One Cohort at a Time: A New Perspective on the Declining Gender Pay Gap*

Jaime Arellano-Bover, Nicola Bianchi,
Salvatore Lattanzio, Matteo Paradisi

November 20, 2024

Abstract

This paper studies the interaction between the decrease in the gender pay gap and the stagnation in the careers of younger workers, analyzing data from the United States, Italy, Canada, and the United Kingdom. Our findings highlight the importance of labor-market entry to understand the shrinking of the gender pay gap. The entire decline in the aggregate pay gap originates from (i) newer worker cohorts who enter the labor market with smaller-than-average gender pay gaps and (ii) older worker cohorts who exit with higher-than-average gender pay gaps. Convergence at labor-market entry originates primarily from younger men's positional losses in firms' hierarchies and the overall pay distribution. We propose an explanation by which a larger supply of older workers can crowd out younger workers from a limited number of top-paying positions. These negative career spillovers disproportionately affect the career trajectories of younger men because they were more likely than younger women to hold higher-paying jobs at baseline. Consistent with this aging-driven crowd-out interpretation, younger men experience the largest positional losses within the hierarchies of firms that are more exposed to the aging of the workforce. These findings hold after controlling for alternative explanations for the progressive closure of the gender pay gap at labor-market entry. Finally, we document that labor-market exit has been the sole contributor to the decline in the gender pay gap after the mid-1990s, indicating that without structural breaks, the closure of the gender pay gap is unlikