.

interest are β_k, which measure the differential impact of hiring a high- or low-quality worker relative to hiring an average-quality worker on wages in each period k.[29]

Worker-level Event Study

Empirical design In a very similar fashion we also study the impact on the wages of movers of joining peer groups of different quality. We have highlighted in section 2.2 how the variation in the change in peer quality experienced by movers is generally greater than that experienced by stayers. We therefore examine how moving into a high-quality or low-quality peer group relative to an average-quality peer group impacts on wages of the mover. We use a very similar empirical strategy to that outlined before, but we match on workers' characteristics, rather than firms', such that, conditional on observables, the only difference between movers is the peer group they end up joining. We follow movers for a total of 10 years, i.e. between 4 years before the move and 5 years after, focusing on workers moving once and not changing employer in the 4 years before the move and in the 5 years after, as described by the scheme in Figure 2.5. As before, we compare movers into high-quality (treatment 1) and low-quality (treatment 2) peer groups with movers into average-quality peer groups (control), where the comparison of own quality and peer quality is based again on the worker and average peer fixed effect estimated in section 2.3.4.

Propensity score matching We perform ex-ante propensity score matching on the following set of workers' observables at time $t = -3$: the number of weeks worked, age, gender, tenure, occupation (blue-collar or white-collar), the decile of the AKM worker fixed effect, the AKM firm-time fixed effect, log firm size, sector dummies and province dummies. Besides, we match workers on log weekly wages at time $t = \{-4, -3, -2\}$. As before, we use a nearest neighbor matching without replacement (and caliper equal to 0.001). Table B.3 reports the mean difference and the p-values of the differences between treated and control workers, separately for treated workers that move into high-quality peer groups (columns 1-4) and treated workers that move into low-quality peer-groups (columns 5-8) for both unmatched and matched samples. The numerous differences between treated and control workers in the unmatched samples disappear when we perform matching. Furthermore, we keep only workers in the common support of the propensity score, which, as shown in Figure B.4, displays a wide overlap for both groups of workers.

After performing matching we are left with a sample of 31,102 workers (15,551 in the treatment and 15,551 in the control group) when studying movers into high-quality peer groups, for a total of 310,220 person-year observations. When we study movers into low-

[29] Our unit of observation is the coworker-year pair. If we use average outcomes at the firm-year level, excluding the new hire or the leaver, we obtain very similar results.

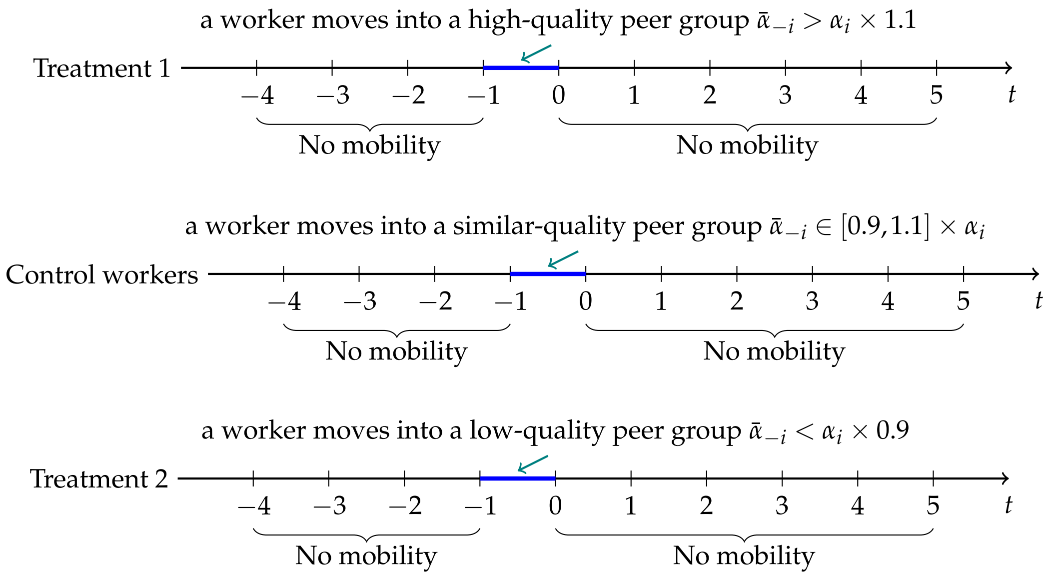

Fig. 2.5 Treatment and control groups in the event-study analysis, mover design

Notes. The diagram shows the empirical design behind the worker-level event study estimated in equation (2.7). *Treatment 1* refers to mobility into high-quality peer groups, whose average peer effect is more than 10% larger than the worker effect of the mover. *Treatment 2* refers to mobility of low-quality peer groups, whose average peer effect is more than 10% smaller than the worker effect of the mover. *Control workers* refer to mobility of workers into average-quality peer groups, whose average peer effect is between −10% and +10% of the worker effect of the mover.

quality peer groups we have a sample of 25,556 (12,778 in the treatment and 12,778 in the control group), for a total of 255,560 person-year observations.

Event study On the matched samples, we estimate the following event study regression:

$$w_{i,t} = \eta_t + \alpha_i + \sum_{k \neq -1} \gamma_k (Treat_i \times \mathbf{1}\{t = k\}) + \epsilon_{i,t}, \tag{2.7}$$

where $w_{i,t}$ is the log weekly wage of worker i in period t. η_t and α_i are year and individual fixed effects, respectively. $\epsilon_{i,t}$ is an error term. $Treat_i$ is a dummy variable for treated workers (either movers into high-quality peers or movers into low-quality peers). The coefficients of interest are γ_k, which measure the differential impact of moving into a high- or low-quality peer group relative to an average-quality peer group on wages in each period k.

2.4.2 Results

Coworker-level Event Study

When a high- or low-quality worker *enters* Figure 2.6 reports the event-study coefficients β_k for each $k \in \{-4, ..., +3\}$, for both treatment groups (coworkers in firms hiring a high-quality worker and coworkers in firms hiring a low-quality worker), relative to the control group (coworkers in firms hiring an average-quality worker). The pre-event parallel trend assumption holds as the effect before the event is small and statistically insignificant.

The post-event effects are quite different for the two treatments. Compared to firms hiring a similar-quality worker, peers in firms, which hire a high-quality worker, see a positive and significant effect on future wages. One year after the high-quality worker's entry, his or her coworkers' wage is 1.9 percent higher than that in control firms. The effect persists in the following years and reaches 3.1 percent after three years. In contrast, there is no effect in period 0. In other terms, the high- quality worker's entry does not impact the coworkers' wage immediately, but it takes some time for the peer effect to diffuse and be reflected in wages. On the other hand, when a firm hires a low-quality worker, the effect on his or her coworkers' future wages is slightly negative, but statistically insignificant. The knowledge spillover may play a role in explaining our findings. A high-quality worker, when joining a new firm, would be able to transmit knowledge to his or her coworkers, and therefore eventually drive up their wages in the following years. On the contrary, when a low-quality worker enters, the amount of knowledge she can transmit is much more limited and, therefore, it is less relevant for coworkers' future wages. As already discussed, learning from high-quality coworkers is not necessarily the only mechanism at play. The entry of a high-quality worker could affect coworkers' wages through alternative channels, e.g. peer pressure(e.g., Mas and Moretti, 2009)[30] or better network.[31]

To gather further insights into the mechanisms that determine our findings, we explore heterogeneous effects across different peer groups. Recall that we have defined the peer group in Section 2.3 as all the workers employed in the same firm with the same occupation. Figure 2.7A shows the effect of a high-quality worker's entry on his peer group and non-peer group (i.e. on coworkers in the same and different occupation, respectively). For the peer group, the effect is almost identical to the one in Figure 2.6: on average, the effect in the post-event window is 1.8 percent. On the other hand, there is no significant effect for the non-peer group, although the precision of the estimates is lower (the difference-in-differences point estimate is 0.4 percent with a standard error of 0.7). This finding runs against the hypothesis that network effects play a big role in determining the peer effect on wages, as better networks should benefit all workers equally in the firm. Instead, the finding is consistent with both knowledge spillover and peer pressure being important mechanisms through which better peers influence workers in the workplace. Figure 2.7B shows the same heterogeneous effects for coworkers of low-quality movers. We find a small and non-significant negative effect for coworkers in the same peer group and a positive (but

[30] For instance, the high-quality worker's coworkers may feel obliged to put more effort because of the higher competition coming from their peers.

[31] A high-quality worker has generally more connection to other firms (e.g., Caldwell and Harmon, 2019), which will help the firm expand business and, as a result, wages of all workers in the firm.

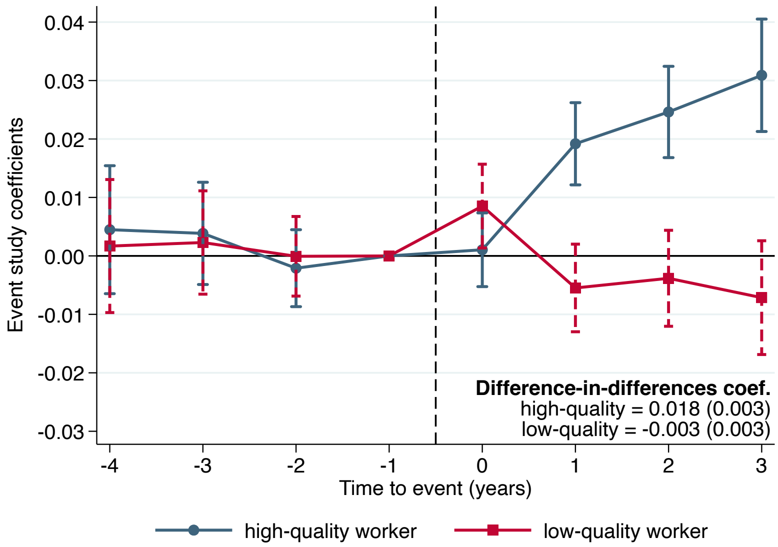

Fig. 2.6 The effect of a high-/low- quality worker's entry on coworkers' future wages

Notes. The figure reports estimates of β_k from equation (2.6), separately for firms hiring high-quality and low-quality workers relative to firms hiring average-quality workers. The dependent variable is the log weekly wage of coworkers in the destination firm. Vertical bars are 95% confidence intervals, obtained from cluster-robust standard errors at the firm level.

still imprecise) effect for non-peers. We cannot draw many conclusions as the estimates are hardly significant, but the positive coefficient on non-peers could signal some within-firm organizational changes that follow the low-quality hire and allow the promotion of better skilled workers to higher paying occupations.

We separately study the effect for high-wage and low-wage workers in Figure 2.7C and 2.7D. We define high-wage workers as those having a wage higher than the median of the sector they belong to. Figure 2.7C reports the estimates from the entry of a high-quality worker and shows that the benefit of working with such worker are equally shared between high-wage and low-wage workers. On the opposite, Figure 2.7D seems to suggest different wage trajectories for high-wage and low-wage workers as they collaborate with the new low-quality hire. While high-wage workers do not experience any wage change, low-wage workers experience some wage loss, although the estimates are quite imprecise and hardly significant.

When a high- or low-quality worker *leaves* We estimate again equation 2.6 on the matched sample of coworkers in firms where high-quality and low-quality workers leave compared to coworkers in firms where an average-quality worker leaves. The outcome variable is in this case the coworkers' wage in the origin firm, hence we measure the impact of workers' leaving a firm on their coworkers' wages. Figure 2.8 reports the event-study coefficients β_k for each $k \in \{-4, ..., +3\}$.

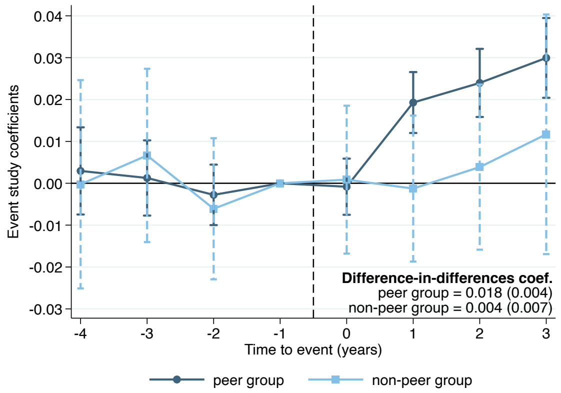

(A) High-quality entry, peers vs non-peers

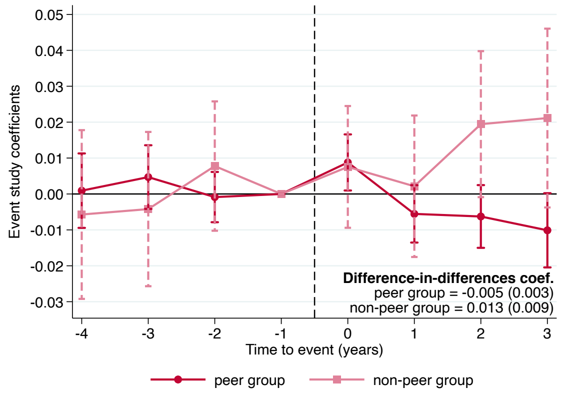

(B) Low-quality entry, peers vs non-peers

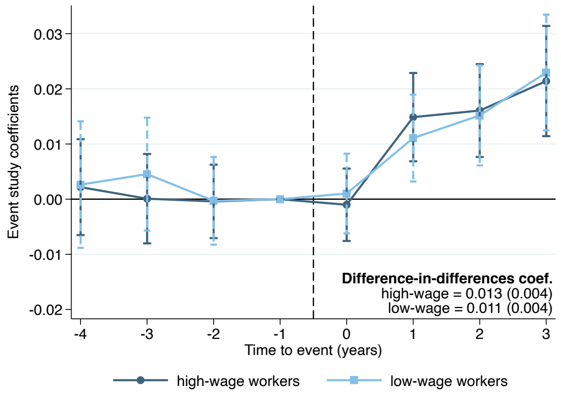

(C) High-quality entry, high- vs low-wage

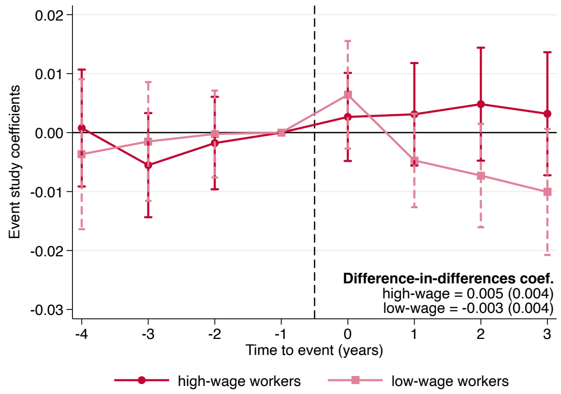

(D) Low-quality entry, high- vs low-wage

Fig. 2.7 Heterogeneous effects for peers and non-peers (top panels) and high-wage and low-wage coworkers (bottom panels), hire design

Notes. The figure reports estimates of β_k from equation (2.6) for different groups of workers. Panel (A) and (B) compare high- and low-quality hires, respectively, to average-quality hires, distinguishing the effect for workers belonging to the same or different peer group (i.e. to the same occupation). Panel (C) and (D) compare high- and low-quality hires, respectively, to average-quality hires, distinguishing the effect for high- and low-wage workers (i.e. above or below the median wage in the 2-digit sector they belong to). The dependent variable is the log weekly wage of coworkers in the destination firm. Vertical bars are 95% confidence intervals, obtained from cluster-robust standard errors at the firm level.

Our findings are somewhat symmetric with respect to those we find on the entry of workers. When a low-quality worker leaves, coworkers' wages increase by 2.4 percent on average, whereas the departure of a high-quality worker depresses coworkers' wages by -1.3 percent. When a firm separates from a low-quality worker, the average peer quality in a firm increases: this likely makes knowledge spillover easier in the firm, which eventually increases future wages. On the other hand, when a high-quality worker leaves, there are potentially two (opposite) effects. First, the leave of a high-quality worker will make the overall peer quality smaller, thus decreasing knowledge spillover efficiency. Second, the high-quality worker's human capital "left" into the firm might play a persistent effect over

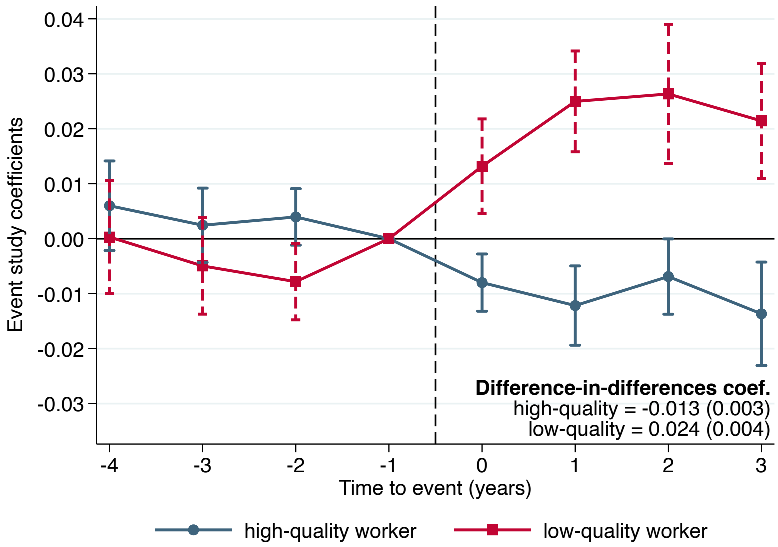

Fig. 2.8 The effect of a high-/low-quality worker's leave on coworkers' future wages

Notes. The figure reports estimates of β_k from equation (2.6), separately for firms separating from high-quality and low-quality workers relative to firms separating from average-quality workers. The dependent variable is the log weekly wage of coworkers in the origin firm. Vertical bars are 95% confidence intervals, obtained from cluster-robust standard errors at the firm level.

the next few years, which could help boost wage growth. Overall, the result suggests that the first channel might exceed the second, resulting in an overall negative effect.

We again explore heterogeneous effects distinguishing peers and non-peers in Figure 2.9A and 2.9B, respectively for high-quality and low-quality leaves. While the departure of a high-quality worker affects in a similar way peers and non-peers, when a low-quality worker leaves the firm, peers seem to benefit more than non-peers, although the difference between groups is not statistically significant: this evidence, nonetheless, suggests again that when a low-quality worker leaves knowledge transmission may become easier, especially among coworkers in the same peer group (i.e. in the same occupation).

When we distinguish between high-wage and low-wage workers (defined as before, based on the median weekly wage in the sector), we do not find significant differences between the two groups of workers when a high-quality worker leaves the firm in Figure 2.9C. On the contrary, low-wage coworkers gain more from the departure of low-quality workers than high-wage coworkers, especially in the first two years following mobility. If knowledge spillover becomes easier within the firm then low-wage workers may have more to gain from it with respect to high-wage workers. At the same time, within-firm competition for promotion may become lower after one competitor leaves, indicating that part of the effect may reflect occupational upgrading.[32]

[32]Unfortunately, we cannot study promotions as we lack detailed information in the data on the specific job title of each worker's contract.

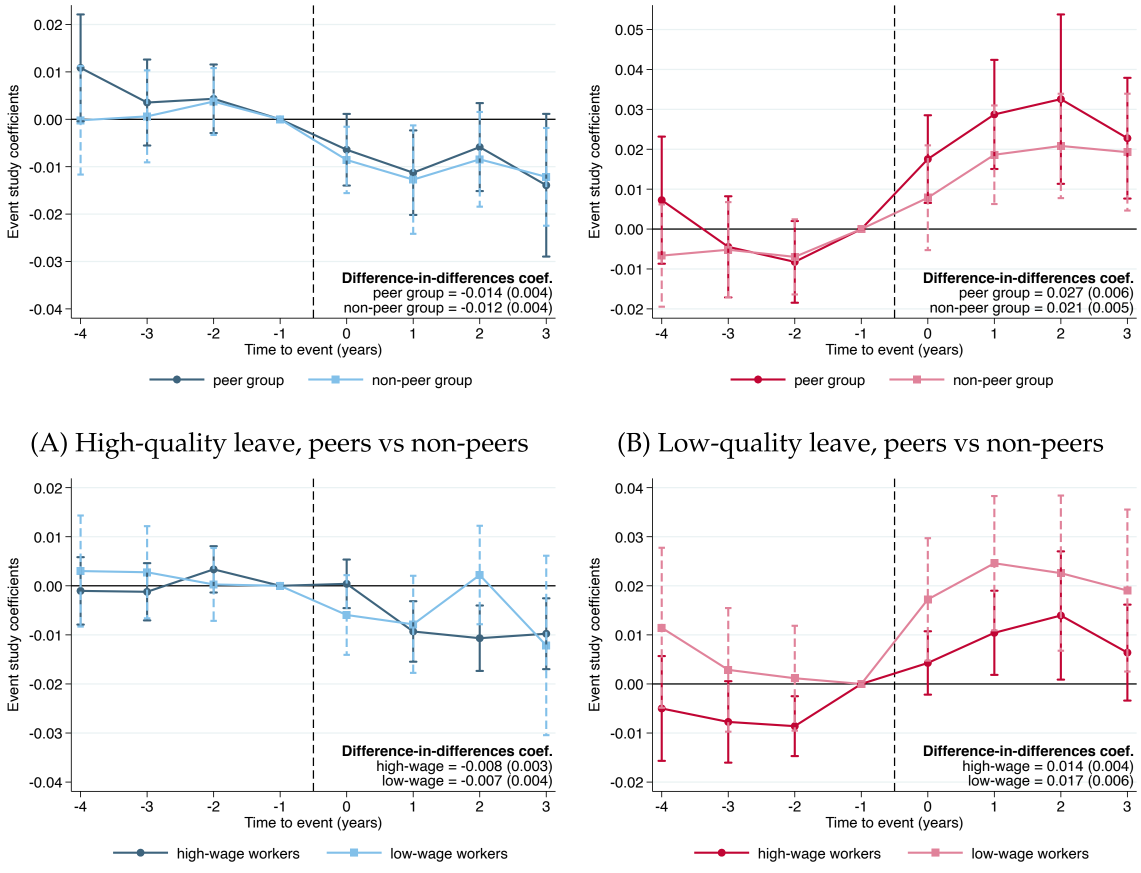

(A) High-quality leave, peers vs non-peers

(B) Low-quality leave, peers vs non-peers

(C) High-quality leave, high- vs low-wage

(D) Low-quality leave, high- vs low-wage

Fig. 2.9 Heterogeneous effects for peers and non-peers (top panels) and high-wage and low-wage coworkers (bottom panels), leave design

Notes. The figure reports estimates of β_k from equation (2.6) for different groups of workers. Panel (A) and (B) compare separations from high- and low-quality workers, respectively, to average-quality separations, distinguishing the effect for workers belonging to the same or different peer group (i.e. to the same occupation). Panel (C) and (D) compare high- and low-quality separations, respectively, to average-quality separations, distinguishing the effect for high- and low-wage workers (i.e. above or below the median wage in the 2-digit sector they belong to). The dependent variable is the log weekly wage of coworkers in the origin firm. Vertical bars are 95% confidence intervals, obtained from cluster-robust standard errors at the firm level.

Worker-level Event Study

We report in Figure 2.10 the estimates from equation (2.7), comparing the wage trajectories of movers into high- and low-quality peers relative to movers into average-quality peers for event period $k = -4, ..., +5$. The event-study analysis suggests that, before the move, there are no significant differences in the wage trajectories between high- and low-quality peers relative to average-quality peers. In the post-event window, estimates indicate that moving into high-quality peers represents an important and substantial driver of wage growth. On average, the weekly wages of such movers increase by 3.9 percent in the post-

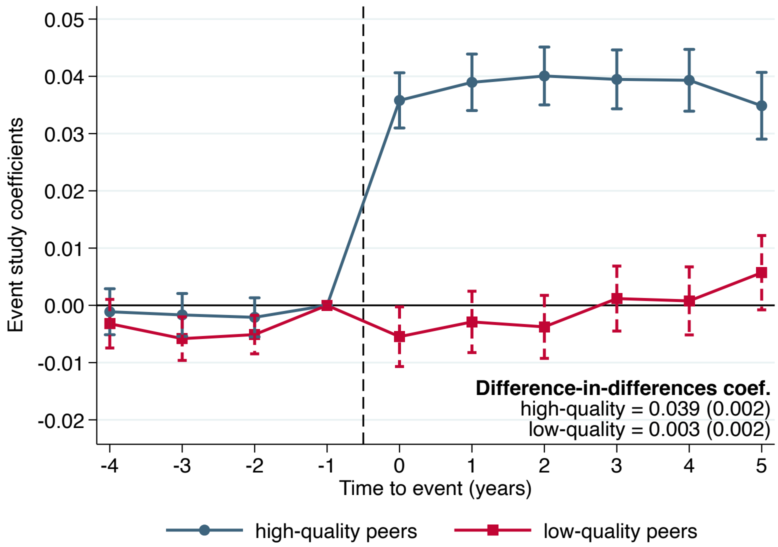

Fig. 2.10 The effect of moving into a high-/low-quality peer group on movers' future wages

Notes. The figure reports estimates of γ_k from equation (2.7), separately for workers moving into high-quality and low-quality peer groups relative to those moving into average-quality peer groups. The dependent variable is the log weekly wage of the mover. Vertical bars are 95% confidence intervals, obtained from cluster-robust standard errors at the individual level.

event period. On the contrary, moving into low quality peers does not affect the trajectory of wages. Moreover, the effect for movers into high-quality peers materializes on impact and then remains approximately constant throughout the five years after mobility. This finding highlights the importance, in terms of wage progression, of working in a high-quality workplace. High-quality peers have a positive impact on workers' performance (by either transmitting knowledge, putting pressure on new entrants or establishing a better network to rely on). Instead, there are no such effects from low-quality peers.

We explore heterogeneous effects by workers' occupation in the year before the move. That is, we distinguish among blue-collar and white-collar movers in Figure 2.11A and 2.11B for movers into high-quality and low-quality peers, respectively. The positive effect of high-quality peers is larger for white-collar workers compared to blue-collar ones: the average difference-in-differences coefficient in the post-event window is 5.7 percent vs 3.6 percent, although we observe an increasing trend in wages in the pre-event window for white-collar workers (but a clear trend break, nonetheless). Interestingly, white-collar workers gain also when moving into low-quality peer groups (Figure 2.11B), as they earn 1.8 percent higher wages on average after mobility relative to white-collar movers into average-quality peers. White-collar workers are probably more skilled,[33] hence they can anyway gain from moving into a new workplace, because of the ability to extract rents or because they can

[33] As anecdotal evidence on this point, the AKM worker fixed effect for white-collar workers is around 20 log points higher than that for blue-collar workers.

boost productivity of the workplace (something that is supported by the increasing trend in the wage gain in the post-event period). The net gain for white-collar workers from moving into high-quality peers relative to low-quality peers is anyway 3.9 percent ($5.7 - 1.8$).

Figure 2.11C and 2.11D report heterogeneous effect for high-wage and low-wage movers (defined, as before, with respect to the median wage in the sector). High-quality peers benefit particularly low-wage workers, for which knowledge transmission from better peers is probably a more important driver of wage growth. On the contrary, they are penalized when moving into low-quality peers. This pattern is consistent with low-wage workers having worse human capital accumulation and outside options which materialize in lower wage growth when moving into low-quality workplaces.

Overall, this analysis helps us understand the mechanisms that identify the peer effect in section 2.3.4. In general, we conclude that changes in peer quality for job stayers determined by the hire of a high-quality workers or the separation from a low-quality worker are the most important in determining wage growth for job stayers, where the positive effects are generally stronger for workers in the same peer group (i.e. in the same occupation) of the new hire or the separated worker. For job switchers, moving into high-quality peer groups contribute to raising the wage level on impact, but moving into low-quality peers does not affect the trajectory of wages.

2.4.3 Robustness Checks

Firm-level shocks One concern with our mobility design is that the entry or leave of a worker may be correlated with firm-level shocks, even after matching on a rich set of observables. On the one hand, the entry of a high-quality peer may be correlated with expectations of future sale growth. On the other, the departure of a high-quality peer may be correlated with expectations of negative shocks to firm sales. We attempt to address this concern by, first, comparing coworkers' wages in firms that hire (or separate from) workers of high- or low-quality relative to firms that hire (or separate from) workers of average-quality and, second, by matching firms on observables before mobility happens. We also inspect the evolution of firms' sales around mobility episodes. Due to data limitations, we only have information on sales (and value added, as a further robustness) for a subset of our data over the period 1996-2001.[34] We re-estimate equation (2.6) in the matched sample using log sales (value added) per worker as a dependent variable and we weight regressions by firm size. Given the sample restrictions and the data limitations we cannot have as many pre-periods as in the main analysis, hence we limit the period before mobility to two years. Figures B.5 and B.6 report the event study estimates for the hire and leaver designs, respectively. Both figures show results for log sales per worker in panel (A) and log value added per worker

[34]Specifically, in the matched sample, we have information on 325 firms in the hire design (13,029 person-year observations) and 338 firms (14,639 person-year observations) in the leaver design.

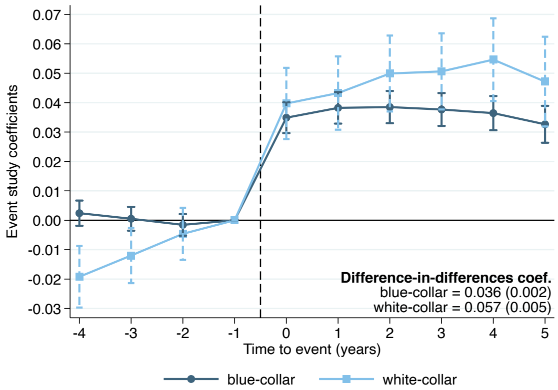

(A) High-quality peers, blue- vs white-c. mover

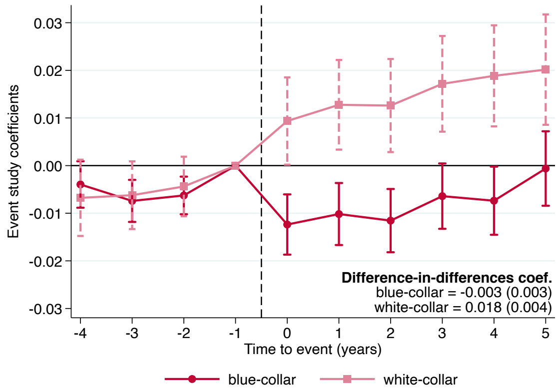

(B) Low-quality peers, blue- vs white-c. mover

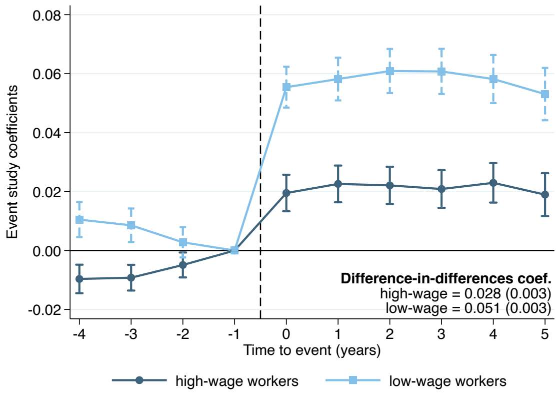

(C) High-quality peers, high- vs low-wage mover

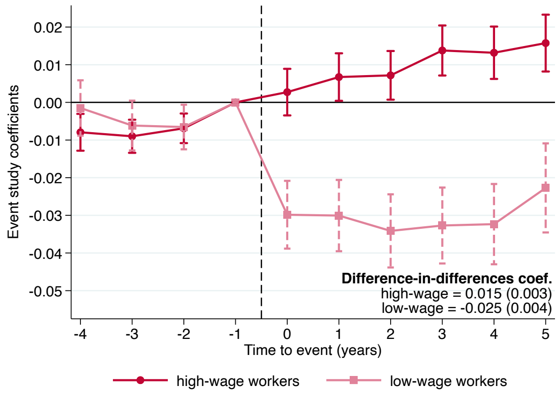

(D) Low-quality peers, high- vs low-wage mover

Fig. 2.11 Heterogeneous effects for blue- and white-collar (top panels) and high-wage and low-wage movers (bottom panels), worker-level design

Notes. The figure reports estimates of γ_k from equation (2.7) for different groups of workers. Panel (A) and (B) show the effects for movers into high- and low-quality peers, respectively, relative to average-quality peers, distinguishing movers that in the year before the move are employed as blue- or white-collar workers. show the effects for movers into high- and low-quality peers, respectively, relative to average-quality peers, distinguishing high- and low-wage movers (i.e. above or below the median wage in the 2-digit sector they belong to). The dependent variable is the log weekly wage of movers. Vertical bars are 95% confidence intervals, obtained from cluster-robust standard errors at the individual level.

in panel (B) and indicate that, in the the subsample for which data is available, there is no significant correlation between the quality of workers who move between firm and firm sales or value added. Moreover, before the move we do not observe different patterns in sale and value added evolution in different groups. This evidence, albeit descriptive and limited to a subset of data, indicates that, after matching, the different groups of firms are comparable in terms of their sale and value added growth. At the same time, this analysis does not really say whether workers' mobility is determined by *expectations* of sale growth, but only that it is not correlated with *realizations* of firm growth. The patterns shown in

Figures B.5 and B.6 are consistent with movers having inaccurate expectations about sales, something that due to data limitations we are unable to capture.

Continuous treatment We replicate our analysis using a continuous treatment, rather than dividing firms into those hiring (separating from) high-, low- or average-quality workers and movers into those joining high-, low- or average-quality peer groups. Specifically, in the coworker-level event study, we replace the dummy $Treat_j$ in equation (2.6) with the AKM worker effect of the worker joining or leaving the firm, i.e. we estimate:

$$w_{-i,j,t} = \eta_t + \psi_j + \sum_{k \neq -1} \widetilde{\beta}_k (\widehat{\alpha}_i \times \mathbf{1}\{t = k\}) + \sum_{k \neq -1} \theta_k (X_j^{pre} \times \mathbf{1}\{t = k\}) + \epsilon_{-i,j,t}, \tag{2.8}$$

where $\widehat{\alpha}_i$ is the pre-estimated worker effect of the joiner or leaver and the other variables are defined as before. We include controls for pre-mobility firm characteristics in X_j^{pre}, interacted with time event dummies: average weekly wages in all pre-periods, the AKM firm effect, the average worker effect, the average age of employees, the share of female and blue-collar employees, firm size and age measured at time -3. $\widetilde{\beta}_k$ measure in this case the dynamic effects on coworkers' wages of a one percent increase in the joiner/leaver quality.

In the worker-level event study, we replace $Treat_i$ in (2.7) with the average AKM peer effect of coworkers in the destination firm. In other terms, we estimate:

$$w_{i,t} = \eta_t + \alpha_i + \sum_{k \neq -1} \widetilde{\gamma}_k (\overline{\widehat{\alpha}}_{-i} \times \mathbf{1}\{t = k\}) + \sum_{k \neq -1} \kappa_k (X_i^{pre} \times \mathbf{1}\{t = k\}) + \epsilon_{i,t}, \tag{2.9}$$

where $\overline{\widehat{\alpha}}_{-i}$ is the average pre-estimated worker effect of coworkers in the destination firms and the other variables are defined as before. We include controls for pre-mobility worker characteristics in X_i^{pre}, interacted with time event dummies: weekly wage, number of weeks worked, decile of worker fixed effect distribution, age, gender, tenure, occupation, AKM firm effect and firm size, all measured at time -3. $\widetilde{\gamma}_k$ measure the dynamic effects on the mover's wage of a one percent increase in the peer quality of the destination firm.

Results are reported in Figure B.7, which shows a very similar pattern to that reported in the main analysis. A 10 percent increase in the quality of a new hire increases coworkers' wage by 0.9 percent on average in the period after mobility, whereas a 10 percent increase in the quality of separating workers does not affect coworkers' wages in the origin firm, although the dynamic effect shows some negative adjustment in the year following mobility. For a mover, a 10 percent increase in the quality of peers at the destination firm increases wages by 2 percent on average. Overall, these results confirm the findings from the event study with discrete treatment groups.

2.5 Conclusion

This paper explores a critical driver of wage growth: coworker's quality. We find that the quality of coworkers plays an important role in increasing future wages. Using an innovative estimation method in the classical AKM framework, we show that a 10 percent increase in coworker's quality raises the next year's wage by 1.8 percent, which is almost equivalent to the size of the return to working with college educated peers (Nix, 2020). The peer effect gradually decreases in importance over time, but we find that after five years a 10 percent increase in past peer quality still determines 0.5 percent higher wages. When exploring heterogeneous effects, we find that the peer effect is larger for movers, workers with low tenure and young workers, whereas we do not find differential effects for different firm sizes. The peer effect is likely due to a combination of knowledge spillover, peer pressure and better network opportunities. The finding that after five years a peer effect is still present and that it is stronger for movers let use believe that knowledge acquired from past peers plays an important role in shaping future wages.

Furthermore, the peer effect is identified, for job stayers, by changes in peer quality when a worker enters or leaves the firm and, for job switchers, by moving into a new peer group. We separately analyse each of these channels in the second part of the paper, by setting up an event study analysis around mobility of workers. After carefully matching firms in the pre-mobility period, we find that if a firm hires a high-quality worker, coworkers' wages increase by 1.8 percent relative to coworkers' wages in firms hiring an average-quality worker. We do not find, instead, significant effects in firms hiring a low-quality worker. The opposite effect is found in firms separating from a worker, as we find a positive effect for coworkers' wages in firms separating from low-quality workers and a moderate negative effect for coworkers' wages in firms separating from high-quality workers.

We also explore the wage trajectories of workers who move into peer groups of different quality and find that moving into high-quality peers is an important driver fo wage growth, with wages being on average 3.9 percent higher than those of workers moving into average-quality peer groups. On the contrary, the wage trajectory is unaffected for workers moving into low quality peers.

Overall, our findings suggest that hiring high-quality workers, separating from low-quality workers and moving into high-quality peers generates higher wages than counterfactual scenarios where peer quality does not change. This indicates that knowledge spillover, but also pressure from better peers, are both mechanisms that may help explain why peers matter in the workplace for wage growth.

Future research should focus on opening the black box of the mechanisms behind the contribution of peers to workers' wage growth. A separate analysis of the knowledge spillover effect from the peer pressure and the network effect would help us in understanding how interaction in the workplace benefits workers. Increased availability of administrative data, coupled with either structural or reduce-form models, or laboratory experiments, will help reach a definitive answer to this question.

Chapter 3

Sorting Robots: How Automation Shapes the Allocation of Workers Across Firms

Abstract This chapter studies the worker-, firm- and sector-level adjustment to robots. Combining detailed matched employer-employee data for Italy over the period 1994-2018 with robot counts by industry in the manufacturing sector, we show that automation adoption expands employment opportunities and reduces labour market transitions. At the worker level, those who are either high-skilled, white-collar, or employed in more productive firms experience employment and earnings gains. Meanwhile at the firm-level, sales and value added increase, while employment outcomes are highly heterogeneous between ex-ante more and less productive firms; with the former increasing employment of all workers, irrespective of their skill level, and the latter reducing it. These changes in labour force composition are further inspected at the sector-level, where an event study approach following spikes in automation adoption reveals a negative effect of automation on labour market sorting.

Keywords: Automation, Employment, Wages, Matched Employer-Employee Data

***JEL* codes:** J23, J30, O33

3.1 Introduction

The impact of automation on the labour market is one of the more widely discussed topics in the academic and policy debate. The fear is that automation can have disruptive consequences on the organization of labour, leading to a surge in unemployment of workers performing tasks that could be replaced by robots. Theoretical work on this subject (Acemoglu and Restrepo, 2019) highlights that there are two main economic effects of automation on the labour market: a displacement effect and a productivity effect. The former has a negative effect on employment and wages as workers replaced by machines are displaced to other firms or to unemployment. The latter has a positive effect as firms become more productive and expand output and, therefore, employment and wages. Which effect prevails is ultimately an empirical question.

In this paper, we contribute to the large and growing strand of the literature that attempts at giving answers to this question, by showing the worker-, firm- and sector-level adjustment of labour markets to robots. We focus on industrial robots, defined as "an automatically controlled, reprogrammable, multipurpose manipulator" (IFR, 2020), which are mainly deployed in the manufacturing sector. The surge of industrial robots in production processes is evident in all economies, as the number of operational robots increased by a factor of 4.9 in the world as a whole.

We focus on Italy as a case study. Italy is an interesting country to analyse the impact of automation for two main reasons. First, it is the second largest robot market in Europe and the sixth worldwide. As shown in Figure 3.1, Italy experienced a growth in the number of industrial robots per 1000 workers that is larger than the average of other European countries and second to Germany, only. The surge in industrial robots in Italy is related to the weight of the manufacturing sector – and of the automotive industry, in particular – in its economy. Second, Italy has a relatively rigid labour market and high employment protection legislation. The automation effects may be considerably different in countries with different labour market institutions.[1]

We explore the impact of automation on the labour market along three lines. First, we investigate the worker-level adjustment to robots. Specifically, we follow Dauth et al. (2021) and, using detailed administrative linked employer-employee data, we compare the earnings, employment and wage trajectories of workers with similar observable characteristics that are initially employed in sectors with different automation exposure over the period 1994-2018. We deal with endogeneity in the relationship between labour market outcomes and automation by exploiting a standard instrumental variable strategy. We use the surge in robot exposure across sectors in other countries as an instrument for the national surge in

[1]For example, the evidence has so far found negative effects of automation on employment in the United States Acemoglu and Restrepo (2019) and positive effects in European countries Dauth et al. (2018); Graetz and Michaels (2018). These differences may partly be due to differences in labour market institutions.

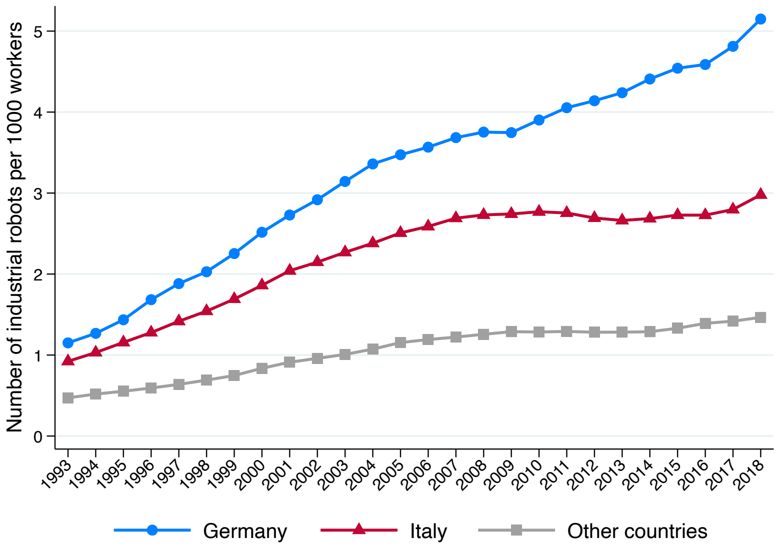

Fig. 3.1 Number of industrial robots per 1000 workers in selected countries

Notes. The figure shows the number of robots per 1000 workers exploiting data from the International Federation of Robotics on the stock of operational robots for each country and year. *Other countries* is an average of the stock of robots in Finland, France, Norway, Spain, Sweden and United Kingdom.

robots, building upon the pioneering work of Autor et al. (2014). This analysis reveals that, as in Germany but differently from the United States (Acemoglu and Restrepo, 2019), the impact of automation on employment is positive, with more exposed workers increasing their labour supply and experiencing lower rates of job churning, while we find no significant effects on earnings or wages. These effects, however, are highly heterogeneous and configure the presence of "winners" and "losers" from automation. Specifically, incumbent workers experience increased employment and earnings levels, whereas workers displaced to other firms – in manufacturing or outside of it – experience worse labour market outcomes. Moreover, worker and firm types matter for the magnitude of the adjustment: blue-collar and low-skill workers experience significant penalties, as well as workers initially employed in less productive firms, that is, firms in the bottom tercile of the value added per worker distribution.[2] These results are confirmed if, instead of looking at employment as the main outcome, we use the number of days spent in unemployment or short-time work, for which we find results of opposite sign.

Second, we investigate the firm-level adjustment to robots, by using a subsample of the data with balance-sheet information over the period 1994-2012. We focus on firms to evaluate how robots affect firm-level employment, wages and performance and to examine, through different lens, how automation impacts reallocation of workers with different skill levels

[2] Less productive firms are also less likely to automate (Koch et al., 2021), which may help explain why employees in such firms do not benefit from automation.

across firms with different productivity levels. Using a similar strategy to the one adopted for the worker-level analysis, we find that firms in sectors with higher automation exposure expand employment and labour costs and witness significant increases in value added and sales. We then investigate how firms change their labour force composition and we find that firms, in sectors with higher automation exposure, with low productivity levels decrease employment of high-skilled workers who move to more productive firms, contributing to higher assortative matching in the labour market. At the same time, more productive firms also increase employment of low- and middle-skilled workers, highlighting how the productivity effect of automation benefits all workers, but only in a limited subsample of firms with high productivity levels, contributing to reducing assortativeness in the labour market. The overall effect on labour market sorting is, therefore, ambiguous.

Finally, to further inspect the relationship between robots and sorting, we investigate how automation changes the allocation of workers across firms by adopting an event study strategy around automation spikes (in the spirit of Bessen et al., 2019), that is, around events in which the number of industrial robots increases in a sector by a high margin relative to other sectors. We investigate how such "automation shocks" translate into labour market outcomes and specifically in how matching and sorting changes within each sector. We measure sorting as the correlation between AKM (Abowd et al., 1999) worker effects and measures of firm types, such as, log value added per worker and log sales per worker (as in Bombardini et al., 2019).[3] Overall, we find that there is a short-run negative and significant effect of automation on labour market sorting and a medium-run insignificant result. This implies that, at least under this point of view, automation does not contribute to between firm earnings inequality, determined by high-wage workers being employed by more productive firms (Card et al., 2013; Song et al., 2018), but it contributes to between firm inequality as there is an increasing dichotomy between low-productive and high-productive firms and to between workers inequality in light of the differential effects for different worker types.

This paper contributes to a rapidly expanding literature on the effects of automation on the labour market, which shows that the effects of automation on employment and wages are highly heterogeneous and crucially depend on the context in which they materialize. Acemoglu and Restrepo (2019) report large negative employment and wage effects across US local labour markets. For Europe, the evidence is far less pessimistic. The pioneering study by Graetz and Michaels (2018) finds positive productivity effects and no employment losses. Similar findings are also reported in Klenert et al. (2020), who additionally do not find displacement effects for low-skilled workers. Dauth et al. (2021) and Dottori (2020) are the most closely related papers to ours. The first one provides worker-level evidence on the

[3] An alternative is to measure sorting as the correlation between AKM worker and firm effects. However, as highlighted in Bonhomme et al. (2019); Eeckhout and Kircher (2011), firm effects can be poor proxies of true firm types and the correlation with worker effects could be severely biased due to low mobility of workers across firms (Andrews et al., 2008). To overcome these problems we choose to use more reliable measures of firm types, exploiting balance sheet administrative data on firms.

effects of robots for Germany, a country that has many points in common to Italy in terms of the structure of its labour market, which heavily depends on its manufacturing sector, and on labour market institutions, such as wage bargaining and employment protection legislation. The effects of automation are found to be positive on German employees, especially on incumbent workers. The second one replicates these analyses for Italy. The conclusion of Dottori (2020) are widely similar to those of Dauth et al. (2021) for Germany and are in line with our results. This paper considerably extends on these two works, by providing additional evidence on the type of firms where displaced workers end up being employed and on the effects on unemployment, short-time work take-up and sick leaves. Moreover, we contribute by providing evidence on the firm- and sector-level effects of automation, exploiting rich data on balance sheets of firms, showing how the effects of automation crucially depend on ex-ante characteristics of firms and workers and on how labour force composition changes.

One limitation of the present analysis (and of the strand of the literature that uses the same data source for robots adoption) is that it exploits automation variation at the sector-level. Following the examples of Autor et al. (2014) and Dauth et al. (2021), we analyse how different types of workers and firms within sectors are affected and, exploiting rich administrative matched employer-employee data, which allow to follow workers' careers over time, we are able to track the contribution of different reallocation mechanisms to the overall effect of automation. However, there is also a growing set of contributions in the field that attempt to measure firm-level adoption of automation (Aghion et al., 2020; Bessen et al., 2019; Bonfiglioli et al., 2020; Dixon et al., 2020) and how this translates into firm-level outcomes, such as productivity and employment. The findings are contrasting and generally highlight the presence of selection of firms into adopting automated technologies. For this reason, we control for a rich set of firm and industry-level controls in order to alleviate the selection effect.

The rest of the paper is organized as follows. Section 3.2 describes the data. Section 3.3 analyses the worker-level adjustment to automation. Section 3.4 reports the firm-level results on the effects of robots. Section 3.5 discusses the implications for labour market sorting, taking a sector-level perspective. Finally, Section 3.6 concludes.

3.2 Data

There are three main data sources. On the worker-firm side, we use LoSaI (Longitudinal Sample Inps) records, a matched employer-employee dataset that contains a random sample of the universe of workers in the Italian non-agricultural private sector. The data covers approximately 7% of the universe of employees over the period 1985-2018. The dataset comprises different archives. In the worker archive, we observe the entire working history of each individual and, specifically, we have information on annual gross earnings,[4] the number of weeks and days worked in a year, the type of contract (full-time or part-time and permanent or temporary) and broad occupation categories (apprentice, blue-collar, white-collar, middle manager, executive). In a separate archive, the dataset contains demographic information on each employee, such as year of birth, gender and region of residence. Finally, the firm archive records total firm size in discrete brackets and firm's sector.

Using firm identifiers, we match firm records with balance sheet information from Cerved, a dataset covering the period 1994-2012. Specifically, we retain information on firms' sales, value added, total assets, debts and profits. Moreover, over the same period, the administrative data records total employment and its breakdown into employment of blue- and white-collar workers. Additionally, we have information on total labour costs, distinguishing again between blue- and white-collar workers.

On the automation side, data on robot adoption is provided by the International Federation of Robotics, which records the stock of industrial robots per sector in a number of different countries over the period 1993-2018. For the purpose of this paper, we will focus on the stock of industrial robots per sector in Italy as well as seven European countries (Germany, France, Spain, Sweden, Denmark, Finland and the UK). The sectors included in the IFR data mostly comprise manufacturing and are less detailed for non-manufacturing sectors. Therefore, we will focus mostly on manufacturing, but will also include agriculture, forestry and fishing, mining and quarrying, electricity, gas and water supply. Table C.1 in the appendix details the sectors included in the final sample. Given that LoSaI data contains sector information at the 2-digit level, we group sectors into 19 groups (as shown in the third column of Table C.1), corresponding, approximately, to 2-digit sectors. We will retain information on services to capture potential re-allocation of workers outside of manufacturing.

3.3 Worker-level Adjustment

We begin by investigating the effects of automation on individual workers, by selecting a cohort of employees and following their cumulative outcomes over the period 1993-2018.

3.3.1 Empirical Strategy

Regression model We closely follow Autor et al. (2014) and Dauth et al. (2021) and focus on the labour market outcomes of workers employed in manufacturing in 1994. We focus specifically on workers with high-labour force attachment, described shortly, and present results for all workers for robustness. We select workers in the data that, in 1994, are between

[4]The measure of earnings is gross of labour income taxes and pension contributions on the side of the employee.

21 and 39 years old, have at least 2 years of tenure with their firm and have worked at least 26 weeks on average per year in the period 1990-1994. We then retain only workers employed in a firm in manufacturing in 1994. On this sample of workers, we run the following specification:

$$Y_{ij} = \alpha + \beta \cdot \widehat{\triangle \text{robots}}_j + \gamma \cdot X'_{ij} + \delta \cdot Z'_j + \varepsilon_{ij} \tag{3.1}$$

Y_{ij} is a cumulative outcome (earnings, days worked, or average wages) over the period 1994-2018 for worker i initially employed in sector j, with $j = \{1, ..., 19\}$. $\widehat{\triangle \text{robots}}_j$ is the sector-level change in robot exposure between 1994 and 2018, instrumented with the corresponding change in seven other European countries (Finland, France, Germany, Norway, Spain, Sweden, United Kingdom). X_{ij} controls worker-level observables: birth cohort dummies, a dummy for women, region of residence dummies, tenure dummies (0-3 years, 4-6 years, 7 years or more), terciles of conditional wages,[5] and firm size dummies (brackets of workers: 1-10, 11-20, 21-50, 51-100, 101-500, 500 or more employees). Z_j contains sector-level controls: dummies for broad manufacturing sectors (food, beverages and tobacco; consumer goods, industrial goods and capital goods),[6] the change in net exports from China and Eastern Europe[7] and the change in ICT exposure at the sector-level.[8] The coefficient of interest is β which measures the effects of one additional robot per 1000 workers between 1994 and 2018 on the outcome. This empirical strategy allows to compare outcomes of workers with similar observable characteristics in the baseline year across sectors within broad industries that experienced different changes in automation penetration and crucially controls for other aggregate trends in technology adoption (via the inclusion of the change in ICT) and globalization (via the change in trade exposure).

Instrument validity Our empirical strategy rests on the assumption that we are able to identify exogenous changes in robot adoption. To do so, we use a standard instrumental variable strategy, by predicting robot exposure based on the change in robot penetration in other high-income countries. The validity of the instrument requires the usual conditions of relevance and exogeneity. As to the first one, we provide extensive evidence that robot penetration in high-income countries is highly correlated with robot penetration in Italy. The

[5]Conditional wages are computed as the residual of a regression in 1994 of full-time equivalent daily wages on a dummy for female workers, a part-time dummy, quadratic polynomials in age and tenure, 20 dummies for region of residence, dummies for firm size and sector dummies.

[6]These dummies group together NACE Rev. 2 sectors 10 to 12 (food, beverages and tobacco); 13 to 18 and 30 to 33 (consumer goods); 19 to 25 (industrial goods); 26 to 29 (capital goods).

[7]We use data on trade volumes by country and sector from the OECD to compute the change in total Italian net exports (exports minus imports) over the period 1994-2018. We use exports to and imports from China, Azerbaijan, Belarus, Bulgaria, Czech Republic, Estonia, Georgia, Hungary, Kazakhstan, Kyrgyzstan, Latvia, Lithuania, Moldova, Poland, Romania, Russia, Slovak Republic, Slovenia, Tajikistan, Turkmenistan, Ukraine, Uzbekistan.

[8]We use data from the EU-KLEMS database on gross fixed capital formation in computing and communications equipment. Because of data availability, we compute the change over the period 1996-2016.

second one requires that robot adoption in other high-income countries captures international technology advancement in each industry that are uncorrelated with labour demand shocks in domestic industries. This requirement fails if, for example, labour demand shocks are correlated across countries, so that instrumental variables estimates would be biased by correlation between robot adoption and unobserved components of labour demand. One example of such correlation could be migration of high-skill workers in countries with higher technological potential that boosts robot adoption through greater complementarity with worker types. Additional concerns about instrument exogeneity are related to the possibility that robot adoption may happen as a consequence of other labour demand shocks that shift employment and wages and therefore technology investment decisions by firm. Examples of such shocks are growth in imports from China and Eastern Europe or the growth in computerization. Although there is no easy solution to properly account for such potential biases, we report estimates controlling for a rich set of baseline covariates in each industry, among which we include changes in import from China and Eastern Europe and in ICT equipment.

Descriptive statistics Table 3.1 reports descriptive statistics for the main sample of workers with high-labour force attachment (as described before) in columns 1 and 2 and the full sample of workers in columns 3 and 4. Panel A reports the cumulated outcomes over the period 1994-2018. We measure cumulative earnings over the 25 years following 1994 and we normalise them with respect to base year earnings. Hence, in the absence of earnings growth, the average worker would have 25-fold his or her earnings in 1994. This is indeed the case for workers in the main sample (column 1). Workers in the full sample experience slightly larger earnings growth. The cumulative number of days worked by the average worker is 5,444 in column 1 and 5,218 in column 3, i.e. approximately 83% and 80% of the number of working days available. The number of labour market transitions is larger for workers with lower labour force attachment. Panel B reports descriptive statistics about the worker-, firm- and sector-level controls included in the regressions. Panel C shows that the average worker experiences a change in automation exposure of approximately 8.8 additional robots per 1000 workers, with a large variation between workers at the 10th or 25th percentile of the robot exposure distribution, who experience an almost null change and workers at the 90th or 75th percentile who experience a change of 20 and 15 additional robots per 1000 workers. The number of workers in the main sample is 136,736.

3.3.2 Results

First stage We begin by investigating the strength of the first-stage relationship between the endogenous change in robot exposure across Italian sectors and the change in robot exposure across European sectors. Table C.2 in the Appendix reports the estimates under different

Table 3.1 Summary statistics, worker-level sample

	(1)	(2)	(3)	(4)
	Main sample		Extended sample	
	mean	(sd)	mean	(sd)
Panel A: Outcomes, cumulated values over 1994-2018				
Earnings / Base year earnings	24.62	(14.81)	29.46	(55.91)
Days	5,444.46	(2396.03)	5,218.32	(2497.85)
Average daily wage / Base year wage	1.20	(0.34)	1.20	(0.36)
Transitions	2.30	(2.24)	2.58	(2.45)
Panel B: Control variables				
Female	0.33	(0.47)	0.33	(0.47)
Age	30.15	(5.10)	29.74	(5.20)
Avg. earnings, 1990-94	22,163.28	(9438.31)	20,257.59	(10197.96)
Avg. daily wage, 1990-94	79.46	(28.05)	78.29	(28.31)
North-West	0.39	(0.49)	0.38	(0.49)
North-East	0.28	(0.45)	0.28	(0.45)
Centre	0.17	(0.38)	0.17	(0.38)
South	0.16	(0.37)	0.17	(0.38)
Tenure $\leq$ 3 years	0.25	(0.43)	0.40	(0.49)
Tenure 4-6 years	0.34	(0.47)	0.28	(0.45)
Tenure $\geq$ 7 years	0.41	(0.49)	0.33	(0.47)
Firm size $\leq$ 10	0.22	(0.42)	0.24	(0.43)
Firm size 11-20	0.15	(0.35)	0.15	(0.35)
Firm size 21-50	0.17	(0.37)	0.16	(0.37)
Firm size 51-100	0.11	(0.31)	0.10	(0.30)
Firm size 101-500	0.19	(0.39)	0.18	(0.38)
Firm size $>$ 500	0.17	(0.37)	0.16	(0.37)
Food products, beverages and tobacco	0.07	(0.25)	0.08	(0.27)
Consumer products	0.38	(0.49)	0.38	(0.49)
Industrial goods	0.30	(0.46)	0.30	(0.46)
Capital goods	0.18	(0.39)	0.18	(0.38)
$\triangle$ net exports / wagebill in %	-20.18	(120.22)	-19.90	(119.86)
$\triangle$ ICT equipment in Euro per worker	123.49	(184.12)	123.43	(181.77)
Panel C: Robot exposure				
$\triangle$ robots per 1000 worker	8.75	(13.57)	8.81	(13.46)
p10-p90 interval	[0.13-19.86]		[0.13-19.86]	
p25-p75 interval	[0.33-14.89]		[0.33-14.89]	
Sample size	136,736		174,405	

Notes. The table reports summary statistics from the high labour force attachment worker sample used in the main analysis (columns 1 and 2) and the full sample of workers (columns 3 and 4). Panel A reports means and standard deviations of outcomes, cumulated over the period 1994-2018. Panel B report control variables in the baseline year 1994. Panel C shows means, standard deviations and quantiles of the distribution of robot exposure, measured as the change in robots per worker by sector between 1994 and 2018.

specifications. In all specifications, the first-stage relationship is strong and statistically significant with an F-statistic equal to 114 in column (6), which controls for worker, firm and sector characteristics.

Overall impact of automation on workers Table 3.2 reports the estimates of equation (3.1) using cumulative earnings over 1994-2018 under different model specifications. Panel A reports the main sample that will be used throughout the rest of the worker-level analysis, i.e. the high labour force attachment group. Panel B reports the estimates for the full sample of workers. For both samples, the table also reports the p-values from the Hansen-Sargan test, which provides evidence against the rejection of the null of valid instruments. Column (1) reports OLS estimates of β from equation (3.1), i.e. the effect of one additional robot per worker on cumulative earnings in percent of baseline earnings in 1994. The only additional covariates included in this specification are birth year dummies (from 1955 to 1972). The effect is positive but not statistically significant in panel A and is higher in magnitude and statistically significant in panel B. Column (2) reports 2SLS estimates of the same specification. Comparing OLS and 2SLS coefficients suggests a slightly negative bias in OLS. The 2SLS estimate, moreover, is statistically significant in both panels A and B. Additional demographic controls (gender and region of residence) are included in column (3). The inclusion of such controls considerably reduces the magnitude of the estimate, which becomes insignificant in panel A and only marginally significant in panel B. The inclusion of further employment controls (tenure dummies, log earnings in 1994 and terciles of adjusted wages) in column (4), of industry dummies and firm size dummies in column (5), and of the change in net exports and ICT capital in column (6) further reduces the magnitude of the coefficient, which becomes statistically insignificant in both panels. For the remainder of the analysis, we will use the specification of column (6) and restrict the focus to the sample used in panel A, only.

Table 3.3 reports the estimates of equation (3.1) for cumulative earnings (which replicates column (6) of Table 3.2), cumulative days worked, average daily wages (computed as the ratio of cumulative earnings and cumulative days worked) and cumulative labour market transitions (to other firms and to non-employment). Panel A reports the estimates for the treatment period 1994-2018. The impact of robots is positive and statistically significant on cumulative days worked: the coefficient suggests that moving from the 25th to the 75th percentile of the change in automation exposure translates into additional 111 ($= 5.631 \times [19.86 - 0.13]$) days worked. At the same time, column (3) suggests there is no significant effect on average daily wages, but there is a significantly lower job churning rate (column (4)): workers initially employed in more exposed industries experience a lower number of labour market transitions. Specifically, comparing again workers at the 25th and 75th percentile of automation adoption, the latter experiences 0.11 less transitions ($= -0.572/100 \times [19.86 - 0.13]$).

Table 3.2 Robot exposure and cumulative earnings, 1993-2018: OLS and 2SLS estimates

	(1) OLS	(2)	(3)	(4) 2SLS	(5)	(6)
Panel A: High labour force attachment group						
$\triangle$ robots per 1000 worker	6.129	7.669*	3.963	3.318	2.044	2.302
	(4.031)	(4.275)	(2.649)	(2.333)	(1.938)	(1.679)
Hansen-Sargan test		0.756	0.627	0.731	0.398	0.819
Kleibergen-Paap F-statistic		140.5	135.2	135.7	138.1	114.3
Observations	136,736	136,736	136,736	136,736	136,736	136,736
Panel B: All workers						
$\triangle$ robots per 1000 worker	13.038**	14.874***	10.415***	10.399*	4.341	3.672
	(5.113)	(5.519)	(3.232)	(5.306)	(4.640)	(5.020)
Hansen-Sargan test		0.754	0.757	0.386	0.349	0.356
Kleibergen-Paap F-statistic		128.6	121.0	121.7	127.9	111.1
Observations	174,405	174,405	174,405	174,405	174,405	174,405
Birth year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls			Yes	Yes	Yes	Yes
Employment history				Yes	Yes	Yes
Industry and firm controls					Yes	Yes
$\triangle$ Net exports and $\triangle$ ICT						Yes

Notes. The table reports coefficients from the estimation of equation (3.1). Column (1) reports OLS estimates. Columns (2) to (6) report 2SLS estimates. Columns (1) and (2) include birth year dummies. Column (3) controls for workers' gender and region of residence (20 dummies). Column (4) includes tenure dummies (4-6 years, ≥ 7 years), log earnings in 1994 and terciles of adjusted wages. Column (5) includes five firm size dummies (11-20, 21-50, 51-100, 101-500, 500 or more employees) and four dummies for broad manufacturing sectors (food, beverages and tobacco; consumer goods, industrial goods and capital goods). Column (6) further controls for the change in net exports form China and Eastern Europe over 1994-2018 and the change in ICT capital over 1996-2016. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.3 Robot exposure and cumulative earnings, cumulative days, years with positive earnings, earnings per year and per day, 1993-2018 and 1990-1993: 2SLS estimates

	(1) Cumulative earnings	(2) Cumulative days	(3) Average daily wage	(4) Cumulative transitions
Panel A: treatment period 1994-2018				
△ robots per 1000 worker	2.302	5.631***	0.008	-0.572***
	(1.679)	(1.264)	(0.065)	(0.162)
Panel B: pre-period 1990-1994				
△ robots per 1000 worker	0.113	0.298**	0.006	-0.092**
	(0.071)	(0.129)	(0.006)	(0.044)

Notes. $N = 136{,}736$. The table reports estimates of β from equation (3.1). Outcomes are cumulated values over 1994-2018 in panel A and over 1990-1994 in panel B. Columns (1) to (4) report the estimates for cumulative earnings $\times 100$, cumulative days, average daily wage $\times$ 100 (computed as the ratio of cumulative earnings over cumulative days), cumulative labour market transitions $\times 100$. All regressions include the full set of controls from column (6) of Table 3.2. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

One concern about the estimates is the potential selection of workers into industries with more resources to automate production processes, even after controlling for a rich set of worker-level and firm-level observables. We test this hypothesis by running the same regressions on outcomes cumulated over the period 1990-1994. Results are reported in panel B. They suggest some selection for cumulative days and labour market transitions (columns (2) and (4), respectively). In other terms, workers in more exposed industries were increasing their labour supply and were subject to lower job churning also before the treatment period. However, Figure 3.2 reassures against the magnitude of such "pre-trends". The figure plots dynamic effects for cumulative earnings, cumulative days and cumulative labour market transitions. Specifically, the figure plots the coefficient of a regression of cumulative outcomes up to the year reported on the x-axis, on the change in automation exposure between 1994 and 2018, so to evaluate how the overall effect amasses over time.[9] The figure shows that pre-trends in cumulative days and transitions are generally very small and flat, whereas only starting from 1994 the figure shows a clear change in the slope of the dynamic effects. Moreover, the figure also suggests that the effect for earnings is positive and statistically significant before 2011, but the estimates become less precise (although coefficients become larger) in more recent years.

Decomposition The overall effect of automation exposure on workers is the sum of the effects for incumbent workers and job changers. The latter can be further decomposed into the effect for movers to another firm within the same 2-digit industry in manufacturing,

[9] For years prior to 1994, we cumulate outcomes between 1990 and the year reported on the x-axis.

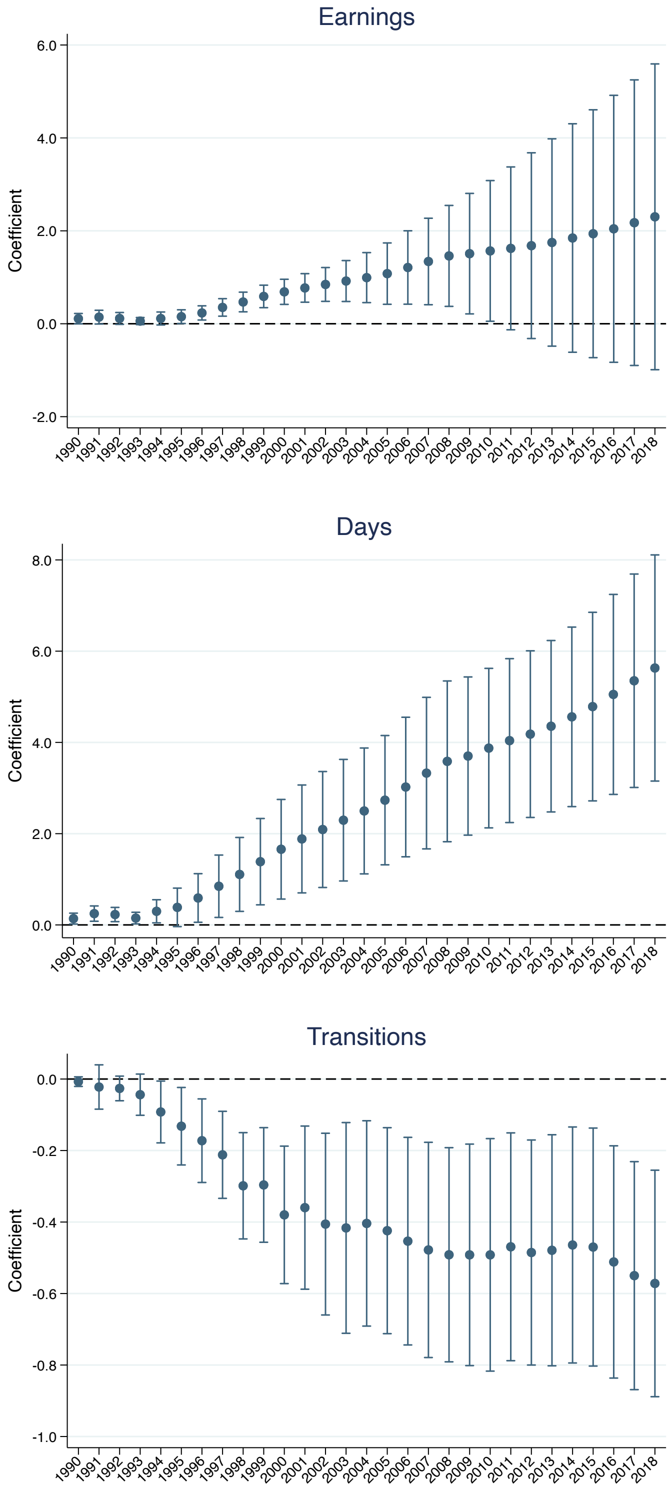

Fig. 3.2 Cumulative earnings, days and labour market transitions since 1990

Notes. $N = 136{,}736$. The figure reports dynamic estimates of β from equation (3.1). Outcomes are cumulated values between 1994 and the year reported on the x-axis for years after 1994, and between 1990 and the year on the x-axis for years before 1994. The top figure reports the estimates for cumulative earnings $\times 100$. The central figure reports the estimates for cumulative days and the bottom figure reports the estimates for cumulative labour market transitions $\times 100$. All regressions include the full set of controls from column (6) of Table 3.2. Vertical bars are 95% confidence intervals, from cluster-robust standard errors by 36 two-digit NACE Rev.2 industries.

movers to other 2-digit industries in manufacturing, and movers in the non-manufacturing sector. Table 3.4 reports such decomposition. Panels A, B and C use cumulative earnings, cumulative days and average daily wages as outcomes, respectively. Column (1) shows the overall effect, replicating results from Table 3.3. Column (2) shows the effect for workers staying at their original firm, using as dependent variables only cumulated outcomes for workers remaining with their original employer. The table suggests that such workers experience significant gains in terms of both earnings and days worked. Comparing the implied effect at different points in the distribution of automation exposure, incumbent workers at th 75th percentile experience close to 150% ($7.390 \times [19.86 - 0.13] = 145.8$) larger cumulative earnings growth and approximately one additional year of work ($16.049 \times [19.86 - 0.13] = 316.6$) than a worker at the 25th percentile. The effect on average daily wages is positive but not statistically significant. The positive effect for incumbent workers is compensated by negative effects for workers changing employer: column (3) report the effects for workers moving to another firm within the same 2-digit industry in manufacturing. Such workers experience both income and days losses, but only the latter are significant. This evidence suggests the presence of displacement effects, as theorized by Acemoglu and Restrepo (2019), that produce earnings losses for workers losing their job following the rise in automation exposure, and it configures the presence of "winners" and "losers" from technology adoption that the aggregate effect partially hides. Column (4) of the table reports the effect for workers moving to another employer and 2-digit industry within manufacturing, whereas column (5) reports the estimates for workers moving to non-manufacturing. The coefficients are negative for all outcomes in both specifications, but statistically significant only in the latter. In particular, workers displaced to the non-manufacturing sector experience earnings, days of work and wage losses of considerable magnitude, that are almost as large as the positive overall effect for all workers. This evidence is similar to the one provided for Germany by Dauth et al. (2021), but the magnitude of the effects is larger in Italy.

Table C.3 in the Appendix further decomposes the effects for workers changing employers into the effects for workers moving to larger or smaller firms than the firm of origin. Columns (4) and (5) show that the negative effect for workers changing employer within industry are statistically significant for workers moving to smaller firms only, but the estimates are very imprecise and it is impossible to draw conclusions on differences between different sets of coefficients. Similar differences between mobility to larger versus smaller firms are observed for workers moving into the non-manufacturing sector, too (columns 10-11).

Heterogeneous effects by worker and firm characteristics The results presented so far have already highlighted how the effects of automation on the labour market are highly heterogeneous for incumbent workers relative to job changers, with the former benefiting from increased labour supply and earnings and the latter experiencing worse labour market

Table 3.4 Exposure to robots and earnings and employment by industry and firm, 1993-2018: 2SLS estimates

	(1) all employers	(2)	(3) manufacturing	(4)	(5) non-manuf.
Same 2-digit industry		Yes	Yes	No	No
Same firm		Yes	No	No	No
Panel A: Cumulative earnings					
$\triangle$ robots per 1000 worker	2.302	7.390***	-1.968	-1.069	-2.051**
	(1.679)	(1.094)	(1.241)	(1.780)	(1.034)
Panel B: Cumulative days					
$\triangle$ robots per 1000 worker	5.631***	16.049***	-5.036**	-1.319	-4.062**
	(1.264)	(2.930)	(2.509)	(3.372)	(1.570)
Panel C: Average daily wage					
$\triangle$ robots per 1000 worker	0.008	0.058	-0.115	-0.027	-0.166**
	(0.065)	(0.039)	(0.081)	(0.119)	(0.065)

Notes. $N = 136,736$. The table reports estimates of β from equation (3.1). Outcomes are cumulated values over 1994-2018. Panels A, B and C shows results for cumulative earnings, cumulative days and average daily wages, respectively. Column (1) cumulates outcomes over all employers. Column (2) cumulates outcomes only at the original employer in the baseline year 1994. Column (3) cumulates outcomes only at different employers in the same 2-digit industry. Column (4) cumulates outcomes at different employers in different 2-digit industries. Column (5) cumulates outcomes at employers outside of manufacturing. All regressions include the full set of controls from column (6) of Table 3.2. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

outcomes from technological progress. The polarization in labour market outcomes may depend also on worker characteristics, such as occupation and skills, and firm characteristics, such as productivity. In this paragraph, we investigate the presence of heterogeneous effects by exploiting the richness of the LoSaI administrative data. We analyse three dimensions of heterogeneity. First, Figure 3.3 reports heterogeneous effects by terciles of AKM worker fixed effects, which we use as proxies of workers' skills.[10] Specifically, we estimate on the full sample over 1985-2018 a two-way fixed effects regression (Abowd et al., 1999) of log daily wages on worker and firm effects, controlling for cubic polynomials in age and tenure, occupation dummies, a part-time dummy – in levels and interacted with a dummy for female workers – and year effects. We then classify workers into terciles of the estimated worker effects.[11] The figure shows that the benefits from automation are highly heterogeneous across workers' skill levels. Panel A shows that only workers in the top

[10] Unfortunately, the dataset does not record education level of workers. Hence, the worker fixed effect will capture the portable component of earnings across employers, which will be correlated with ability and education.

[11] The estimation of the AKM regression requires to restrict the sample to the largest connected group (Abowd et al., 2002) of workers and firms linked by worker mobility. We use the full sample between 1985 to 2018 in

tercile of worker effects benefit from automation exposure with all employers, while both workers in the bottom and middle terciles experience earnings losses. Furthermore, the figure shows that the earnings losses for the latter two worker groups are concentrated among employers different than the original one. In contrast, high-skill workers have larger gains if they remain with their employer, but still do not experience loss in earnings if they relocated to other employers. Panels B and C show how the positive effects for workers in the top tercile materialize. The cumulative number of days worked is similar across skill group at all employers (panel B), while there is a small skill gradient for workers changing employer. On the opposite, high-skill workers experience gains in average daily wages of approximately 0.5 percent across all employers (i.e. approximately 9.9 percent higher wages for a high-skill worker at the 75th percentile of the automation adoption distribution relative to a worker at the 25th percentile). Moreover, except for workers moving out of manufacturing, the figure suggests that high-skill workers benefit from wage increases both at their original employer and at other employers within the same 2-digit industry or within the manufacturing sector. Overall, this evidence suggest that earnings inequality between high- and low-skill workers likely increases as a consequence of higher automation exposure, as only the former experience earnings gains. One finding that the aggregate effect would not be able to signal. The reason why high-skill workers have higher gains from automation may be related to their higher complementarity with new technologies with respect to low-skill workers. At the same time, they are shielded against earnings losses when moving to different employers, likely because of better outside options relative to other worker types.

Figure 3.4 reports heterogeneous effects by occupation in 1994. One concern about the potential negative effects of automation in manufacturing is that it can replace manual tasks performed by workers employed in routine occupations. Blue-collar workers are at higher risk of being negatively affected by automation as they are more likely to perform such routine manual tasks. In fact, the figure shows that automation has a negative overall effect on cumulative earnings of blue-collar employees, while the effect is strongly positive and statistically significant for white-collar workers (panel A). The negative effects for blue-collar workers are entirely concentrated among displaced workers, while both blue- and white-collar have positive earnings effects when staying at their origin employer. Panel B shows that the effect on cumulative days is positive across all employers and is larger for blue- relative to white-collar workers employed in the same firm throughout the period 1994-2018. Blue-collar workers experience, instead, a negative and statistically significant effect on days worked when displaced to the non-manufacturing sector. However, as highlighted before for workers' skill levels, differences by occupation are determined by different wage gains for blue- and white-collar workers. While the former experience a decrease in wages

order to maximize the size of this group. Therefore, we retain observations for 131,961 workers in the baseline year 1994 (97% of the number of workers in the main analysis).

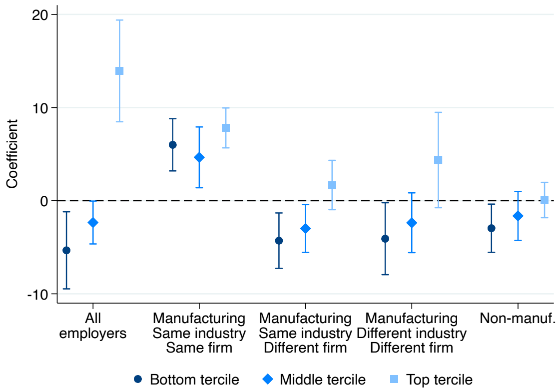

(A) Cumulative earnings

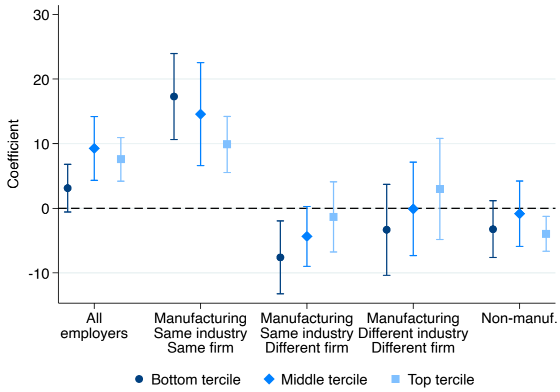

(B) Cumulative days

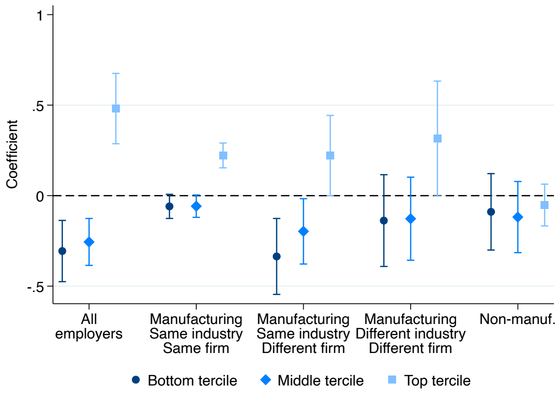

(C) Average daily wage

Fig. 3.3 Heterogeneous effects by terciles of AKM worker fixed effects

Notes. $N = 131{,}96$. The figure reports heterogeneous effects in the effects of robots on cumulative earnings (panel A), cumulative days (panel B) and average daily wages (panel C) by terciles of AKM worker fixed effects. The figures replicate the specifications in Table 3.4. Vertical bars are 95% confidence intervals from cluster-robust standard errors by 36 two-digit NACE Rev.2 industries.

both when staying at their employer and when moving to different employers, for the latter the estimates suggest an opposite result. Except for white-collar workers displaced in a different firm within the same 2-digit industry, the other effects are positive and statistically significant, signalling on the one hand the complementarity between workers employed in white-collar occupations with automation in affected firms and industries and the better employment opportunities of such workers when displaced to other employers or industries.

We further investigate an additional source of heterogeneity, related to firm productivity, measured as value added per worker. The information on firms' balance sheet data is available only for a subset of observations used in the main analysis.[12] However, even in

[12] The number of observations drops to 45,766, approximately 33% of the total number of observations used in the main analysis (136,736).

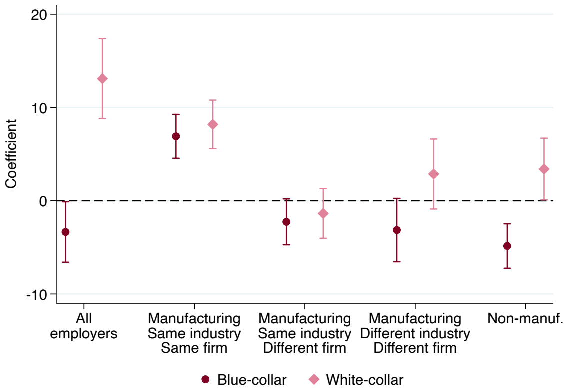

(A) Cumulative earnings

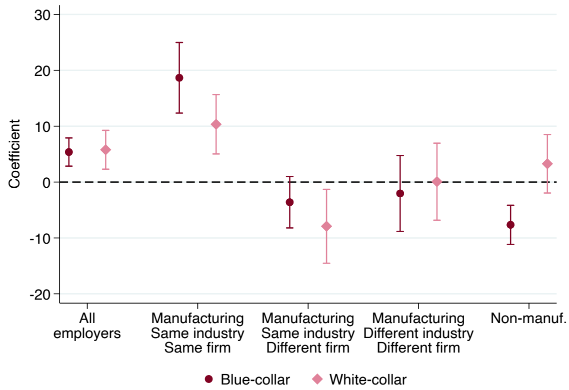

(B) Cumulative days

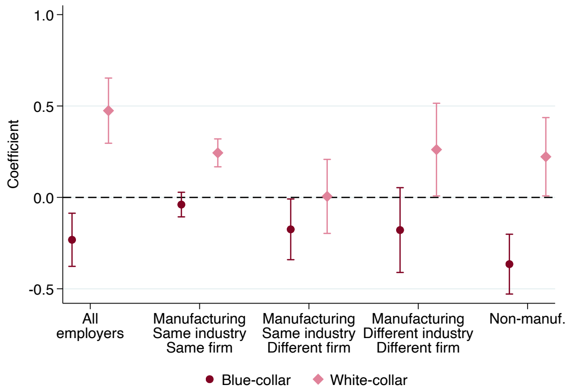

(C) Average daily wage

Fig. 3.4 Heterogeneous effects by occupation

Notes. $N = 136{,}736$. The figure reports heterogeneous effects in the effects of robots on cumulative earnings (panel A), cumulative days (panel B) and average daily wages (panel C) by occupation (blue- or white-collar). The figures replicate the specifications in Table 3.4. Vertical bars are 95% confidence intervals from cluster-robust standard errors by 36 two-digit NACE Rev.2 industries.

this restricted sample the estimated effects of automation are remarkably similar to those in the full sample, as Table C.4 in the Appendix shows. Figure 3.5 shows heterogeneous effects for firms in terciles of the distribution of value added per worker. Firms with higher productivity levels are, on the one hand, more likely to automate – as the evidence for other countries suggests (Koch et al., 2021) – and, on the other hand, they may be as a consequence more likely to preserve employment relationships and share their rents with workers. Panel A shows that there exists a gradient in the cumulative earnings gains across all employers based on firms' value added. This gradient is mainly determined by positive effects for workers employed in firms in the middle and top tercile of the distribution who stay with their employer and a negative (although imprecisely estimated) effect for workers employed in bottom tercile firms. The effects on earnings are mainly determined by positive labour

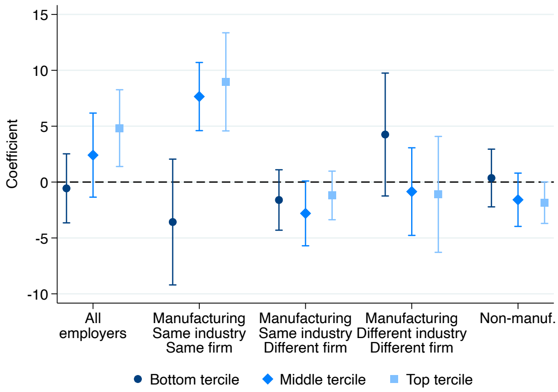

(A) Cumulative earnings

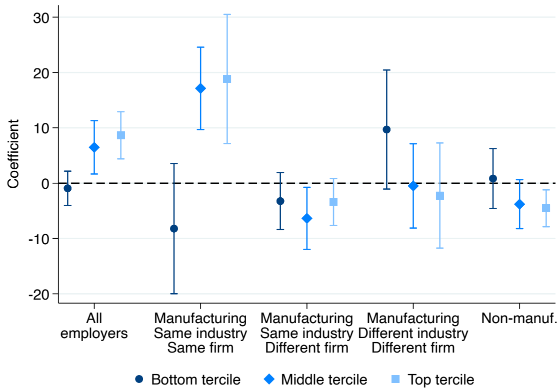

(B) Cumulative days

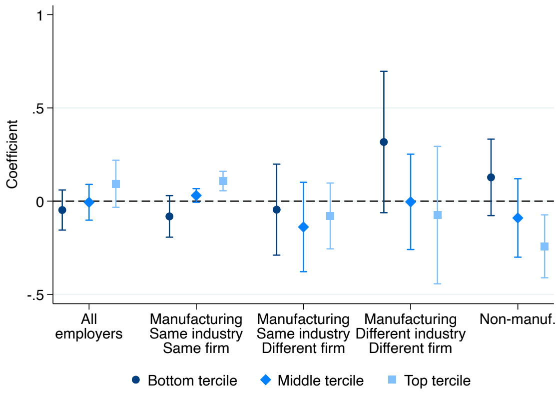

(C) Average daily wage

Fig. 3.5 Heterogeneous effects by terciles of value added per worker

Notes. $N = 45,766$. The figure reports heterogeneous effects in the effects of robots on cumulative earnings (panel A), cumulative days (panel B) and average daily wages (panel C) by terciles of firms' value added per worker. The figures replicate the specifications in Table 3.4. Vertical bars are 95% confidence intervals from cluster-robust standard errors by 36 two-digit NACE Rev.2 industries.

supply effects, as patterns for cumulative days worked closely resemble those for earnings (panel B), rather than by increased daily wages (panel C). Hence, workers employed in highly productive firms increase their labour supply with the firm, by working close to 20 days more in response to one additional robot per worker. This is an important signal of the complementarity between automation and human work and, at the same time, of the productivity effect of automation which likely creates the need for new jobs that workers can take. However, there are striking differences between workers initially employed by more productive and less productive firms, with the latter experiencing worse consequences from technology adoption than the former.

Short-time work, unemployment, sick leaves Table 3.5 investigates the effects of automation on a different set of outcomes, related to unemployment and health of employees. Columns (1) to (3) reports the effects of automation on the cumulative days of short-time work, unemployment and sick leaves. Short-time work programs provide subsidies for hours reductions of workers in firms experiencing temporary shocks. The variation we am exploiting in automation exposure is at the sector-level, which implies that within sectors there are adopters and non-adopters. If non-adopters are "left behind" by the process of technological innovation they may experience temporary shocks before they are able to adapt to new production processes or more permanent shocks if they eventually exit the market. In both cases, these firms may recur to short-time work to preserve employment relationships. The coefficient in panel A suggests a small positive effect which, however, is not statistically significant at conventional levels. Column (2) reports the effect on the cumulative days of unemployment (computed as the number of days of unemployment benefit receipt). The coefficient is negative and statistically significant, suggesting that a worker at the 75th percentile of the automation adoption distribution experiences 23 ($= -1.180 \times [19.86 - 0.13]$) days less of unemployment insurance receipt relative to a worker at the 25th percentile. Considering that the maximum duration of unemployment benefit receipt is 8 months for most of the years in the sample, this is a sizeable effect (approximately 9.5% of maximum duration). Column (3) shows the effect on the cumulative number of sick leave days taken by workers. Automation can have ex-ante positive effects on workers' health as robots can perform tasks that would be dangerous for workers or that would cost physical effort or negative effects if the fear of losing the job increases stress and psychological conditions of workers (see, e.g., Schwabe and Castellacci, 2020). The estimates suggest that neither of these effects translates into additional leave of absence from work, as on average we estimate an almost precise null effect on the number of sick leave days taken by workers.

We further investigate the presence of heterogeneous effects for such outcomes in Figure 3.6. Panel A shows heterogeneity by terciles of worker fixed effects. The null aggregate effect on short-time work masks important heterogeneity by worker type. While low- and middle-skill workers experience additional days of short-time work in response to automation exposure (although the effect is not statistically significant at 95% level), high-skill workers experience a lower amount of days spent in short-time work. We find a similar gradient for the cumulative number of days receiving unemployment benefits and for the number of days spent in sick leave. This last finding, in particular, may suggest that workers who are more likely to suffer displacement effects from automation are also those more likely to have worse health outcomes. This can be due by either increased stress coming from the fear to end up in unemployment or by displacement to non-adopting firms that could use old production processes that might be more physically intensive and may therefore increase health risks for individuals. Results by occupation in panel B show striking differences between blue- and white-collar employees in terms of the cumulative number of days spent

Table 3.5 Robot exposure and cumulative days in short-time work, unemployment and sick leave, 1993-2018 and 1990-1993: 2SLS estimates

	(1) Cumulative STW	(2) Cumulative unemployment	(3) Cumulative sick leave
Panel A: treatment period 1994-2018			
$\triangle$ robots per 1000 worker	0.658	-1.180***	0.003
	(0.710)	(0.229)	(0.021)
Panel B: pre-period 1990-1994			
$\triangle$ robots per 1000 worker	0.305	-0.042	-0.003
	(0.192)	(0.046)	(0.006)

Notes. $N = 136{,}736$. The table reports estimates of β from equation (3.1). Outcomes are cumulated values over 1994-2018 in panel A and over 1990-1994 in panel B. Columns (1) to (3) report the estimates for cumulative days of short-time work, cumulative days of unemployment benefit receipt and cumulative sick leave days. All regressions include the full set of controls from column (6) of Table 3.2. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

in short-time work, which suggests a positive (and imprecisely estimated) effect for blue-collar workers and a negative effect for white-collar, and of the cumulative number of days in sick leave, which again suggests that automation negatively affects health of workers whose employment prospects are more likely to be harmed by automation. There are, instead, non-significant differences in terms of unemployment insurance receipt. Panel C reports heterogeneous effects by terciles of firms' value added per worker. Workers in more productive firms experience a lower amount of days in short-time work, unemployment and sick leave (although for the latter the effect is not statistically significant).

3.4 Firm-level Adjustment

This section analyses the firm-level consequences of automation, by investigating the impact on firms' productivity, revenues, employment and the composition of the labour force. We use data from LoSaI records aggregated at the firm-level, matched with balance sheet data from Cerved over the period 1994-2012.

3.4.1 Empirical Strategy

The unit of observation is a firm and the outcomes are cumulated over the period 1994-2012. For each outcome, we run the following specification

$$Y_{fj} = \omega + \xi \cdot \triangle\widehat{\text{robots}}_j + \zeta \cdot W'_{fj} + \theta \cdot P'_j + \eta_{fj} \tag{3.2}$$

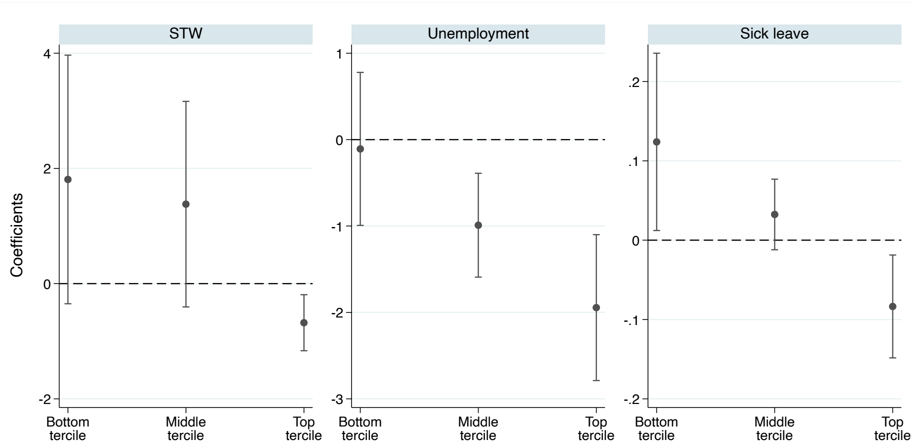

(A) By terciles of worker FEs

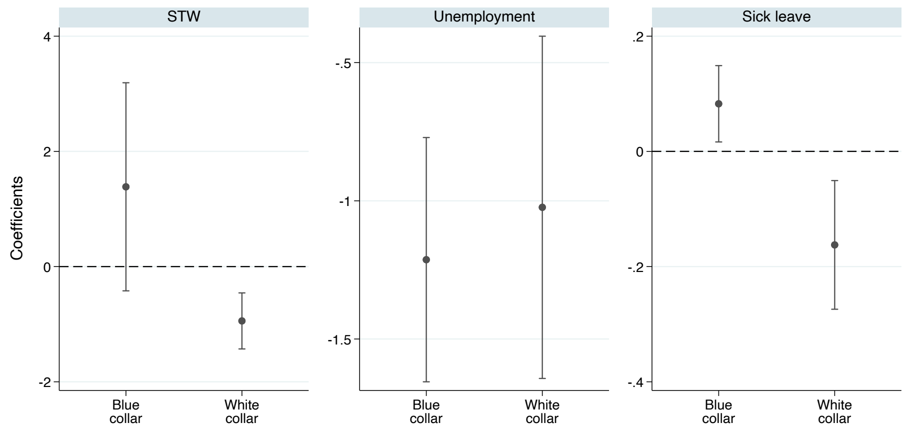

(B) By occupation

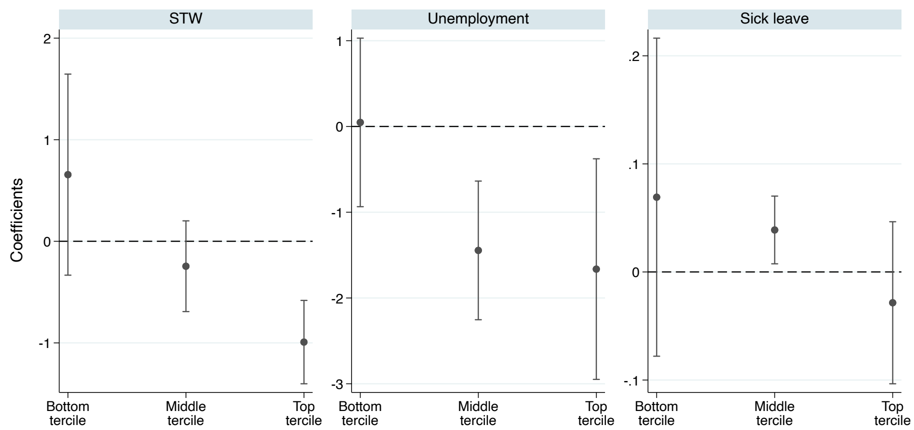

(C) By terciles of VA per worker

Fig. 3.6 Heterogeneous effects for cumulative days of short-time work, unemployment and sick leave

Notes. $N = 136{,}736$. The figure reports heterogeneous effects in the effects of robots on cumulative days of short-time work (left panel), unemployment benefit receipt (central panel) and sick leave (right panel) by terciles of worker fixed effects (panel A), occupation (panel B) and terciles of value added per worker (panel C). All regressions include the set of controls included in column (6) of Table 3.2. Vertical bars are 95% confidence intervals from cluster-robust standard errors by 36 two-digit NACE Rev.2 industries.

where Y_{fj} are cumulated outcomes for firm f in sector j over 1994-2012. ω is a constant. $\triangle\widehat{\text{robots}}_j$ is the predicted change in robot exposure in sector j after instrumentation with the predicted change in other European countries (as before, Finland, France, Germany, Norway, Spain, Sweden and United Kingdom). W_{fj} contains firm-level covariates, such as cohort of birth of the firm, region of location, log average annual earnings and sales in 1994, and firm size dummies (1-10, 11-20, 21-50, 51-100, 101-500, 500 or more employees). P_j contains sector-level controls: dummies for broad manufacturing sectors (food, beverages and tobacco; consumer goods; industrial goods; capital goods), the change in net exports form China and Eastern Europe over the period 1994-2012 and the change in ICT capital between 1996 and 2012.

Given that we only observe a sample of firms in the data, i.e. those that employ one of the workers sampled in the main worker-level data, it is difficult to assess whether the disappearance of a firm from the sample is due to firm closure or to the fact that workers initially sampled in those firms stop working. Therefore, we define two samples. One that uses the full sample of firms, which ignores this potential selection problem, and another one that limits the observations to firms that are either observed for the whole period 1994-2012 or that have a non-missing closing date that is at most two years larger than the last year of firm's appearance in the dataset. We will report results for both samples, which comprise 24,578 and 19,804 observations, respectively.

It is important to highlight that estimates of ξ will be the sum of two effects: one direct effect of firms that automate and the indirect effect of firms that do not automate but are in sectors with high automation adoption. Hence, these effect have to be interpreted with caution as they do not represent the direct effect for adopters, but rather the average effect between adopters and non-adopters in sectors with relatively higher automation exposure.

Table 3.6 reports descriptive statistics from the restricted and full sample in columns 1-2 and 3-4, respectively. Panel A shows the outcomes, which are slightly larger in most cases in the restricted sample on average. Overall, in the restricted sample there is a 19.6 fold increase in the wage bill, an 18.3 fold increase in employment, an 18.9 fold increase in value added and a 20.6 fold increase in sales. In per capita terms, value added is only 2% larger, while sales per worker grow by 12%. The firms included in the sample are in majority born in the 1980s, located in the North of the country and with firm-size between 21 and 50 employees (panel B). Panel C shows that the average firm experiences a change in automation exposure of approximately 6 additional robot per worker

The next paragraph discusses the results from the estimation of equation (3.2).

3.4.2 Results

Table C.5 reports the first stage regressions for the firm-level restricted and full samples in panels A and B, respectively. In both cases the first stage F-statistic is large, signalling

Table 3.6 Descriptive statistics, firm-level sample

	(1)	(2)	(3)	(4)
	Restricted sample		Full sample	
	mean	(sd)	mean	(sd)
Panel A: Outcomes, cumulated values over 1994-2012				
Wage bill / Base year wage bill	19.66	(17.54)	17.53	(17.02)
Employment / Base year employment	18.30	(15.36)	16.38	(14.99)
Value added / Base year value added	18.85	(18.05)	16.87	(17.49)
Value added per worker	1.02	(0.34)	1.03	(0.34)
Sales / Base year sales	20.64	(20.02)	18.41	(19.37)
Sales per worker	1.12	(0.43)	1.12	(0.42)
Exit rate	0.50	(0.50)	0.51	(0.50)
Panel B: Control variables				
Cohort $<$ 1970	0.08	(0.27)	0.08	(0.27)
Cohort 1970	0.34	(0.47)	0.33	(0.47)
Cohort 1980	0.38	(0.48)	0.38	(0.49)
Cohort 1990	0.20	(0.40)	0.21	(0.40)
North-West	0.36	(0.48)	0.36	(0.48)
North-East	0.37	(0.48)	0.37	(0.48)
Centre	0.14	(0.34)	0.14	(0.34)
South	0.13	(0.33)	0.14	(0.34)
Log avg. daily wage, 1994	9.61	(0.24)	9.59	(0.25)
Log avg. sales, 1994	15.37	(1.37)	15.24	(1.39)
Firm size $\leq$ 10	0.17	(0.37)	0.20	(0.40)
Firm size 11-20	0.22	(0.41)	0.23	(0.42)
Firm size 21-50	0.35	(0.48)	0.33	(0.47)
Firm size 51-100	0.16	(0.36)	0.14	(0.35)
Firm size 101-500	0.10	(0.31)	0.09	(0.29)
Firm size $>$ 500	0.00	(0.06)	0.00	(0.06)
Food products, beverages and tobacco	0.06	(0.24)	0.06	(0.24)
Consumer products	0.36	(0.48)	0.36	(0.48)
Industrial goods	0.37	(0.48)	0.37	(0.48)
Capital goods	0.17	(0.37)	0.16	(0.37)
$\triangle$ net exports / wagebill in %	-27.09	(211.94)	-32.39	(233.06)
$\triangle$ ICT equipment in Euro per worker	18.80	(165.68)	16.51	(161.92)
Panel C: Robot exposure				
$\triangle$ robots per 1000 worker	6.05	(12.18)	5.87	(11.97)
p10-p90 interval	[0.05-7.84]		[0.05-7.84]	
p25-p75 interval	[0.24-7.24]		[0.24-7.24]	
Sample size	19,804		24,578	

Notes. The table reports summary statistics from the restricted (columns 1 and 2) and full (columns 3 and 4) firm-level samples (see text for details on the definition of restricted and full samples). Panel A reports means and standard deviations of outcomes, cumulated over the period 1994-2012. Panel B reports control variables in the baseline year 1994. Panel C shows means, standard deviations and quantiles of the distribution of robot exposure, measured as the change in robots per worker by sector between 1994 and 2012.

there is no problem of weak instruments. Table 3.7 reports the estimates from equation (3.2) under different specifications and using cumulative value added as an outcome. Column (1) shows OLS estimates of ξ with the inclusion of firm birth cohort dummies as the only additional controls and signals the presence of a positive and statistically significant effect of automation adoption on value added in both restricted and full samples (panels A and B, respectively). Column (2) reports the 2SLS estimates, which are very close to OLS results. Additional controls are included in columns (3) to (6), such as region dummies, baseline log annual wages and sales, industry and firm size dummies and the change in net exports and ICT capital. Column (6) is the preferred specification that compares the outcomes of firms with similar observable characteristics and indicates that the effect of automation adoption at the sector level on cumulative value added is positive but statistically significant only in the full sample. To give an idea of the economic magnitude of this estimate, comparing a firm at the 25th percentile of the automation exposure to one at the 75th percentile, the latter experiences 27.6% ($= 3.942 \times [7.24 - 0.24]$) higher value added growth.

Table 3.8 shows results for different outcomes using the same specification as the one reported in column (6) of Table 3.7. Firms in sectors more exposed to automation increase employment by 4.3-4.9%, depending on the sample (column 1), and increase the wage bill[13] (column 2). Average earnings – computed as the ratio between cumulative wage bill over cumulative employment – do not increase. Column (4) reports results from Table 3.7 on value added. While there is positive effect on value added, column (5) shows that firms' productivity, measured as value added per worker does not increase. Sales (column 6) increase by 5.1% and 6.1% in the restricted and full sample, respectively. Similarly to value added, sales per worker do not increase (column 7) indicating that higher revenues go hand in hand with increased employment.

Table C.6 in the Appendix reports estimates for additional outcomes. Column (1) shows that even using balance sheet information on labour costs (rather than administrative data on annual wages and employment) does not change the conclusion that labour costs tend to increase in firms exposed to higher automation exposure. Columns (2) and (3) show that firms more exposed to automation do not increase assets and debt levels (for the latter, there is a marginally significant increase in the full sample). Profits decrease (column 4), although the estimates are not statistically significant. There is no statistically significant effect on the number of years with positive profits and on firm exit rate (columns 5 and 6).

Heterogeneity Figure 3.7 reports heterogeneous effects for blue- and white-collar workers. The administrative data records average annual earnings and employment by broad occupation. We therefore compute the occupation-specific wage bill as the product of earnings and employment for each occupation. The figure shows that the increase in the wage bill mainly stems from an increase for blue-collar workers (left panel), determined by an increase

[13] The wage bill is computed as the product of average annual wages with total firm size.

Table 3.7 Robot exposure and firm cumulative value added, 1994-2012: OLS and 2SLS estimates

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS			2SLS		
Panel A: Restricted sample						
△ robots per 1000 worker	6.480**	6.686**	6.646**	6.458***	1.740	2.894
	(2.664)	(2.743)	(2.705)	(2.240)	(1.689)	(1.758)
Hansen-Sargan test		0.232	0.247	0.181	0.515	0.525
Kleibergen-Paap F-statistic		171.7	173.4	171.5	293.5	378.0
Observations	19,802	19,802	19,802	19,802	19,802	19,802
Panel B: Full sample						
△ robots per 1000 worker	7.297***	7.431***	7.344***	7.083***	2.675	3.942**
	(2.321)	(2.401)	(2.377)	(1.967)	(1.631)	(1.666)
Hansen-Sargan test		0.228	0.241	0.195	0.559	0.425
Kleibergen-Paap F-statistic		161.1	162.9	161.0	283.6	365.0
Observations	24,575	24,575	24,575	24,575	24,575	24,575
Cohort dummies	Yes	Yes	Yes	Yes	Yes	Yes
Region dummies			Yes	Yes	Yes	Yes
Baseline sales and wages				Yes	Yes	Yes
Industry and firm size dummies					Yes	Yes
△ Net exports and △ ICT						Yes

Notes. The table reports estimates of ξ from equation (3.2). Column (1) reports OLS estimates. Columns (2) to (6) report 2SLS estimates. Columns (1) and (2) include cohort dummies. Column (3) controls for firms' region of location (20 dummies). Column (4) includes log average sales and wages in 1994. Column (5) includes five firm size dummies (11-20, 21-50, 51-100, 101-500, 500 or more employees) and four dummies for broad manufacturing sectors (food, beverages and tobacco; consumer goods, industrial goods and capital goods). Column (6) further controls for the change in net exports form China and Eastern Europe over 1994-2012 and the change in ICT capital over 1996-2012. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.8 Robot exposure and firm outcomes, 1994-2012: OLS and 2SLS estimates

	(1) Empl.	(2) Wage bill	(3) Avg. earn-ings	(4) Value added	(5) Value added per worker	(6) Sales	(7) Sales per worker
Panel A: Restricted sample							
$\triangle$ robots per 1000 worker	4.307**	3.745*	-0.033	2.894	-0.037	5.066***	-0.009
	(1.782)	(1.938)	(0.024)	(1.758)	(0.060)	(1.788)	(0.053)
Panel B: Full sample							
$\triangle$ robots per 1000 worker	4.899***	4.589**	-0.026	3.942**	-0.031	6.157***	0.005
	(1.682)	(1.786)	(0.022)	(1.666)	(0.050)	(1.705)	(0.046)

Notes. The table reports estimates of ξ from equation (3.2). Outcomes are cumulated values over 1994-2012. Columns (1) to (7) report the estimates for cumulative employment in percent of baseline employment in 1994, cumulative wage bill in percent of baseline wage bill in 1994, average earnings (computed as the ratio of cumulative wage bill over cumulative employment), cumulative value added in percent of baseline value added in 1994, value added per worker (computed as the ratio of value added over employment), cumulative sales in percent of baseline sales in 1994, sales per worker (computed as the ratio of sales over employment). All regressions include the full set of controls from column (6) of Table 3.7. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

in employment (central panel). However, only white-collar workers experience earnings increases (right panel), consistently with results from the worker-level analysis.

We further analyse how firms change their labour force composition, by investigating how employment of low-, middle- and high-skilled workers changes in all firms and in terciles of the value added per worker distribution in Table 3.9.[14] Panel A shows results for the full sample of workers for which we have information on workers' skills, i.e. 22,537 firms (out of a total of 24,575 observations in the main sample). This sample restriction does not affect the estimate on aggregate employment, which is remarkably similar to the one reported in Table 3.8 (4.356 vs 4.307). When decomposing this effect by worker skills, in columns (2) to (4) we find that it is especially employment of low- and middle-skilled workers that is increasing. This results may be surprising but it is actually very similar to the one found in the worker-level analysis, where the within-firm change in labour supply is actually larger for less skilled workers. This result may also suggest that the productivity effect of automation compensates the negative displacement effect in sectors that are most

[14] We compute the number of low-, middle- and high-skill individuals by combining worker-level data on a sample of workers for each firm and administrative data on total employees for each firm. Specifically, we compute the share of workers in terciles of AKM worker fixed effects employed at each firm and multiply these shares by the total number of employees to get the number of employees in each skill group. We corroborate this approach by using the same procedure to build estimates of the number of blue-collar and white-collar employees, for which we have administrative counts, too. This procedure yields an across person-years average number of blue-collar employees of 34.9 against a true value of 32.2 and an average number of white-collar employees of 20.3 against a true value of 20.8.

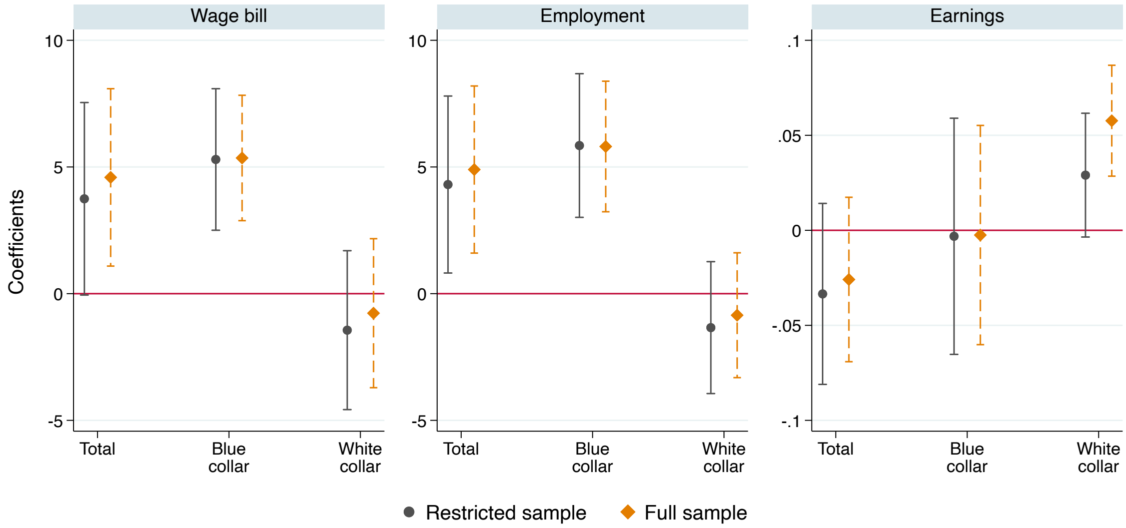

Fig. 3.7 Robot exposure and firm-level outcomes for blue-collar and white-collar workers, 1994-2012: 2SLS estimates

Notes. The figure reports heterogeneous effects by occupation in the estimates of ξ from equation (3.2). Outcomes are cumulated values over 1994-2012. Left panel reports the estimates for cumulative wage bill in percent of baseline wage bill in 1994. Central panel reports the estimates for cumulative employment in percent of baseline employment in 1994. Right panel reports the estimates for average earnings (computed as the ratio of cumulative wage bill over cumulative employment). All regressions include the full set of controls from column (6) of Table 3.7. Vertical lines are 95% confidence intervals from cluster-robust standard errors by 36 two-digit NACE Rev.2 industries.

affected by automation. It must be interpreted, anyway, as a partial equilibrium effect. Panel B decomposes these effects by firms' productivity, as measured by terciles of value added per worker in the baseline. Column (1) shows that the positive effect on employment is almost entirely driven by highly productive firms (which are also more likely to be adopters, as shown in Koch et al., 2021). Firms in the top tercile of productivity increase employment by 12% relative to the baseline, i.e. by 85% ($= 12.205 \times [7.24 - 0.24]$) if we compare firms at the 25th and 75th percentile of automation exposure. When further decomposing the effect for each tercile of firm's productivity into terciles of workers' skills, columns 2-4 of panel B suggest that firms in the top tercile increase employment of all worker types, whereas firms in the bottom and middle terciles reduce employment of high-skill workers (although for the latter the effect is not statistically significant). This result may signal a reallocation of high-skill workers to more productive firms, which however may not lead to increased assortative matching in the labour market as more productive firms increase employment of other worker types, too. Therefore, these results seem to suggest that highly productive firms increase demand of highly skilled workers, probably displacing less skilled workers. But the productivity effect compensates for this initial displacement as demand for all worker types increases at more productive firms. It remains open the question of whether these

Table 3.9 Robot exposure and firm employment by worker skills and firms' productivity, 1994-2012: 2SLS estimates

	(1) Employment	(2) Employment by worker skill	(3)	(4)
		Bottom tercile	Middle tercile	Top tercile
Panel A: Full sample				
△ robots per 1000 worker	4.356**	2.276*	2.103***	-0.397
	(1.767)	(1.215)	(0.569)	(0.945)
Panel B: By VA/Empl tercile				
△ robots bottom tercile	-2.128	1.797	-0.517	-3.111*
	(3.229)	(1.931)	(1.267)	(1.660)
△ robots middle tercile	2.640*	1.525	1.890***	-1.011
	(1.467)	(1.083)	(0.726)	(0.830)
△ robots top tercile	12.205***	3.645***	4.693***	2.726**
	(2.699)	(1.047)	(1.211)	(1.339)
Observations	22,537	22,537	22,537	22,537

Notes. The table reports estimates of ξ from equation (3.2) in the sample of firms with non-missing information about workers' skills (see text for details). Outcomes are cumulated values over 1994-2012. Panel A reports results for the full sample. Panel B decomposes the effect by terciles of firms' value added per worker. Column (1) reports the estimates for cumulative employment in percent of baseline employment in 1994, whereas columns (2) to (4) report the estimates for cumulative employment of workers in the bottom, middle and top terciles of worker skills. All regressions include the full set of controls from column (6) of Table 3.7. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

reallocation effects increase assortative matching in the labour market, as there are two competing forces of opposite sign: the reallocation of high-skill workers to more productive firms (which increases sorting) and the increase in employment of low- and middle-skill workers in more productive firms (which decreases sorting). The question is investigated in the next section by taking a close look at the sector-level intricacies.

3.5 Sector-level Adjustment

In this section, we further analyse how the allocation of workers across firms changes by investigating how the sector-level correlation of worker and firm types changes around automation spikes in an event study framework.

3.5.1 Empirical Strategy

The objective is to investigate what happens to the allocation of workers across firms in a sector in the aftermath of automation adoption. To identify automation episodes, we follow a strategy similar to the one used by Bessen et al. (2019) and identify automation spikes, i.e. positive changes in the number of robots per worker in a sector in a given year that are at least 3 times larger than the average change in other sectors in the same year. We use the year in which such automation spike happens for a sector as the event year. Figure 3.8 provides a visual inspection of the automation spike relative to event time periods (i.e. the difference between calendar and event years, which correspond to the years when the spikes happen). Before the automation spike, sectors have similar changes in robot per worker. In the event year, there is indeed a spike which then fades out after approximately three years. We exploit variation in treatment timing – where treatment is defined as the period after the automation spike – and run an event study analysis to understand how workers allocation across firms is affected by automation spikes. Specifically, we compute for each sector-year the correlation between worker fixed effects and log value added per worker or log sales per worker, which we use as proxies of the degree of assortative matching in a given sector. We do not use other standard measures of sorting, such as the correlation between AKM worker and firm fixed effects for two reasons: first, as highlighted by Eeckhout and Kircher (2011), firm fixed effects represent poor proxies of true firm types; second, as pointed by Andrews et al. (2008); Bonhomme et al. (2019) the correlation between worker and firm effects is likely to provide a downwardly biased estimate of the true correlation in the presence of limited mobility bias, which is particularly worrisome in this case as we only observe a sample of workers for each firm. Hence, we follow Bombardini et al. (2019) and use administrative data to measure firm types rather than AKM fixed effects.

The estimation of the dynamic effects requires an event-study methodology. The staggering design and the absence of a never treated group (as all sectors at some point experience the automation spike in the sample) raises concerns in the presence of treatment effect heterogeneity. In particular, as highlighted by Callaway and Sant'Anna (2020); de Chaisemartin and D'Haultfoeuille (2020); Goodman-Bacon (2018); Sun and Abraham (2020),[15] treatment effects can be severely biased if early-treated units are used as control groups for later-treated units. To overcome this issue, we adopt a stacked dynamic differences-in-differences approach (Cengiz et al., 2019; Deshpande and Li, 2019; Vannutelli, 2020), where only yet-to-be treated units are used as controls for already treated units. In other words, the cohort of sectors treated in year t, uses sectors treated from $t+1$ onwards as control groups, the cohort of sectors treated in year $t+1$ uses sectors treated from $t+2$ onwards as control group, and so on and so forth until the last treated cohort which acts as control group only (as it would not have a valid yet-to-be treated unit that could act as control). This stacked design

[15] See Baker et al. (2021) for a review of the methods developed in the literature in recent times.

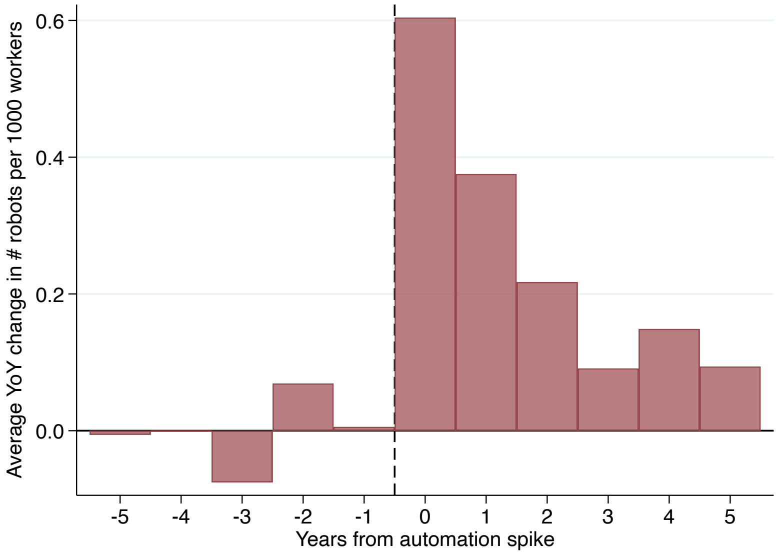

Fig. 3.8 Change in the number of robots by event time

Notes. The figure reports the average year on year change in the number of robots per worker across sectors by time relative to the automation spike, i.e. the year in which a sector increases the number of robots per worker by three times more than the average change of other sectors.

allows units (sectors) to appear more than once in the data and therefore allows for the inclusion of unit fixed effects on top of a treatment dummy (as there is variation in treatment within units). The stacked sample consists of 1,011 observations. We therefore estimate the following regression model:

$$\begin{aligned} Y_{jt} &= \alpha_j + \delta_t + \beta_0 Treat_j + \sum_{k \neq -1} \beta_k Treat_j \times \mathbb{1}(t - E = k) \\ &+ \sum_{k \neq -1} \gamma_k \times \mathbb{1}(t - E = k) + \sum_{k \neq -1} \eta_k X_j^{pre} \times \mathbb{1}(t - E = k) + \iota f(s_{jt}) + \varepsilon_{jt} \end{aligned} \tag{3.3}$$

where Y_{jt} is the correlation between worker fixed effects and firm type (measured with log value added per worker or log sales per worker). α_j and δ_t are sector and year fixed effects. $\mathbb{1}(t - E = k)$ are event time dummies (where E is the event year and k are relative time indicators). $Treat_j$ is a dummy equal to one for sectors after the automation spike. X_j^{pre} are covariates measured in the year prior to the automation spike, namely average log sales per worker, average log earnings, the share of blue-collar workers, the share of female workers, the share of workers below age 40 and the share of firms with more than 100 employees. $f(s_{jt})$ is a quadratic polynomial in the number of sampled workers in each sector from the original worker-level data. ε_{jt} is an error term.

We also report results from a static specification, where event time dummies are replaced with one indicator for the post-treatment period:

$$
\begin{aligned}
Y_{jt} = \alpha_j + \delta_t + \beta_0 Treat_j + \beta_{DD} Treat_j \times Post_t \\
+ \sum_{k \neq -1} \gamma_k \times \mathbb{1}(t - E = k) + \eta X_j^{pre} \times Post_t + \iota f(s_{jt}) + \varepsilon_{jt}
\end{aligned}
\tag{3.4}
$$

where the only difference with equation (3.3) is $Post_t$ which is a dummy taking value one in years after the automation spike.

Table 3.10 reports descriptive statistics of the sector-level data. Panel A reports descriptives about the outcomes. The average correlation between worker effects and log value added per worker is 0.36, whereas the correlation with log sales per worker is 0.32. These values are considerably larger than the correlation between AKM worker and firm effects which is around 0.08 (or 0.09 when grouping the fixed effects into deciles). The difference in the measure of sorting when using firm effects or value added and sales likely signals that AKM firm effects are poor proxies for true firm types, which likely results in a downward bias in their correlation with worker effects. For this reason we report results only for the correlation of worker effects with firm's balance sheet variables, which are likely to provide better measures of the underlying firm type (as suggested, for example, in Bombardini et al., 2019). These outcomes must therefore used with caution (Bombardini et al., 2019; Eeckhout and Kircher, 2011). Panel B reports descriptives about control variables. The majority of workers are blue-collar, male and older than 40 years old. The majority of firms have more than 100 employees. The bottom part of the table reports sample size, which is 361 in the original sample (19 sectors over the period 1994-2012) and 1,011 in the stacked sample, where sectors can appear more than once (as treated units or as yet to be treated units, which act as controls).

3.5.2 Results

Figure 3.9 reports the estimates of equation (3.3) with and without the inclusion of additional control variables. Each figure also reports the static difference-in-differences estimate from equation (3.4) for both specifications with and without controls. Panel A reports the results for the correlation between worker fixed effects and log value added per worker. Before the automation spike the estimates are not significantly different from zero. After the automation spike, the correlation decreases and such decrease is statistically significant up to the fourth year after the spike in the specification without controls and up to the second year after the spike in the specification with controls. A similar conclusion is reached by looking at panel B, that reports the correlation of worker effects with log sales per worker: in this case, the estimates are statistically significant for a shorter period of time, but the trajectory of

Table 3.10 Descriptive statistics, sector-level data

	(1) mean	(2) (sd)
Panel A: Outcomes		
Correlation worker effects and log value added per worker	0.356	(0.140)
Correlation worker effects and log sales per worker	0.320	(0.139)
Correlation worker and firm effects	0.083	(0.173)
Correlation worker and firm effect deciles	0.088	(0.147)
Panel B: Control variables		
Log value added per worker	10.945	(0.427)
Log sales per worker	12.217	(0.536)
Log earnings	10.163	(0.225)
Share blue-collar	0.610	(0.158)
Share female	0.224	(0.102)
Share younger than 40 years old	0.472	(0.085)
Share firms w/ $>$ 100 empl.	0.764	(0.116)
Original sample size	361	
Stacked sample size	1,011	

Notes. The table report means and standard deviations of the sector-level sample. Panel A reports the measures of labour market sorting, i.e. the correlation between worker effects and firm types. Panel B reports summary statistics about control variables. The bottom part of the table shows sample sizes of the original and stacked sample, the latter being used in the estimation.

the coefficients is comparable to that shown in Panel A. The static difference-in-differences estimates suggest a negative effect of -6.6 to -4.8 log points in panel A and of -3.8 to -3.7 log points in panel A, i.e. between -18.5% and -13.5% and between -11.9% and -11.6% of the average correlations across sectors, respectively. This result resolves the ambiguity of the firm-level analysis: there is a decrease in assortative matching in the immediate aftermath of a surge in automation exposure, confirming that firms adjust their workforce by increasing demand of all worker types (Table 3.9). This can be a consequence of a number of different mechanisms: frictions in the labour market that impede timely reallocation of workers across firms, employment protection legislation that protects low-skill workers' employment in more productive adopting firms or the prevalence of a productivity effect that compensates the displacement effect induced by automation. Figure C.1 in the appendix replicates the event study using the methodology proposed by de Chaisemartin and D'Haultfoeuille (2020) and finds very similar results to those reported in the main text, reassuring on the validity of the stacked design.

Overall, these results suggest that there is no increase in assortativeness in the labour market in response to automation spikes. This result is consistent with the firm-level analysis which suggests the presence of a re-allocation effect of high-skill workers from low- to high-

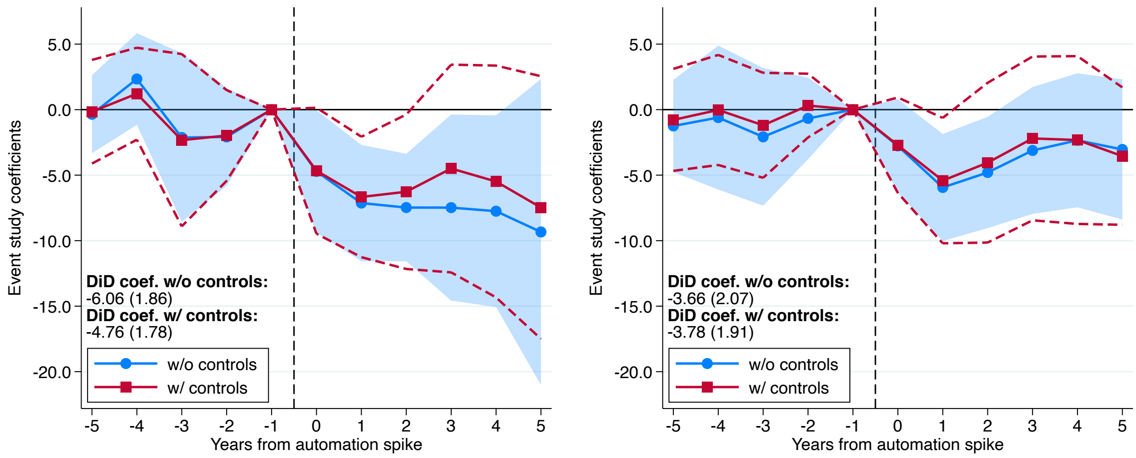

(A) Correlation worker effect and value added (B) Correlation worker effect and sales

Fig. 3.9 Event study of the effect of automation spike on assortative matching in the labour market

Notes. The figure reports the estimates of the dynamic effects β_k from equation (3.3) and the static difference-in-differences coefficient β_{DD} from equation (3.4). The estimates are reported for two specifications: without controls – which only includes year fixed effects and a quadratic polynomial in the number of sampled workers in each sector from the original worker-level data as additional covariates – and with controls – which includes the following control variables, measured in the year prior to the automation spike: average log sales per worker, average log earnings, the share of blue-collar workers, the share of female workers, the share of workers below age 40 and the share of firms with more than 100 employees. The shaded area and dashed lines are 95% confidence intervals from cluster-robust standard errors by 19 sectors (aggregated as in Table C.1 in the appendix).

productivity firms. Nevertheless, the latter group of firms also increases employment of low- and middle-skill workers.

3.6 Conclusion

This paper provides a comprehensive analysis of the effects of automation on the labour market, by investigating the consequences of automation on workers, firms and the allocation of workers across firms at the sector-level. The main conclusion is that the effects of automation are on aggregate null or even positive on workers' employment and earnings prospects and on firms' productivity. At the sector level, there is evidence that labour market sorting decreases after automation spikes in the short-run. However, there is large variation in the effects of automation across different workers and firms. Workers with lower skill levels, in blue-collar occupations and in firms with lower productivity levels experience worse labour market outcomes in terms of cumulative earnings, stemming from either a lower number of days worked or lower wages. Moreover, these groups of workers are more likely to spend additional time in unemployment, in short-time work compensation schemes and in sick leave.

At the firm-level, we find that firms in sectors with higher automation exposure increase employment, labour costs, revenues and value added, but do not significantly increase value added and sales per worker. The effects are extremely heterogeneous according to firm's productive levels as measured by firm's value added per worker. Firms in the first tercile of the distribution of value added per worker decrease employment and significantly so for high-skill workers, while firms in the middle and especially in the top of the value added distribution increase employment. Moreover, the latter group of firms increases employment of all worker skills. To further analyse how the allocation of workers across firms is affected by automation, we use sector-level data on the correlation between worker fixed effects and log value added or sales per worker to investigate how sudden increases in automation adoption in a sector affect sorting. We find that automation, coherently with the results of the firm-level analysis, decreases labour market sorting in the short-run.

Overall, these analyses provide evidence on the impact of automation in a country with a strong manufacturing sector (second largest manufacturing in Europe, after Germany) and with a relatively rigid labour market with high employment protection legislation. The evidence on the impact of automation is so far mixed, with research finding large negative results, especially in the US context, and others finding null to moderately positive results, especially in Europe. This paper adds to the literature a set of findings that show how automation has on aggregate positive effects on both employment and earnings of workers in manufacturing. However, when exploring the effects for different worker or firm types, there is a clear distinction between "winners" and "losers", with workers – especially those employed in less skill-intensive occupations – facing bigger losses from technology adoption. More research in the future should also focus on sectors beyond manufacturing and on firm-level decisions on whether to introduce automation, rather than relying on less granular sector-level information. The increased availability of administrative data and on combinations of administrative and survey information will likely provide a fertile soil on which to build future answers to these research questions.

References

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

Acemoglu, D. and Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2):3–30.

Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. (2020). What Are the Labor and Product Market Effects of Automation?: New Evidence from France. Sciences Po publications info:hdl:2441/170cd4sul89, Sciences Po.

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

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

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

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

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

Andrews, M. J., Gill, L., Schank, T., and Upward, R. (2008). High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias? *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 171.

Arcidiacono, P., Foster, G., Goodpaster, N., and Kinsler, J. (2012). Estimating spillovers using panel data, with an application to the classroom. *Quantitative Economics*, 3(3):421–470.

Arellano-Bover, J. and Saltiel, F. (2021). Differences in On-the-Job Learning Across Firms. *Working Paper*.

Argaw, B. A., Maier, M. F., and Skriabikova, O. J. (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., Booth, A. L., and Bryan, M. L. (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., Casarico, A., and Voitchovsky, S. (2018). Top Incomes and the Gender Divide. *The Journal of Economic Inequality*, 16(2):225–256.

Autor, D. H., Dorn, D., Hanson, G. H., and Song, J. (2014). Trade Adjustment: Worker-Level Evidence. *The Quarterly Journal of Economics*, 129(4):1799–1860.

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

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

Baker, A., Larcker, D. F., and Wang, C. C. Y. (2021). How Much Should We Trust Staggered Difference-In-Differences Estimates? Available at ssrn:, European Corporate Governance Institute – Finance Working Paper No. 736/2021.

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

Bartolucci, C., Devicienti, F., and Monzón, I. (2018). Identifying Sorting in Practice. *American Economic Journal: Applied Economics*, 10(4):408–38.

Battisti, M. (2017). High wage workers and high wage peers. *Labour Economics*, 46(February 2015):47–63.

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

Becker, G. (1964). *Human Capital*. Chicago: Univ. Chicago Press.

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).

Ben-Porath, Y. (1967). The Production of Human Capital and the Life Cycle of Earnings. *Journal of Political Economy*, 75(4):352–365.

Bennedsen, M., Simintzi, E., Tsoutsoura, M., and Wolfenzon, D. (2019). Do Firms Respond to Gender Pay Gap Transparency? Working Paper 25435, National Bureau of Economic Research.

Bertrand, M. (2011). New Perspectives on Gender. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 4, chapter 17, pages 1543–1590. Elsevier.

Bertrand, M., Black, S. E., Jensen, S., and Lleras-Muney, A. (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.

Bertrand, M., Goldin, C., and Katz, L. F. (2010). Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors. *American Economic Journal: Applied Economics*, 2(3):228–55.

Bessen, J., Goos, M., Salomons, A., and van den Berge, W. (2019). Automatic Reaction – What Happens to Workers at Firms that Automate? CPB Discussion Paper 390, CPB Netherlands Bureau for Economic Policy Analysis.

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

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

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

Blau, F. D. and Kahn, L. M. (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 Kahn, L. M. (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., Javorcik, B., and Ulltveit-Moe, K. H. (2018). Working Across Time Zones: Exporters and the Gender Wage Gap. *Journal of International Economics*, 111:122–133.

Bombardini, M., Orefice, G., and Tito, M. D. (2019). Does exporting improve matching? Evidence from French employer-employee data. *Journal of International Economics*, 117(C):229–241.

Bonfiglioli, A., Crinò, R., Fadinger, H., and Gancia, G. (2020). Robot Imports and Firm-Level Outcomes. CRC TR 224 Discussion Paper Series 243, University of Bonn and University of Mannheim, Germany.

Bonhomme, S., Lamadon, T., and Manresa, E. (2019). A distributional framework for matched employer employee data. *Econometrica*, 87(3):699–739.

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

Brune, L., Chyn, E., and Kerwin, J. (2020). Peers and Motivation at Work. *Journal of Human Resources*, pages 0919–10416R2.

Bruno, G. S. F., Ciavarella, A., and Linciano, N. (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.

Caldwell, S. and Harmon, N. (2019). Outside Options, Bargaining, and Wages: Evidence from Coworker Networks. *Job Market Paper*.

Callaway, B. and Sant'Anna, P. H. (2020). Difference-in-differences with multiple time periods. *Journal of Econometrics*.

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

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

Card, D., Cardoso, A. R., and Kline, P. (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., Heining, J., and Kline, P. (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. *Quarterly Journal of Economics*, 128(3):967–1015.

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

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The Effect of Minimum Wages on Low-Wage Jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454.

Comi, S., Grasseni, M., Origo, F., and Pagani, L. (2020). Where Women Make a Difference: Gender Quotas and Firms' Performance in Three European Countries. *ILR Review*, 73(3):768–793.

Cornelissen, T., Dustmann, C., and Schönberg, U. (2017). Peer effects in the workplace. *American Economic Review*, 107(2):425–456.

Cortés, P. and Pan, J. (2020). Children and the remaining gender gaps in the labor market. Working Paper 27980, National Bureau of Economic Research.

Coudin, E., Maillard, S., and To, M. (2018). Family, Firms and the Gender Wage Gap in France. IFS Working Papers W18/01, Institute for Fiscal Studies.

Dauth, W., Findeisen, S., Moretti, E., and Suedekum, J. (2018). Matching in Cities. Working Paper 25227, National Bureau of Economic Research.

Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association*.

de Chaisemartin, C. and D'Haultfoeuille, X. (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, 110(9):2964–96.

Del Bono, E. and Vuri, D. (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.

Deshpande, M. and Li, Y. (2019). Who Is Screened Out? Application Costs and the Targeting of Disability Programs. *American Economic Journal: Economic Policy*, 11(4):213–48.

Devicienti, F., Fanfani, B., and Maida, A. (2019). Collective Bargaining and the Evolution of Wage Inequality in Italy. *British Journal of Industrial Relations*, 57(2):377–407.

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

Dixon, J., Hong, B., and Wu, L. (2020). The Robot Revolution: Managerial and Employment Consequences for Firms. Available at ssrn:, NYU Stern School of Business.

Dottori, D. (2020). Robots and employment: evidence from Italy. Questioni di Economia e Finanza (Occasional Papers) 572, Bank of Italy, Economic Research and International Relations Area.

Eeckhout, J. and Kircher, P. (2011). Identifying Sorting—In Theory. *The Review of Economic Studies*, 78(3):872–906.

Engbom, N. and Moser, C. (2020). Firm Pay Dynamics. MPRA Paper 98477, University Library of Munich, Germany.

Falk, A. and Ichino, A. (2006). Clean Evidence on Peer Effects. *Journal of Labor Economics*, 24(1):39–57.

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

Ferrari, G., Ferraro, V., Profeta, P., and Pronzato, C. D. (2018). Do Board Gender Quotas Matter? Selection, Performance and Stock Market Effects. IZA Discussion Papers 11462, Bonn.

Flabbi, L., Macis, M., Moro, A., and Schivardi, F. (2019). Do Female Executives Make a Difference? The Impact of Female Leadership on Gender Gaps and Firm Performance. *The Economic Journal*, 129(622):2390–2423.

Flinn, C. J., Todd, P. E., and Zhang, W. (2018). Personality traits, intra-household allocation and the gender wage gap. *European Economic Review*, 109:191–220.

Gagliarducci, S. and Paserman, D. (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.

Goodman-Bacon, A. (2018). Difference-in-Differences with Variation in Treatment Timing. Working Paper 25018, National Bureau of Economic Research.

Graetz, G. and Michaels, G. (2018). Robots at Work. *The Review of Economics and Statistics*, 100(5):753–768.

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

Gregory, V. (2019). Firms as Learning Environments: Implications for Earnings Dynamics and Job Search. *Job Market Paper*.

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

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

Herkenhoff, K., Lise, J., Menzio, G., and Phillips, G. M. (2018). Production and Learning in Teams. *National Bureau of Economic Research Working Paper Series*, No. 25179.

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

Hong, L. and Mikkel, S. (2020). Estimation and Inference in a Peer Effects Model with Heteroscedasticity. *Mimeo*.

IFR (2020). World Robotics 2020. Technical report.

Jarosch, G., Oberfield, E., and Rossi-Hansberg, E. (2021). Learning From Coworkers. *Econometrica*, 89(2):647–676.

Jewell, S. L., Razzu, G., and Singleton, C. (2020). Who Works for Whom and the UK Gender Pay Gap. *British Journal of Industrial Relations*, 58(1):50–81.

Kleibergen, F. and Paap, R. (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics*, 133(1):97–126.

Klenert, D., Fernandez-Macias, E., and Anton, J.-I. (2020). Do robots really destroy jobs? Evidence from Europe. JRC Working Papers on Labour, Education and Technology 2020-01, Joint Research Centre (Seville site).

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

Kline, P., Saggio, R., and Sølvsten, M. (2019). Leave-out Estimation of Variance Components. NBER Working Papers 26244, National Bureau of Economic Research, Inc.

Koch, M., Manuylov, I., and Smolka, M. (2021). Robots and Firms. *The Economic Journal*.

Lengermann, P. A. (2002). Is it Who You Are, Where You Work, or With Whom You Work? Reassessing the Relationship Between Skill Segregation and Wage Inequality. Technical report.

Lopes de Melo, R. (2018). Firm Wage Differentials and Labor Market Sorting: Reconciling Theory and Evidence. *Journal of Political Economy*, 126(1):313 – 346.

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

Manski, C. F. (1993). Identification of Endogenous Social Effects: The Reflection Problem. *The Review of Economic Studies*, 60(3):531–542.

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

Morchio, I. and Moser, C. (2019). The Gender Gap: Micro Sources and Macro Consequences. Technical report.

Nix, E. (2020). Learning Spillovers in the Firm. *Working Paper*.

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

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

Papay, J. P., Taylor, E. S., Tyler, J. H., and Laski, M. E. (2020). Learning Job Skills from Colleagues at Work: Evidence from a Field Experiment Using Teacher Performance Data. *American Economic Journal: Economic Policy*, 12(1):359–88.

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

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

Ponthieux, S. and Meurs, D. (2015). Gender Inequality. In Atkinson, A. B. and Bourguignon, F., editors, *Handbook of Income Distribution*, volume 2 of *Handbook of Income Distribution*, chapter 12, pages 981–1146. Elsevier.

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

Rogerson, W. P. (1985). The First-Order Approach to Principal-Agent Problems. *Econometrica*, 53(6):1357–67.

Rozada, M. G. and Yeyati, E. L. (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.

Sandvik, J. J., Saouma, R. E., Seegert, N. T., and Stanton, C. T. (2020). Workplace Knowledge Flows. *The Quarterly Journal of Economics*, 135(3):1635–1680.

Schwabe, H. and Castellacci, F. (2020). Automation, workers' skills and job satisfaction. *PLOS ONE*, 15(11):1–26.

Serafinelli, M. (2019). "Good" firms, worker flows, and local productivity. *Journal of Labor Economics*, 37(3):747–792.

Song, J., Price, D. J., Guvenen, F., Bloom, N., and von Wachter, T. (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.

Sun, L. and Abraham, S. (2020). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*.

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.

Vannutelli, S. (2020). From Lapdogs to Watchdogs: Random Auditor Assignment and Municipal Fiscal Performance in Italy. Technical report.

Appendix A

Appendix for Chapter 1

A.1 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:

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

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

where

$$\psi_j^g = \gamma^g \overline{S}_j \tag{A.5}$$

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

Equation (A.4) is consistent with the two-way fixed effects model *á la* Abowd et al. (1999) presented in equation (1.1) in the main text.

A.2 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 (1.1), 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 A.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 gains and losses with each opposite move is better assessed examining the magnitudes of such changes over the entire 4 year

[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 1.5. Means are computed within each year. Data on the mean earnings for all the moves are reported in Table A.6.

Fig. A.1 Mean weekly earnings of movers across firm effects quartiles

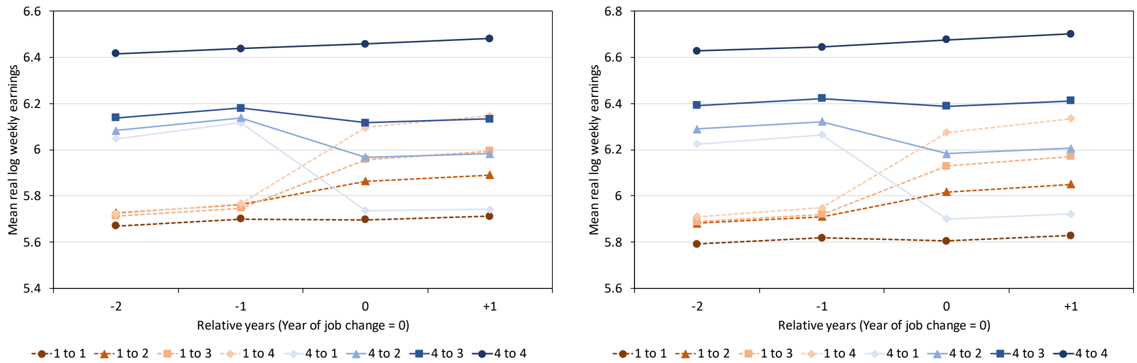

(A) Female sample (B) Male sample

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 A.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.

As an additional check, panels A and B of Figure A.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

[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 $q \neq q'$.

Fig. A.2 Adjusted change in earnings of symmetric job moves across firm effects quartiles

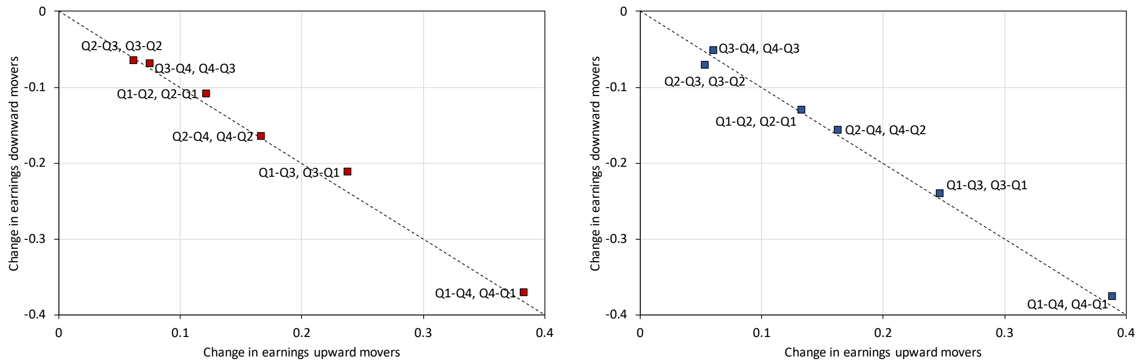

(A) Female sample

(B) Male sample

Fig. A.3 Mean weekly earnings of movers within same firm effects quartiles

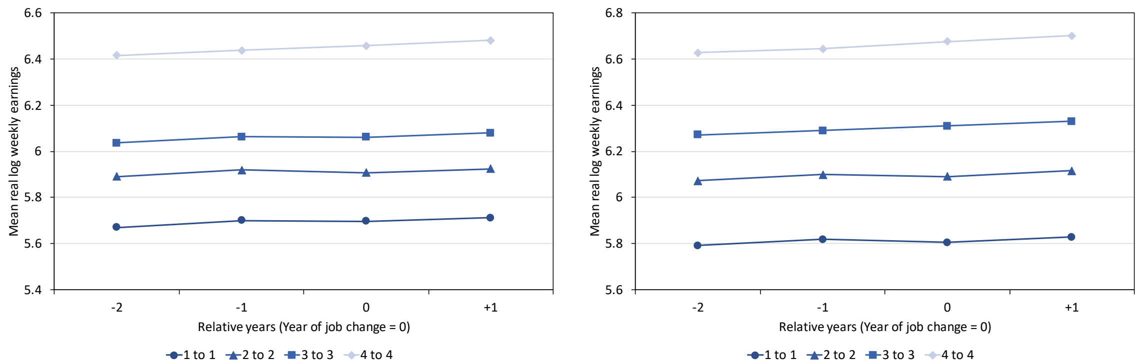

(A) Female sample

(B) Male sample

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 A.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 (1.1) 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

Fig. A.4 Mean AKM residuals across deciles of person and firm effects

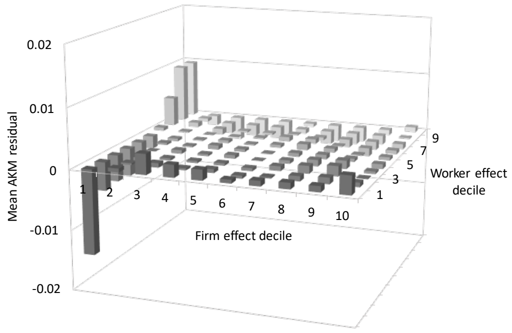

(A) Female sample

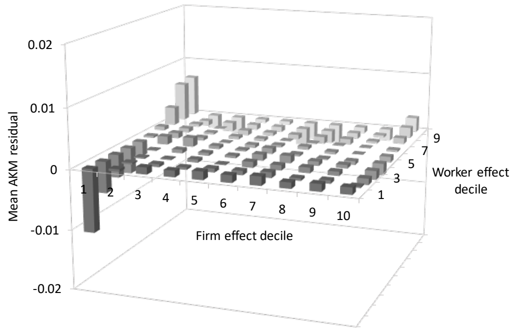

(B) Male sample

residuals in those cells that are most threatened by misspecification. Figure A.4 plots 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.

A.3 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, we follow Card et al. (2016) and set to zero the average firm effect in the food and accommodation sector.

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 1.2). 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 A.5 plots the 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.

We decompose firm effects as in equations (1.4) and (1.5). Results are reported in Table A.1. With this alternative normalisation, the firm contribution to the gender pay gap increases. The difference in firm effects accounts for 38 per cent 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 per cent of the gender pay gap. Besides this, the main conclusions do not change. This alternative normalisation shows that our estimate of the bargaining channel in the main text can be interpreted as a lower bound. However, we prefer the normalisation with respect to the food and accommodation sector because we have information on sectors for *all* firms in our sample, whereas we lose around 20 per cent of person-year observations in the normalisation based on log value added per worker. Since the main conclusions remain qualitatively unchanged, we choose to keep as many observations as possible in the normalisation procedure.

[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.

[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.

Fig. A.5 Firm effects against log value added per worker

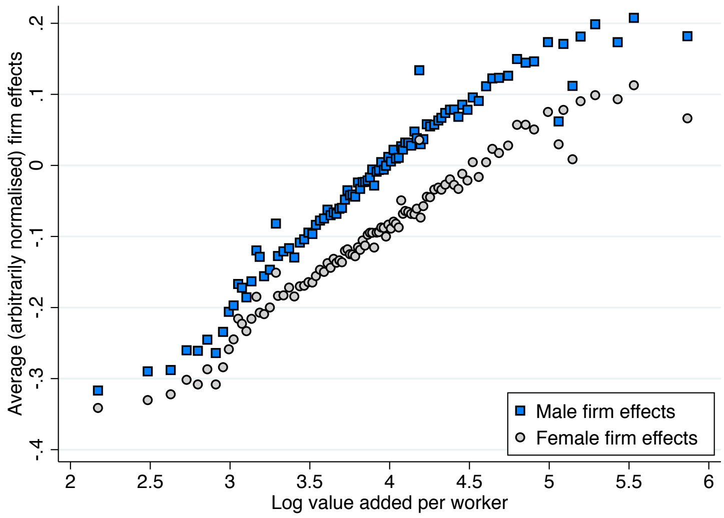

Table A.1 Gender pay gap, firm effects, sorting and bargaining with alternative normalisation

	(1) Log points	(2) Percent of the gender pay gap
Gender pay gap	0.213	
Male firm effect across males	0.246	
Female firm effect across females	0.166	
Firm effects gap	0.081	38.0%
Decomposition:		
Sorting		
Using male coefficients	0.049	22.8%
Using female coefficients	0.044	20.6%
Bargaining		
Using male distribution	0.037	17.4%
Using female distribution	0.032	15.2%
Observations	183,062,102	

Notes. The Table reports results of the Oaxaca-Blinder decomposition of equations (1.4) and (1.5). 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. Column (2) shows the ratio between the quantities reported in column (1) and the gender pay gap.

A.4 Decomposition by occupations and sectors

Occupations Table A.2 reports decompositions of the gender pay gap and firm effects gap 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 per cent of the gender pay gap for blue-collar workers (column 2) can be explained by firm components, mainly due to sorting of women into low-pay firms (roughly 31 per cent of the gender pay gap). A similar result holds for white-collar workers (column 3): the gap in firm effects accounts for 26 per cent of the gender gap in earnings, mainly due to sorting (18-21 per cent) rather than bargaining (5-8 per cent). Since the large majority of workers in our data is either classified as blue- or white-collar (roughly 91 per cent of men and 95 per cent 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 (see Table 1.4). For apprentices and middle managers (columns 1 and 4), 49 per cent and 19.5 per cent 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 per cent of the gender pay gap for apprentices and at least 22 per cent 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 5), 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.

Sectors We investigate the role of sectors in Figure A.6, 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).[9] 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, administration and health.[10] In finance and ICT, on the contrary, differences in bargaining power explain a larger share of the firm effects gap relative to sorting.

[9] We exclude sectors that comprise less than 1 per cent of the total person-year observations.

[10] In particular, it is likely that the results on manufacturing and trade sectors are behind the dominance of sorting in the overall dataset, since they employ more than 50 per cent of all person-year observations in the data.

Table A.2 Decomposition by occupation

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

Notes. The Table reports results of the Oaxaca-Blinder decomposition of equations (1.4) and (1.5). Firm effects are normalised with respect to the average gender-specific firm effects in the food and accommodation sector. Columns (1) to (5) 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.

Fig. A.6 Decomposition by sector

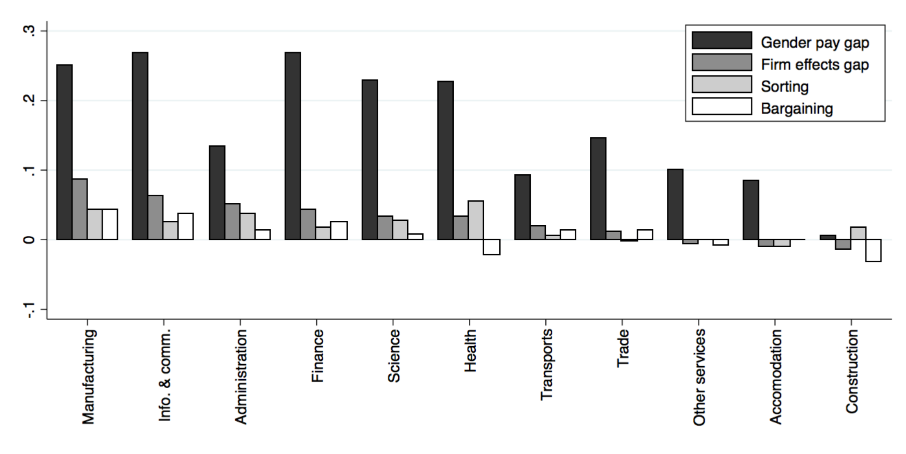

Notes. The Figure reports the gender pay gap, the firm effects gap, sorting and bargaining, by sector. Sectors are defined according to Nace Rev. 2 sectoral codes and ordered according to the gap in firm effects (highest to smallest). We exclude sectors that employ less than 1 per cent of the total person-year observations. The sectors reported in the figure represent 95 per cent of the total person-year observations between 1995 and 2015. Sorting and bargaining are average values, as in equation (1.6).

A.5 Additional Figures and Tables

Fig. A.7 Gender pay gap across the earnings distribution

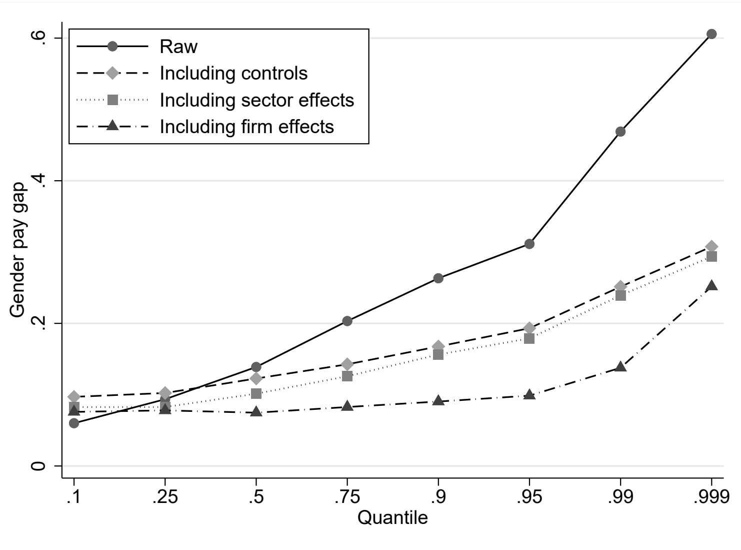

Notes. The Figure plots the coefficients on the male dummy in a quantile regression in four different specifications: without controls ("Raw"); controlling for observable characteristics of workers, i.e. cubic polynomials in age, experience and tenure, a dummy for full-time contract, the number of weeks 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.

Fig. A.8 Sorting and bargaining as a percentage of the gender pay gap in 1995-2000 and 2010-2015

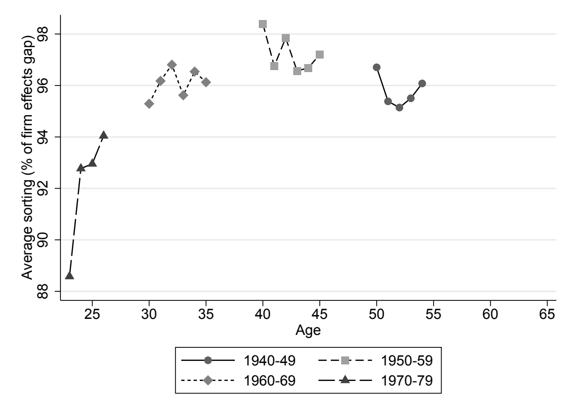

(A) Sorting 1995-2000

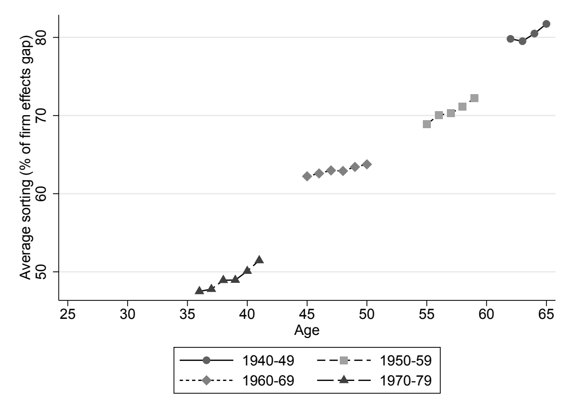

(B) Sorting 2010-2015

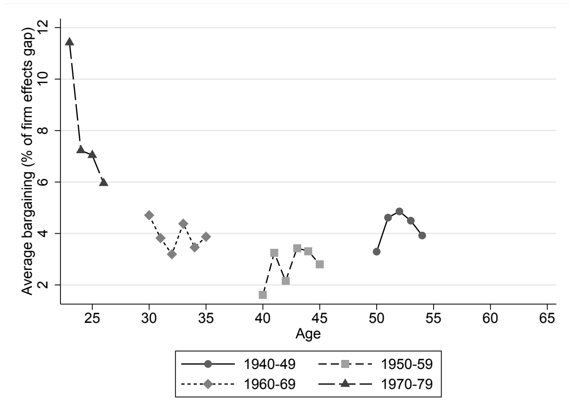

(C) Bargaining 1995-2000

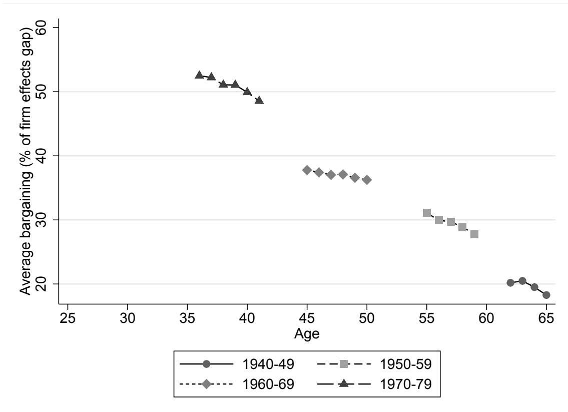

(D) Bargaining 2010-2015

Notes. The Figure plots sorting (panels A and B) and bargaining (panels C and D) as a percentage of the gender pay gap, averaged by age and cohort, from firm effects estimated over the periods 1995-2000 and 2010-2015. The horizontal axis reports age by cohort cells, defined as the mean age for each cohort across the period of time considered. Sorting and bargaining are average values, as in equation (1.6).

Fig. A.9 Mobility rates

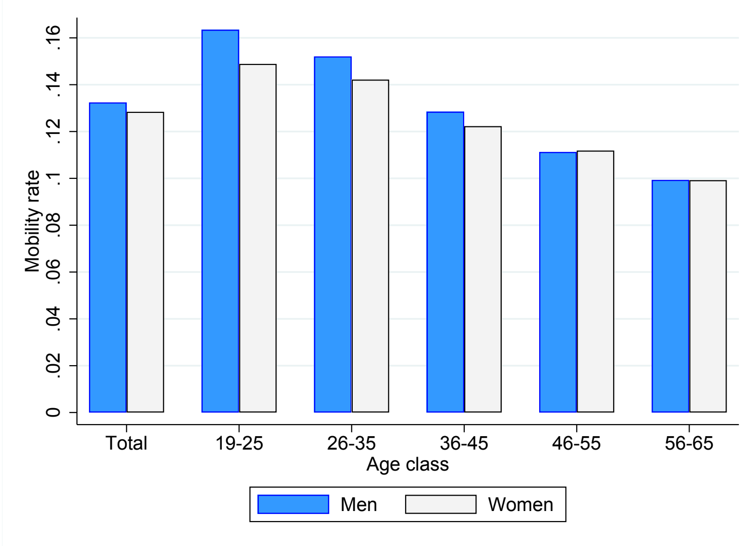

(A) Full sample

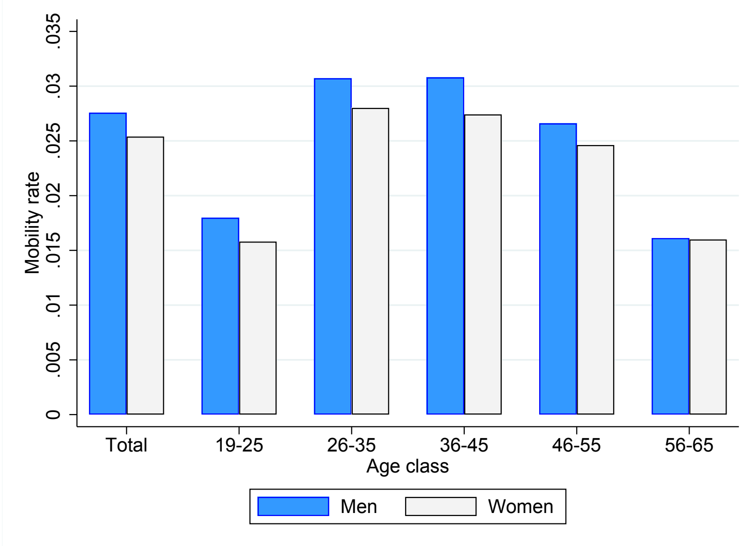

(B) Restricted sample

Notes. The Figure plots the mobility rate of men and women in the full (panel A) and restricted (panel B) sample of moves. The mobility rate is defined as the share of workers changing employer between two *consecutive* years. The full sample considers all moves. The restricted sample is used in section 1.5 and retains only moves such that the worker stays in the destination firm for at least two years after the move. All differences are statistically significant at 1 per cent level.

Fig. A.10 Gender-specific probabilities of moving to higher-quartile firms by age

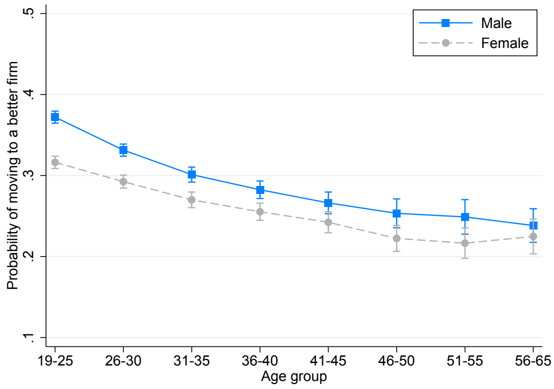

(A) All

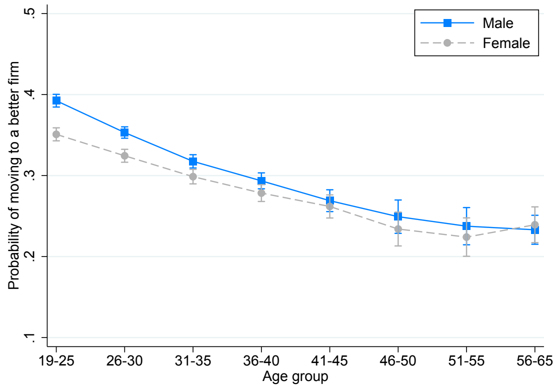

(B) Individual

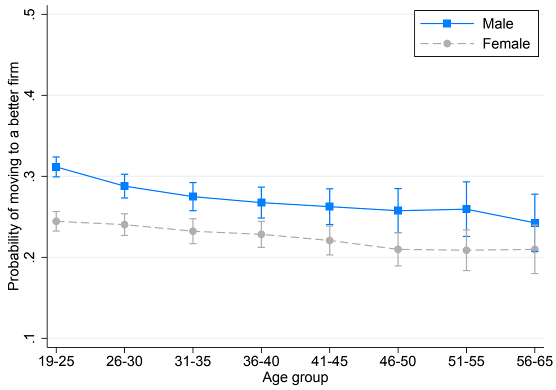

(C) Firm

Notes. The Figure plots the probabilities for men and women of moving to a firm in a higher firm effect quartile for different types of moves and age groups, conditional on the controls included in equation (1.9). Vertical lines are 95 per cent confidence intervals.

Fig. A.11 Average male and female firm effects in Italian provinces

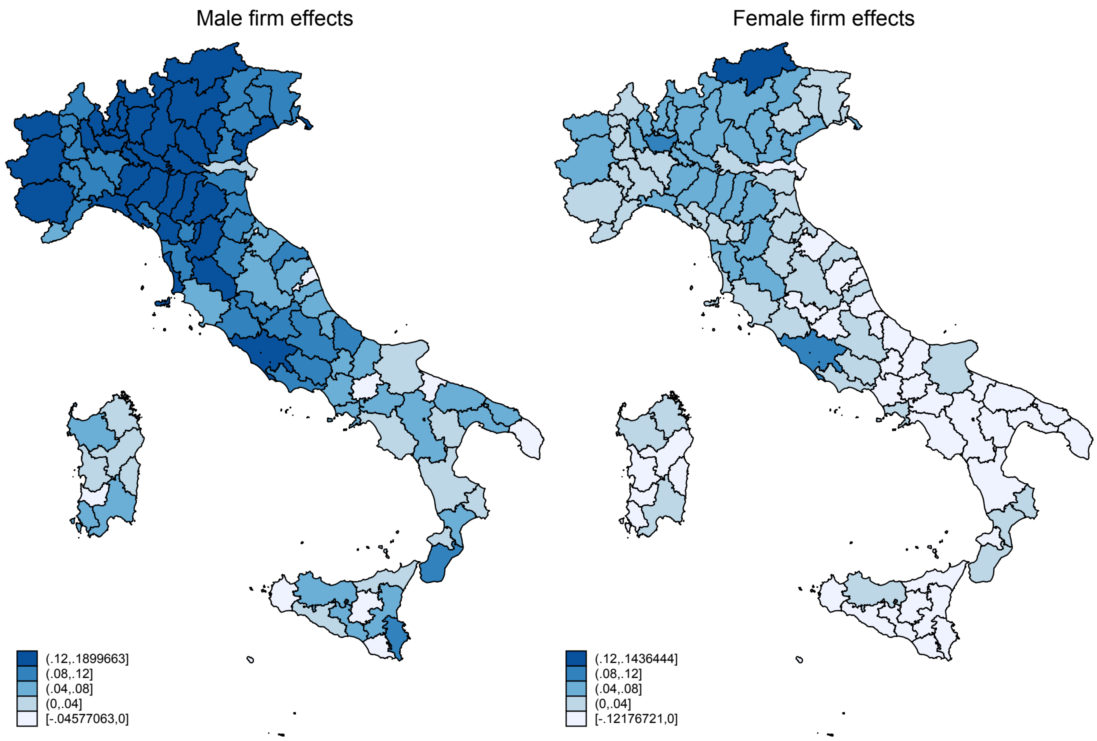

Notes. The Figure plots average male (left panel) and female (right panel) firm effects across Italian provinces. Darker colours indicate higher firm effects. Both graphs are drawn according to the same scale of firm effects.

Fig. A.12 Gender mobility gap within province

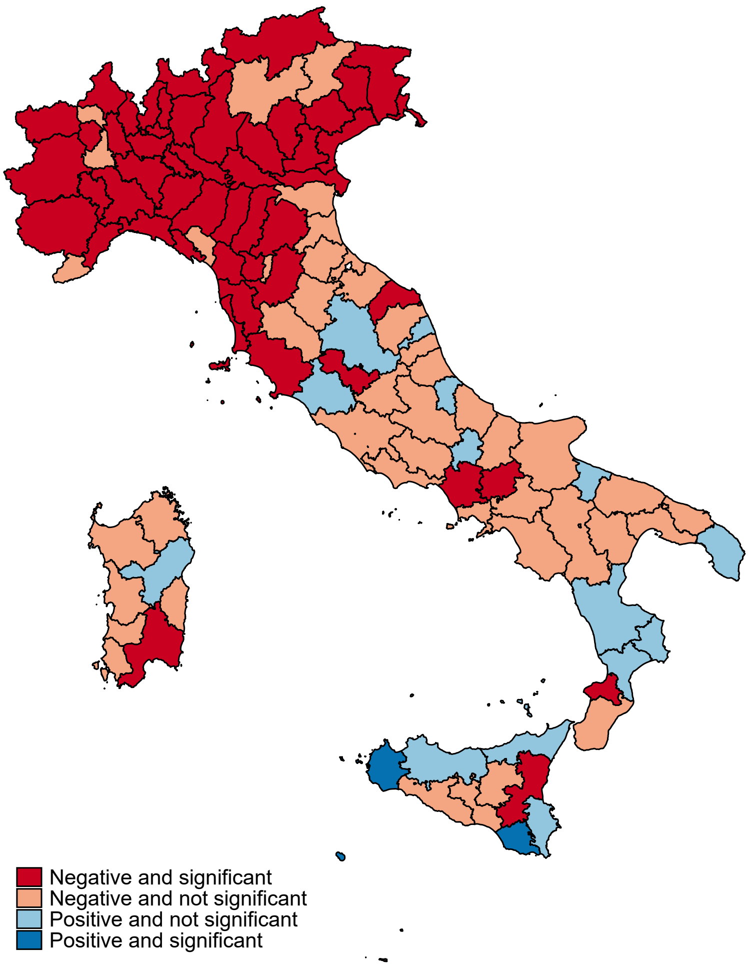

Notes. The Figure plots with different colours the marginal effect of the female dummy in a probit regression as in equation (1.9), estimated for each Italian province. Red (blue) areas denote provinces where the coefficient on the female dummy is negative (positive). Dark (light) areas indicate significant (not significant) coefficients at 95 per cent confidence level.

Fig. A.13 Female vs male firm effects

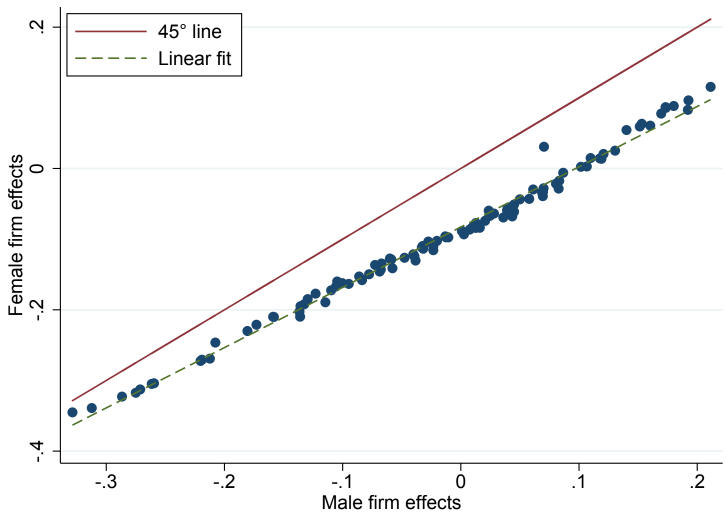

Notes. The Figure plots female against male firm effects averaged across percentile bins of log value added per worker. The slope of the linear fit is 0.85.

Table A.3 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 (1.4) and (1.5) 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 the food and accommodation sector in each period.

Table A.4 Probit model for job moves to a firm in same or higher fixed effect quartile

	(1) All	(2) Firm	(3) Individual
Female	-0.018***	-0.031***	-0.009**
	(0.004)	(0.007)	(0.004)
Age	0.001***	0.001***	0.000
	(0.000)	(0.000)	(0.000)
Change province	-0.033***	-0.027***	-0.029***
	(0.003)	(0.005)	(0.004)
Change occupation	-0.038***	-0.043***	-0.032***
	(0.004)	(0.006)	(0.004)
Change to full-time	-0.027***	-0.029***	-0.025**
	(0.009)	(0.004)	(0.013)
Observations	5,216,076	2,259,559	2,956,517
Sector FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Baseline Probability	0.809	0.839	0.787

Notes. The Table reports average marginal effects from probit regressions where the dependent variable is the probability of moving to a firm in the same or 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 happening because of firm closure. Column (3) shows results for moves not determined by firm closure. Robust standard errors, clustered at the firm level, in parentheses. Significance levels: * $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table A.5 Balance of covariates, before and after matching

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Unmatched Control	Matched Control	Treated	Unmatched Difference	P-value	Matched Difference	P-value
Open-ended contract	0.902	0.939	0.934	0.032	0.000	-0.005	0.614
Part-time	0.094	0.068	0.064	-0.030	0.000	-0.004	0.695
Part-time female	0.212	0.154	0.143	-0.069	0.000	-0.011	0.445
Female hiring rate	0.082	0.068	0.078	-0.004	0.641	0.010	0.442
Age group 35-54	0.645	0.669	0.657	0.011	0.239	-0.012	0.360
Age group 55+	0.102	0.094	0.100	-0.002	0.797	0.006	0.453
Log weekly earnings	6.307	6.588	6.637	0.329	0.000	0.049	0.253
Female log weekly earnings	6.227	6.450	6.473	0.246	0.000	0.023	0.547
Blue-collar	0.444	0.251	0.208	-0.236	0.000	-0.042	0.173
White-collar	0.471	0.532	0.537	0.066	0.000	0.005	0.832
Executives	0.023	0.084	0.103	0.080	0.000	0.018	0.348
Female executives	0.010	0.041	0.050	0.040	0.001	0.009	0.608
Middle managers	0.037	0.118	0.136	0.099	0.000	0.018	0.214
Log value added per worker	5.685	6.438	6.641	0.957	0.000	0.203	0.133
Log sales per worker	7.072	7.564	7.736	0.664	0.000	0.172	0.103
Log firm size	3.617	5.428	5.637	2.020	0.000	0.209	0.351
Industry	0.481	0.476	0.464	-0.017	0.669	-0.012	0.834
Construction	0.047	0.024	0.048	0.001	0.951	0.024	0.216
Trade, transports, accom.	0.276	0.145	0.108	-0.167	0.000	-0.036	0.330
Information & comm.	0.032	0.084	0.090	0.058	0.010	0.006	0.847
Finance & insurance	0.032	0.175	0.139	0.106	0.000	-0.036	0.390
Real estate	0.023	0.006	0.018	-0.005	0.613	0.012	0.249
Prof. and admin. services	0.072	0.078	0.090	0.018	0.421	0.012	0.695
Arts, entertainment & other	0.017	0.012	0.030	0.013	0.323	0.018	0.209
North	0.689	0.765	0.711	0.022	0.531	-0.054	0.298
Centre	0.182	0.169	0.229	0.047	0.154	0.060	0.197
South	0.153	0.072	0.084	-0.069	0.002	0.012	0.716
N. firms	16,040	154	166				

Notes. The Table reports the means of covariates in control and treated groups before and after matching. Column (1) to (3) report means for the unmatched control group, the matched control group and the treated group, respectively. Columns (4) and (5) report the difference in means and the relative p-value between control and treated group in the unmatched sample. Columns (6) and (7) report the same quantities for the matched sample. P-values are obtained from univariate regressions of each covariate on the treatment indicator.

Table A.6 Mean log earnings and frequencies of movers across firm effect quartiles

		Mean Log Real Weekly Earnings				4 Year Change	
Moves	Frequency	-2	-1	0	+1	Raw	Adjusted
Panel A: 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
Panel B: 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 female (panel A) and male (panel B) workers' moves between firm effect quartiles and the mean weekly earnings of the movers during the period between two years prior to the move and 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 $q \neq q'$.

Appendix B

Appendix for Chapter 2

B.1 Theoretical prediction by a principal-agent model

We motivate our empirical analysis by developing a simple two-period principal-agent model, based on Cornelissen et al. (2017). Peer effects affect the worker's future productivity and future wages through knowledge spillover or learning.

Basic Model

There are two periods in the model. In each period, a firm hires N workers. In the theoretical model, we do not consider the endogenous sorting of workers into firms, which our empirical analysis takes into account. Moreover, we assume that the same workers stay in the firm for both periods. Each worker i chooses their efforts to produce output and the firm can only observe the outputs to reward the worker with a wage contract. To simplify the model, we follow the literature and use a linear wage contract, which the firm chooses.

Worker's problem

Each worker is endowed with ability a, which is invariant over time.

Production function. The production function in each period is as follows.

$$q_1 = a + e_1 + L_1(e_1, \bar{a}_{-i}) + \varepsilon_1, \tag{B.1}$$

$$q_2 = a + e_2 + L_2(e_1, \bar{a}_{-i}) + \delta L_1(e_1, \bar{a}_{-i}) + \varepsilon_2, \tag{B.2}$$

where q_t, $t = \{1,2\}$, is the output observed to the firm, e_t is the effort and ε_t is a random productivity shock with zero mean. Both e_t and ε_t are unobserved to the firm. During each period, a worker could also learn from coworkers. Learning, denoted by L_t, depends on worker effort and peer quality. The latter is measured by the peers' average ability. Learning is not only concurrent – it also affects future productivity with a depreciation rate δ. For simplification, we assume that the learning function follows a straightforward expression below.

$$L_t(e_t, \bar{a}_{-i}) = e_t \lambda \bar{a}_{-i}.$$

If there is no learning within the firm, i.e., $\lambda = 0$, then the model reduces to a simplified two-period principal-agent model, as in Rogerson (1985).

Cost function. The effort will induce some disutility characterized by a standard quadratic cost function

$$c(e_t) = ke_t^2,$$

as it is increasingly costly for a worker to use additional effort.

Utility function. We make to simplifying assumptions: (i) the period utility of a worker is linear (risk-neutral); (ii) preferences are time-separable. The total utility across the two periods reads:

$$U = \sum_{t=1,2} [w(q_t) - c(e_t)],$$

where the wage contract $w(.)$ is determined by the firm, as detailed in the next paragraph.

Firm's problem

Assume the firm has a linear contract:

$$\begin{aligned} w_1 &= \alpha + \beta q_1, \\ w_2 &= \alpha + \beta q_2 + \theta q_1. \end{aligned}$$

The firm chooses β to reward the concurrently observed outputs from the worker, and choose θ to reward the first-period effort. In other words, the firm chooses the slopes β and θ to maximize profits, given the incentive constraint (IC) and individual rationality (IR) hold.

The optimal solution and its implications

Since we have not specified the cardinality of α, it is handy to manipulate α so that the IC holds, and we assume that the interior solution to the worker's maximization problem, i.e., IR holds:

$$\frac{\partial \mathbb{E}[w_2]}{\partial a_{-i}} = \beta^* \cdot \left(\frac{\partial q_2}{\partial e_2^*} \frac{\partial e_2^*}{\partial a_{-i}} + \frac{\partial q_2}{\partial e_1^*} \frac{\partial e_1^*}{\partial a_{-i}} \right) + \theta^* \cdot \frac{\partial q_1}{\partial e_1^*} \frac{\partial e_1^*}{\partial a_{-i}},$$

where all terms are positive. Therefore, $\frac{\partial \mathbb{E}[w_2]}{\partial a_{-i}}$ is positive. Thus, the model suggests that an increase in coworkers' quality could lead to wage growth in the future.[1]

B.2 Construction of C

Denote i for each row of the observation, where we suppress the notation of i, t to i. First, let's construct a matrix to indicate the location of coworker. For example, for each i,

$$c_i = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \\ \vdots \\ 1 \\ \vdots \end{pmatrix} - \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 1 \\ \vdots \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \\ \vdots \\ 0 \\ \vdots \end{pmatrix}$$

where i refers to the coworker's index location. The averaging matrix is constructed as follows.

$$\tilde{c}_i = \frac{c_i}{c_i' \cdot \vec{1}} \quad \Rightarrow \quad \tilde{C} = \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix} \quad \Rightarrow \quad C = \begin{pmatrix} \tilde{C} \\ 0 \end{pmatrix}$$

where th auxiliary matrix 0 makes the dimension of C the same as X.

A simple example

As a simple example on how we construct C, suppose we have the following data, where there are only five workers and two peer groups. The first column and second column of the data indicate the indices of worker and peer group, respectively.

$$\text{data} = \begin{pmatrix} 1 & 1 \\ 2 & 1 \\ 3 & 2 \\ 4 & 2 \\ 5 & 2 \end{pmatrix}$$

[1] The above analysis uses linearity as a simplifying assumption, which does not seems to be realistic. We might prove the finding by relaxing the linearity assumptions. Nevertheless, our numeric simulation finds the results still hold if we use non-linear wage contracts and utility functions.

We first construct an averaging matrix $\tilde{C}$ below to detect who is each worker's peer and what weight they are assigned when calculating the average peer quality. One might read $\tilde{C}$ as follows. The first row of $\tilde{C}$ says: 1 is not a coworker of himself, 2 is her coworker, and 3, 4, 5 are not her coworkers. The third row says, 1 and 2 are not 3's coworkers, 3 is not a coworker of herself, but 4 and 5 are her coworkers. Both of them weight half when calculating the average coworker quality.

$$\tilde{C} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 0.5 & 0 & 0.5 \\ 0 & 0 & 0.5 & 0.5 & 0 \end{pmatrix}$$

To make sure C and X have the same dimension, we add an auxiliary matrix 0 to $\tilde{C}$ as a final component of C. That is, $C = [\tilde{C}, \mathbf{0}]$.

B.3 Additional Figures and Tables

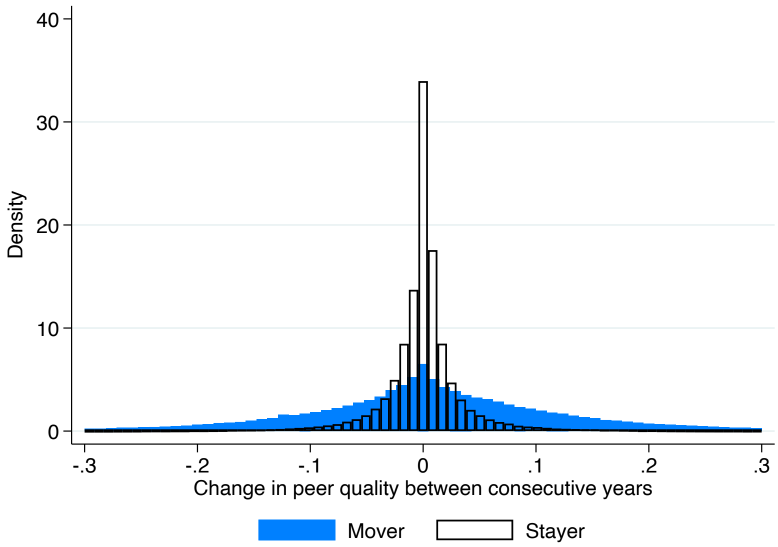

Fig. B.1 Density of the change in peer fixed effect between consecutive years for movers and stayers.

Notes. The figure shows the density of changes in peer quality between consecutive years for movers and stayers. The peer quality is measured as the leave-one out average of AKM worker effects at the peer group level, i.e. $\bar{\alpha}_{-i,t} = \frac{1}{|M_{-it}|}\sum_{k \in M_{-it}} \alpha_k$, where $|M_{-it}|$ is the number of peers of worker i.

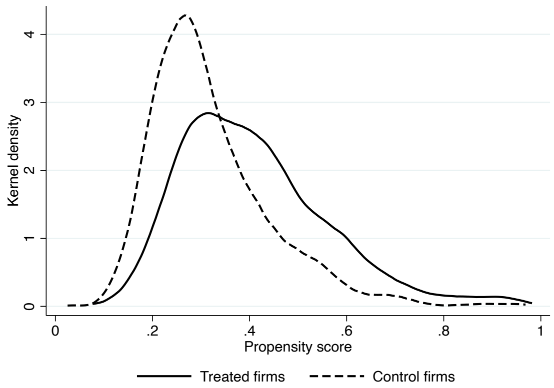

(A) Treatment 1 vs Control group

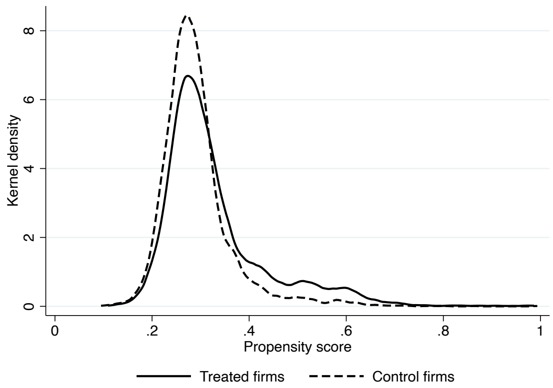

(B) Treatment 2 vs Control group

Fig. B.2 Propensity score density, hire design

Notes. The figure reports the propensity score density for firms hiring high-quality (treatment 1) and average-quality (control) workers in panel (A) and for firms hiring low-quality (treatment 1) and average-quality (control) workers in panel (B).

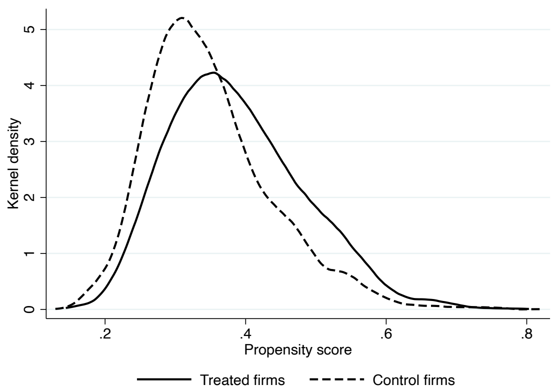

(A) Treatment 1 vs Control group

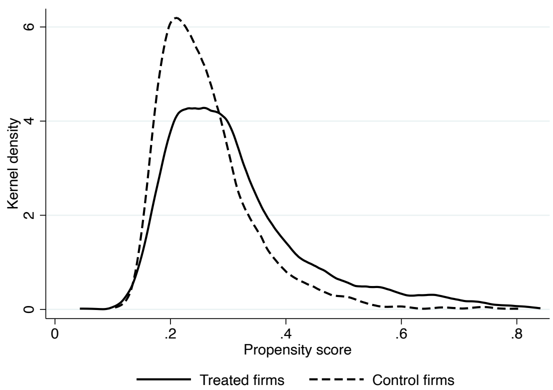

(B) Treatment 2 vs Control group

Fig. B.3 Propensity score density, leaver design

Notes. The figure reports the propensity score density for firms separating from high-quality (treatment 1) and average-quality (control) workers in panel (A) and for firms hiring low-quality (treatment 1) and average-quality (control) workers in panel (B).

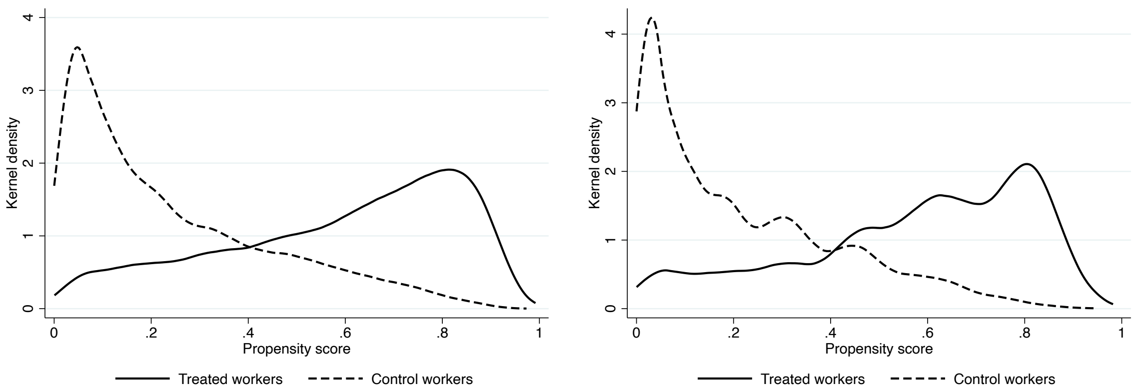

(A) Treatment 1 vs Control group

(B) Treatment 2 vs Control group

Fig. B.4 Propensity score density, mover design

Notes. The figure reports the propensity score density for workers moving into high-quality (treatment 1) and average-quality (control) peer groups in panel (A) and for workers moving into low-quality (treatment 1) and average-quality (control) peer groups in panel (B).

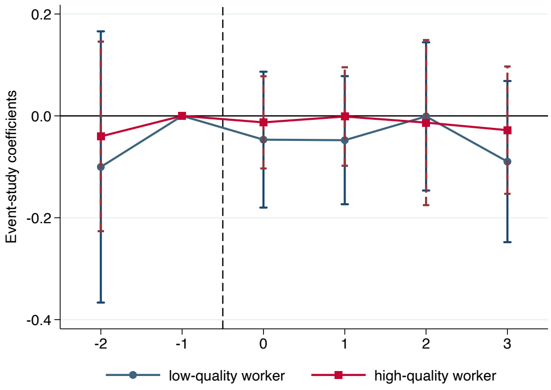

(A) Sales per worker

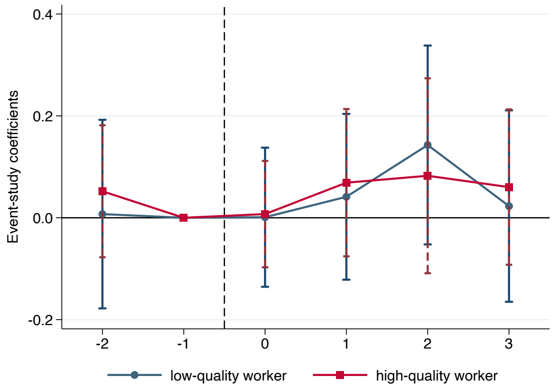

(B) Value added per worker

Fig. B.5 Log value added per worker and log sales per worker around mobility, hire design

Notes. The figure reports estimates of β_k from equation (2.6), separately for firms hiring high-quality and low-quality workers relative to firms hiring average-quality workers. The dependent variable is log sales per worker in panel (A) and log value added per worker in panel (B). Regressions are weighted by firm size. Vertical bars are 95% confidence intervals, obtained from cluster-robust standard errors at the firm level.

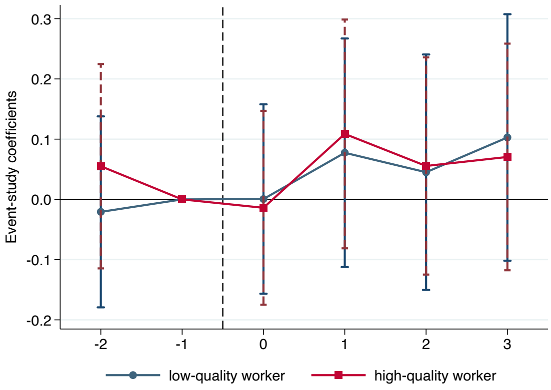

(A) Sales per worker

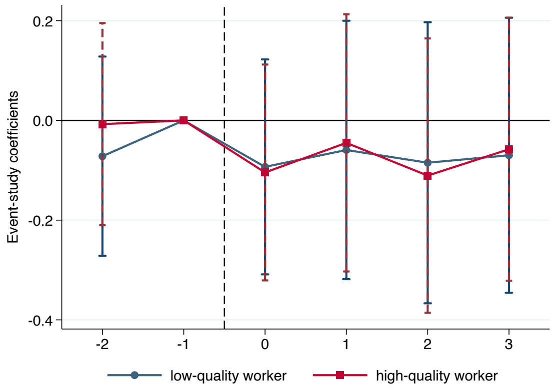

(B) Value added per worker

Fig. B.6 Log value added per worker and log sales per worker around mobility, leaver design

Notes. The figure reports estimates of β_k from equation (2.6), separately for firms separating from high-quality and low-quality workers relative to firms separating from average-quality workers. The dependent variable is log sales per worker in panel (A) and log value added per worker in panel (B). Regressions are weighted by firm size. Vertical bars are 95% confidence intervals, obtained from cluster-robust standard errors at the firm level.

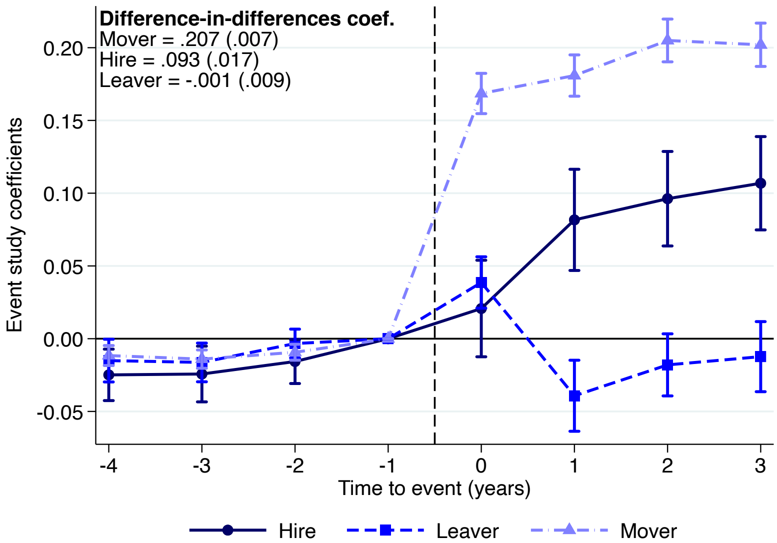

Fig. B.7 Event study, continuous treatment

Notes. The figure reports estimates of $\widetilde{\beta}_k$, i.e. the dynamic effects of a 1 percent increase in the quality of a new hire or a separation on coworkers' wages in the origin (*Hire*) and destination firms (*Leave*) from equation (2.8), and the estimates of $\widetilde{\gamma}_k$, i.e. the dynamic effects of a 1 percent increase in the quality of the peer group a mover joins (*Mover*), from equation (2.9). Vertical bars are 95 percent confidence intervals, obtained from cluster-robust standard errors at the firm (for $\widetilde{\beta}_k$) and individual (for $\widetilde{\gamma}_k$) level.

Table B.1 Balance test of covariates, before and after matching, hiring design

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	High-quality mover				Low-quality mover			
	Unmatched		Matched		Unmatched		Matched	
	Diff.	P-value	Diff.	P-value	Diff.	P-value	Diff.	P-value
Mean wage -4	-21.574	0.000***	-5.932	0.141	7.136	0.124	1.291	0.785
Mean wage -3	-21.511	0.000***	-5.905	0.156	7.951	0.093*	4.570	0.342
Mean wage -2	-21.050	0.000***	-5.580	0.165	10.742	0.026**	4.523	0.356
AKM worker effect	-0.088	0.000***	-0.006	0.178	-0.000	0.955	0.006	0.273
AKM firm effect	-0.001	0.721	-0.006	0.155	-0.025	0.000***	-0.001	0.791
Employees' mean age	3.512	0.000***	-0.177	0.399	0.659	0.000***	-0.165	0.456
Share female	0.006	0.534	0.003	0.811	0.028	0.007***	-0.011	0.383
Share blue-collar	0.014	0.145	-0.007	0.513	-0.049	0.000***	0.011	0.400
Firm size	0.293	0.136	-0.069	0.770	0.236	0.233	0.106	0.642
Firm age	0.492	0.000***	-0.107	0.512	0.280	0.047**	-0.021	0.905
Value added	332.034	0.645	-425.186	0.631	-1.2e+03	0.138	-75.134	0.938
Revenues	94.932	0.895	-420.232	0.634	517.826	0.501	-11.555	0.990
Manufacturing	-0.104	0.000***	0.012	0.425	-0.041	0.002***	0.018	0.263
Construction	0.055	0.000***	-0.006	0.548	0.002	0.823	0.008	0.382
Wholesale	-0.020	0.030**	0.001	0.904	-0.038	0.000***	-0.001	0.927
Accommodation	0.018	0.000***	-0.004	0.519	0.008	0.074*	-0.004	0.462
Transports	0.001	0.771	-0.003	0.574	-0.007	0.125	-0.003	0.630
Finance	-0.012	0.000***	0.001	0.733	0.001	0.809	-0.003	0.633
Services	-0.010	0.058*	-0.002	0.783	-0.010	0.100*	-0.005	0.459
Health	0.038	0.000***	0.010	0.062*	0.047	0.000***	-0.009	0.171
Domestic	0.003	0.289	-0.000	0.897	0.005	0.128	0.001	0.784
Other	0.031	0.000***	-0.010	0.237	0.032	0.000***	-0.003	0.751
Belluno	0.019	0.001***	0.002	0.746	0.001	0.882	-0.004	0.590
Padua	-0.023	0.016**	0.005	0.692	-0.001	0.934	0.006	0.617
Rovigo	-0.002	0.722	0.000	0.940	-0.006	0.275	0.005	0.389
Treviso	-0.003	0.766	0.004	0.723	-0.010	0.335	0.006	0.609
Venice	-0.004	0.686	-0.012	0.254	0.004	0.657	-0.015	0.212
Vicenza	-0.008	0.433	0.003	0.820	-0.003	0.784	-0.002	0.903
Verona	0.020	0.036**	-0.002	0.875	0.015	0.152	0.002	0.861
N. treated	2517		2164		2015		1848	
N. control	4636		2164		4636		1848	

Notes. The table reports a balance test of covariates used for matching in the analysis of the effect of a new hire on coworkers' wages. Columns (1) to (4) report the average difference and the p-value of the difference for each variable in the unmatched (columns 1-2) and matched (column 3-4) samples comparing firms hiring a high-quality worker to those hiring an average-quality worker. Columns (5) to (8) report the average difference and the p-value of the difference for each variable in the unmatched (columns 5-6) and matched (column 7-8) samples comparing firms hiring a low-quality worker to those hiring an average-quality worker. Heteroskedasticity robust p-values are obtained from univariate regressions of each covariate on a dummy for firms hiring high-quality workers (in columns 2 and 4) or low-quality workers (in columns 6 and 8). Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.2 Balance test of covariates, before and after matching, leaver design

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	High-quality mover				Low-quality mover			
	Unmatched		Matched		Unmatched		Matched	
	Diff.	P-value	Diff.	P-value	Diff.	P-value	Diff.	P-value
Mean wage -4	-2.060	0.509	-4.895	0.163	-1.729	0.676	-0.362	0.940
Mean wage -3	-1.298	0.682	-3.758	0.292	-0.208	0.960	-1.934	0.688
Mean wage -2	0.039	0.990	-4.018	0.256	4.700	0.285	-1.859	0.693
AKM worker effect	-0.036	0.000***	-0.004	0.327	-0.036	0.000***	-0.002	0.708
AKM firm effect	0.000	0.940	-0.004	0.264	-0.008	0.061*	-0.002	0.735
Employees' mean age	2.109	0.000***	-0.040	0.811	1.972	0.000***	0.202	0.352
Share female	-0.006	0.470	0.009	0.365	0.048	0.000***	-0.017	0.183
Share blue-collar	0.005	0.539	-0.002	0.840	-0.093	0.000***	0.015	0.260
Firm size	0.642	0.000***	-0.046	0.828	-0.149	0.519	-0.385	0.087*
Firm age	0.124	0.309	-0.072	0.616	0.006	0.965	-0.037	0.836
Value added	39.360	0.952	782.142	0.303	709.148	0.354	209.679	0.825
Revenues	-290.164	0.657	176.462	0.816	363.910	0.627	177.693	0.851
Manufacturing	-0.057	0.000***	-0.002	0.895	-0.105	0.000***	0.014	0.392
Construction	0.045	0.000***	0.003	0.682	0.013	0.091*	0.014	0.139
Wholesale	-0.022	0.007***	-0.001	0.915	-0.027	0.004***	-0.006	0.596
Accommodation	0.017	0.000***	-0.002	0.701	0.003	0.528	0.001	0.860
Transports	0.003	0.462	-0.003	0.557	-0.001	0.796	0.002	0.762
Finance	-0.008	0.016**	0.004	0.250	0.003	0.502	-0.007	0.237
Services	-0.025	0.000***	-0.001	0.880	0.004	0.504	-0.007	0.399
Health	0.017	0.000***	0.005	0.231	0.055	0.000***	-0.003	0.676
Domestic	0.006	0.037**	-0.001	0.846	0.009	0.014**	-0.002	0.634
Other	0.023	0.000***	-0.003	0.628	0.046	0.000***	-0.006	0.535
Belluno	0.013	0.009***	-0.003	0.644	0.010	0.079*	-0.002	0.827
Padua	-0.023	0.008***	-0.001	0.892	-0.018	0.070*	-0.002	0.868
Rovigo	0.006	0.200	0.000	0.949	0.004	0.439	0.005	0.411
Treviso	-0.019	0.029**	0.011	0.250	-0.004	0.708	0.008	0.526
Venice	0.020	0.015**	0.001	0.915	0.002	0.808	-0.001	0.964
Vicenza	-0.003	0.785	-0.008	0.447	-0.008	0.429	0.010	0.467
Verona	0.006	0.489	-0.000	0.972	0.014	0.148	-0.019	0.139
N. treated	3065		2905		2046		1885	
N. control	5374		2905		5374		1885	

Notes. The table reports a balance test of covariates used for matching in the analysis of the effect of a separation on coworkers' wages. Columns (1) to (4) report the average difference and the p-value of the difference for each variable in the unmatched (columns 1-2) and matched (column 3-4) samples comparing firms separating from a high-quality worker to those separating from an average-quality worker. Columns (5) to (8) report the average difference and the p-value of the difference for each variable in the unmatched (columns 5-6) and matched (column 7-8) samples comparing firms separating from a low-quality worker to those separating from an average-quality worker. Heteroskedasticity robust p-values are obtained from univariate regressions of each covariate on a dummy for firms separating from high-quality workers (in columns 2 and 4) or low-quality workers (in columns 6 and 8). Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.3 Balance test of covariates, before and after matching, mover design

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	High-quality peers				Low-quality peers			
	Unmatched		Matched		Unmatched		Matched	
	Diff.	P-value	Diff.	P-value	Diff.	P-value	Diff.	P-value
Wage -4	-0.070	0.000***	0.007	0.513	0.201	0.000***	-0.009	0.658
Wage -3	-0.076	0.000***	0.008	0.490	0.216	0.000***	-0.011	0.578
Wage -2	-0.081	0.000***	0.007	0.567	0.229	0.000***	-0.009	0.681
Weeks worked	-0.420	0.004***	0.105	0.321	0.196	0.017**	0.085	0.618
AKM worker effect decile	-2.682	0.000***	-0.033	0.712	2.881	0.000***	0.094	0.279
Age	6.504	0.000***	-0.303	0.425	0.360	0.407	-0.523	0.068*
Female	0.091	0.000***	0.002	0.957	-0.125	0.000***	-0.023	0.182
Tenure	-0.114	0.038**	0.091	0.074*	-0.012	0.760	0.070	0.328
Blue-collar	0.049	0.001***	0.008	0.658	-0.261	0.000***	0.031	0.209
AKM firm effect	-0.005	0.341	0.002	0.680	-0.003	0.477	0.007	0.231
Log firm size	0.260	0.002***	0.002	0.973	0.319	0.000***	0.033	0.759
Manufacturing	0.015	0.358	0.001	0.964	-0.027	0.052*	-0.019	0.490
Construction	0.004	0.361	0.001	0.867	0.005	0.238	-0.002	0.800
Wholesale	-0.007	0.528	-0.003	0.790	-0.004	0.539	0.012	0.410
Accommodation	-0.000	0.683	0.001	0.259	0.001	0.216	-0.001	0.443
Transports	-0.002	0.706	0.001	0.818	-0.005	0.175	-0.002	0.455
Finance	-0.018	0.235	0.004	0.649	0.010	0.197	0.018	0.161
Services	-0.002	0.605	0.000	0.836	0.002	0.533	-0.002	0.784
Health	0.003	0.267	-0.001	0.409	0.008	0.232	-0.002	0.680
Domestic	0.003	0.148	-0.001	0.286	0.003	0.408	-0.001	0.839
Other	0.003	0.411	-0.003	0.745	0.009	0.147	-0.001	0.945
Belluno	0.001	0.667	0.002	0.512	0.005	0.085*	0.001	0.668
Padua	-0.009	0.429	-0.001	0.963	-0.011	0.078*	0.011	0.571
Rovigo	-0.003	0.408	0.000	0.979	-0.009	0.031**	-0.003	0.494
Treviso	0.000	0.995	0.001	0.936	0.004	0.501	-0.006	0.641
Venice	-0.008	0.510	0.001	0.956	-0.005	0.595	-0.000	0.986
Vicenza	0.018	0.073*	-0.004	0.756	0.019	0.078*	-0.001	0.974
Verona	0.001	0.830	0.001	0.879	-0.002	0.704	-0.003	0.757
N. treated	26194		15511		22547		12778	
N. control	46007		15511		46007		12778	

Notes. The table reports a balance test of covariates used for matching in the analysis of the effect on worker's wages of moving into peer groups of different quality. Columns (1) to (4) report the average difference and the p-value of the difference for each variable in the unmatched (columns 1-2) and matched (column 3-4) samples comparing workers moving into high-quality peers to those moving into average-quality peers. Columns (5) to (8) report the average difference and the p-value of the difference for each variable in the unmatched (columns 5-6) and matched (column 7-8) samples comparing workers moving into low-quality peers to those moving into average-quality peers. Heteroskedasticity robust p-values are obtained from univariate regressions of each covariate on a dummy for workers moving into high-quality peers (in columns 2 and 4) or low-quality peers (in columns 6 and 8). Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Appendix C

Appendix for Chapter 3

C.1 Additional Figures and Tables

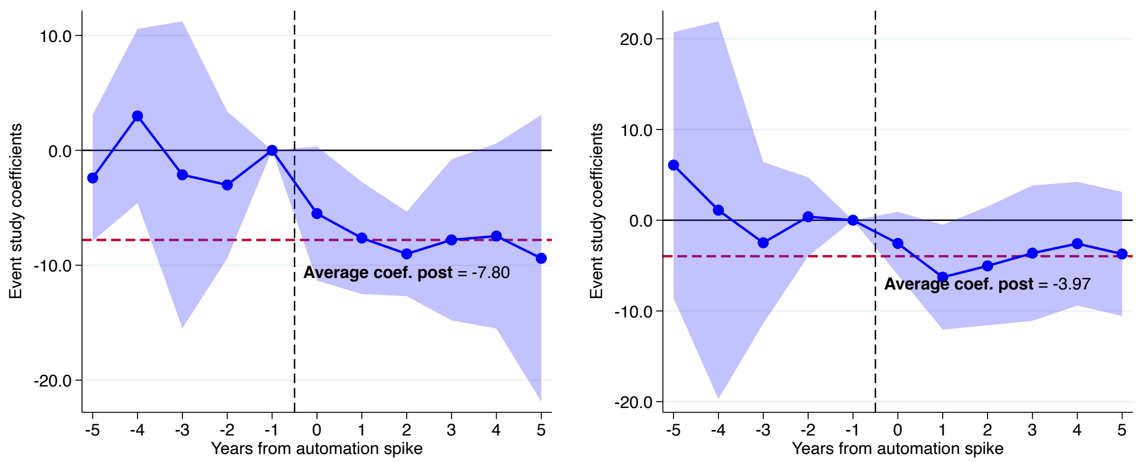

(A) Correlation worker effect and value added

(B) Correlation worker effect and sales

Fig. C.1 Event study of the effect of automation spike on assortative matching in the labour market, de Chaisemartin and D'Haultfoeuille (2020) methodology

Notes. The figure reports event study estimates, following the methodology proposed by de Chaisemartin and D'Haultfoeuille (2020), of the effects of automation spikes on labour market sorting, measured as the correlation between AKM worker fixed effects and log value added per worker in panel A and with log sales per worker in panel B. The regressions include non-parametric controls, measured in the year prior to the automation spike, for average log sales per worker, average log earnings, the share of blue-collar workers, the share of female workers, the share of workers below age 40 and the share of firms with more than 100 employees. The shaded areas are 95% confidence intervals from bootstrapped standard errors.

Table C.1 Sectors included in IFR data

Sector	Nace	Group
Agriculture, forestry, fishing	A-B	1
Mining and quarrying	C	2
Food and beverages	10, 11, 12	3
Textiles	13, 14, 15	4
Wood and furniture	16	5
Paper	17, 18	6
Pharmaceuticals, cosmetics	19	7
Other chemical products n.e.c.	20, 21	8
Rubber and plastic products (non-automotive)	22	9
Chemical products, unspecified	229	9
Glass, ceramics, stone, mineral products (non-automotive)	23	10
Basic metals	24	11
Metal products (non-automotive)	25	12
Electronic components/devices	260	13
Semiconductors, LCD, LED	261	13
Computers and peripheral equipment	262	13
Info communication equipment, domestic and prof.	263	13
Medical, precision, optical instruments	265	13
Household/domestic appliances	275	14
Electrical machinery n.e.c. (non-automotive)	271	14
Electrical/electronics unspecified	279	14
Industrial machinery	28	15
Metal, unspecified	289	15
Automotive	29	16
Other vehicles	30	17
All other manufacturing branches	91	18
Electricity, gas, water supply	E	19

Notes. The table reports the sectors included in IFR data, together with their classification in Nace Rev. 2 and how the sectors have been grouped for the purpose of the present paper.

Table C.2 First stage regressions

	(1)	(2)	(3)	(4)	(5)
Panel A: High labour force attachment group					
$\triangle$ robots Germany	-0.288	-0.289***	-0.289***	-0.276***	-0.358***
	(0.044)	(0.043)	(0.043)	(0.088)	(0.101)
$\triangle$ robots Spain	-0.156	-0.146	-0.143	0.512	0.341
	(0.560)	(0.558)	(0.556)	(0.680)	(0.570)
$\triangle$ robots Finland	-0.575	-0.511	-0.444	5.291	1.731
	(3.800)	(3.795)	(3.790)	(4.450)	(3.573)
$\triangle$ robots France	2.100	2.112***	2.113***	1.393***	1.392***
	(0.492)	(0.490)	(0.488)	(0.406)	(0.409)
$\triangle$ robots Norway	-3.066	-3.163	-3.213	11.783	5.542
	(6.030)	(6.037)	(6.013)	(7.959)	(6.798)
$\triangle$ robots Sweden	4.393	4.374***	4.347***	1.471	4.069***
	(1.436)	(1.431)	(1.428)	(1.750)	(1.346)
$\triangle$ robots United Kingdom	0.983	0.962	0.959	0.627	1.378**
	(0.745)	(0.745)	(0.744)	(0.649)	(0.675)
F-statistic of excluded instruments	140.5	135.2	135.7	138.1	114.3
Observations	136,736	136,736	136,736	136,736	136,736
Panel B: All workers					
$\triangle$ robots Germany	-0.288**	-0.288***	-0.288***	-0.277***	-0.364***
	(0.044)	(0.044)	(0.044)	(0.088)	(0.104)
$\triangle$ robots Spain	-0.159**	-0.147	-0.149	0.545	0.372
	(0.579)	(0.576)	(0.575)	(0.701)	(0.590)
$\triangle$ robots Finland	-0.379**	-0.320	-0.289	5.658	2.009
	(3.990)	(3.982)	(3.975)	(4.737)	(3.827)
$\triangle$ robots France	2.120**	2.133***	2.134***	1.399***	1.385***
	(0.505)	(0.502)	(0.501)	(0.410)	(0.420)
$\triangle$ robots Norway	-3.559**	-3.680	-3.754	11.625	5.289
	(5.942)	(5.964)	(5.944)	(7.894)	(6.692)
$\triangle$ robots Sweden	4.355**	4.337***	4.323***	1.341	4.040***
	(1.500)	(1.493)	(1.490)	(1.840)	(1.413)
$\triangle$ robots United Kingdom	0.965**	0.942	0.942	0.588	1.392*
	(0.787)	(0.789)	(0.787)	(0.689)	(0.710)
F-statistic of excluded instruments	128.6	121.0	121.7	127.9	111.1
Observations	174,405	174,405	174,405	174,405	174,405
Birth year dummies	Yes	Yes	Yes	Yes	Yes
Demographic controls		Yes	Yes	Yes	Yes
Employment history			Yes	Yes	Yes
Industry and firm controls				Yes	Yes
$\triangle$ Net exports and $\triangle$ ICT					Yes

Notes. The table reports the estimates from the first stage regression of the change in robots per 1000 worker in the Italian labour market on the corresponding change in other seven European countries between 1994 and 2018. Panel A focuses on the main high labour force attachment sample used in the analysis. Panel B shows results for the full sample of workers. Column (1) includes birth year dummies. Column (2) controls for workers' gender and region of residence (20 dummies). Column (3) includes tenure dummies (4-6 years, $\geq$ 7 years), log earnings in 1994 and terciles of adjusted wages. Column (4) includes five firm size dummies (11-20, 21-50, 51-100, 101-500, 500 or more employees) and four dummies for broad manufacturing sectors (food, beverages and tobacco; consumer goods, industrial goods and capital goods). Column (5) further controls for the change in net exports form China and Eastern Europe over 1994-2018 and the change in ICT capital over 1996-2016. The table further reports the Kleibergen and Paap (2006) first stage F-statistic of excluded instruments. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C.3 Exposure to robots and earnings and employment by industry, firm and firm size, 1993-2018: 2SLS estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	all employers	manufacturing							non-manufacturing		
Same 2-digit industry		Yes	Yes	Yes	Yes	No	No	No	No	No	No
Same firm		Yes	No	No	No	No	No	No	No	No	No
Larger firm		-	-	Yes	No	-	Yes	No	-	Yes	No
Panel A: Cumulative earnings											
△ robots per 1000 worker	2.302	7.390***	-1.968	-0.958	-1.010**	-1.069	-1.081	0.012	-2.051**	-0.995	-1.055***
	(1.679)	(1.094)	(1.241)	(0.941)	(0.396)	(1.780)	(1.470)	(0.384)	(1.034)	(0.679)	(0.390)
Panel B: Cumulative days											
△ robots per 1000 worker	5.631***	16.049***	-5.036**	-2.752	-2.285***	-1.319	-1.123	-0.196	-4.062**	-1.774*	-2.288***
	(1.264)	(2.930)	(2.509)	(1.975)	(0.760)	(3.372)	(2.799)	(0.712)	(1.570)	(1.022)	(0.661)
Panel C: Average daily wage											
△ robots per 1000 worker	0.008	0.058	-0.115	-0.065	-0.114**	-0.027	-0.019	-0.003	-0.166**	-0.085*	-0.131***
	(0.065)	(0.039)	(0.081)	(0.088)	(0.049)	(0.119)	(0.116)	(0.038)	(0.065)	(0.050)	(0.043)

Notes. $N = 136{,}736$. The table reports estimates of β from equation (3.1). Outcomes are cumulated values over 1994-2018. Panels A, B and C shows results for cumulative earnings, cumulative days and average daily wages, respectively. Column (1) cumulates outcomes over all employers. Column (2) cumulates outcomes only at origin employer in the baseline year 1994. Column (3) cumulates outcomes only at different employers in the same 2-digit industry. Columns (4) and (5) decompose this effect for workers moving to smaller and larger firms, respectively. Column (6) cumulates outcomes at different employers in different 2-digit industries. Columns (7) and (8) decompose this effect for workers moving to smaller and larger firms, respectively. Column (9) cumulates outcomes at employers outside of manufacturing. Columns (10) and (11) decompose this effect for workers moving to smaller and larger firms, respectively. All regressions include the full set of controls from column (6) of Table 3.2. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C.4 Exposure to robots and earnings and employment by industry, firm, 1993-2018: 2SLS estimates, sample with non-missing balance sheet information

	(1) all employers	(2) manufacturing	(3) manufacturing	(4) manufacturing	(5) non-manuf.
Same 2-digit industry		Yes	Yes	No	No
Same firm		Yes	No	No	No
Panel A: Cumulative earnings					
△ robots per 1000 worker	2.910*	5.875***	-1.822	0.161	-1.305
	(1.515)	(1.295)	(1.133)	(2.247)	(0.955)
Panel B: Cumulative days					
△ robots per 1000 worker	5.906***	12.526***	-4.394**	0.972	-3.198*
	(1.530)	(3.523)	(2.184)	(4.170)	(1.698)
Panel C: Average daily wage					
△ robots per 1000 worker	0.032	0.042*	-0.087	0.036	-0.109
	(0.052)	(0.024)	(0.085)	(0.151)	(0.081)

Notes. $N = 45,766$. The table replicates Table 3.4 using observations with non-missing information on firms' value added. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C.5 First stage firms, 1994-2012

	(1)	(2)	(3)	(4)	(5)
Panel A: Restricted sample					
△ robots Germany	-0.005*	-0.004	-0.004	0.219*	0.194
	(0.046)	(0.045)	(0.045)	(0.126)	(0.122)
△ robots Spain	-2.589*	-2.596***	-2.592***	-3.308***	-3.148***
	(0.363)	(0.362)	(0.363)	(0.397)	(0.409)
△ robots Finland	-0.529*	-0.511	-0.492	-1.670	-2.987*
	(0.960)	(0.951)	(0.960)	(1.525)	(1.469)
△ robots France	4.227*	4.237***	4.235***	4.247***	4.037***
	(0.494)	(0.490)	(0.490)	(0.533)	(0.592)
△ robots Norway	-13.214*	-13.434**	-13.482**	-17.239***	-23.570**
	(5.529)	(5.478)	(5.478)	(5.988)	(9.137)
△ robots Sweden	4.457*	4.457***	4.458***	3.916***	5.068***
	(0.469)	(0.465)	(0.466)	(0.546)	(1.118)
△ robots United Kingdom	-1.685*	-1.702**	-1.703**	-2.610***	-2.397***
	(0.713)	(0.708)	(0.708)	(0.613)	(0.666)
F-statistic of excluded instruments	171.7	173.4	171.5	293.5	378.0
Observations	19,802	19,802	19,802	19,802	19,802
Panel B: Full sample					
△ robots Germany	-0.006**	-0.005	-0.005	0.221*	0.193
	(0.045)	(0.045)	(0.045)	(0.125)	(0.124)
△ robots Spain	-2.563**	-2.570***	-2.565***	-3.305***	-3.143***
	(0.377)	(0.375)	(0.377)	(0.406)	(0.422)
△ robots Finland	-0.492**	-0.470	-0.448	-1.715	-3.009**
	(0.943)	(0.933)	(0.940)	(1.522)	(1.459)
△ robots France	4.199**	4.211***	4.209***	4.225***	4.027***
	(0.509)	(0.505)	(0.505)	(0.531)	(0.591)
△ robots Norway	-13.329**	-13.573**	-13.625**	-17.144***	-23.386**
	(5.445)	(5.390)	(5.383)	(5.894)	(9.227)
△ robots Sweden	4.425**	4.427***	4.427***	3.894***	5.041***
	(0.470)	(0.466)	(0.467)	(0.537)	(1.145)
△ robots United Kingdom	-1.657**	-1.678**	-1.679**	-2.585***	-2.366***
	(0.705)	(0.699)	(0.699)	(0.621)	(0.682)
F-statistic of excluded instruments	161.1	162.9	161.0	283.6	365.0
Observations	24,575	24,575	24,575	24,575	24,575
Cohort dummies	Yes	Yes	Yes	Yes	Yes
Region dummies		Yes	Yes	Yes	Yes
Baseline sales and wages			Yes	Yes	Yes
Industry and firm size dummies				Yes	Yes
△ Net exports and △ ICT					Yes

Notes. The table reports the estimates from the first stage regression of the change in robots per 1000 worker in the Italian labour market on the corresponding change in other seven European countries between 1994 and 2012. Panel A focuses on the restricted sample (see text for definitions) used in the firm-level analysis. Panel B shows results for the full sample of firms. Column (1) includes cohort dummies. Column (2) controls for firms' region of location (20 dummies). Column (3) includes log average sales and wages in 1994. Column (4) includes five firm size dummies (11-20, 21-50, 51-100, 101-500, 500 or more employees) and four dummies for broad manufacturing sectors (food, beverages and tobacco; consumer goods, industrial goods and capital goods). Column (5) further controls for the change in net exports form China and Eastern Europe over 1994-2012 and the change in ICT capital over 1996-2012. The Kleibergen and Paap (2006) first-stage F-statistic of excluded instruments is reported in both panel A and B. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C.6 Robot exposure and firm outcomes, 1994-2012: OLS and 2SLS estimates

	(1) Labour costs	(2) Assets	(3) Debts	(4) Profits	(5) Years w/profits > 0	(6) Exit rate
Panel A: Restricted sample						
$\triangle$ robots per 1000 worker	4.527**	-1.413	1.441	-4.805	1.179	-0.038
	(2.040)	(4.340)	(1.224)	(17.756)	(1.092)	(0.056)
Observations	19,771	19,763	19,802	19,684	19,802	19,802
Panel B: Full sample						
$\triangle$ robots per 1000 worker	5.529***	1.643	2.916*	-9.053	1.349	-0.038
	(1.907)	(4.238)	(1.580)	(13.518)	(1.034)	(0.048)
Observations	24,530	24,517	24,575	24,397	24,575	24,575

Notes. The table reports estimates of ξ from equation (3.2). Outcomes are cumulated values over 1994-2012. Columns (1) to (6) report the estimates for cumulative labour costs (from balance sheet data), cumulative assets, cumulative debts, cumulative profits – all in percent of their baseline levels in 1994 –, years with positive profits and the exit rate of firms. Coefficients are multiplied by 100. All regressions include the full set of controls from column (6) of Table 3.7. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C.7 Robot exposure and firm employment by worker skills and firms' productivity, 1994-2012: 2SLS estimates

	(1)	(2)	(3)	(4)
	Employment	Employment by worker skill		
		Bottom tercile	Middle tercile	Top tercile
Panel A: Restricted sample				
△ robots per 1000 worker	4.008**	2.027	2.152***	-0.505
	(1.840)	(1.358)	(0.601)	(1.016)
Panel B: By VA/Empl tercile				
△ robots bottom tercile	-3.188	1.255	-0.696	-3.323**
	(2.960)	(2.265)	(1.386)	(1.635)
△ robots middle tercile	1.862	1.176	1.666**	-1.105
	(1.577)	(1.158)	(0.713)	(0.906)
△ robots top tercile	12.777***	3.741***	5.184***	2.619*
	(2.448)	(1.007)	(1.074)	(1.350)
Observations	18,564	18,564	18,564	18,564

Notes. The table reports estimates of ξ from equation (3.2) in the restricted sample, focusing on firms with non-missing information about workers' skills (see text for details). Outcomes are cumulated values over 1994-2012. Panel A reports results for the full sample. Panel B decomposes the effect by terciles of firms' value added per worker. Column (1) reports the estimates for cumulative employment in percent of baseline employment in 1994, whereas columns (2) to (4) report the estimates for cumulative employment of workers in the bottom, middle and top terciles of worker skills. All regressions include the full set of controls from column (6) of Table 3.7. Robust standard errors, clustered by 36 two-digit NACE Rev.2 industries, in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.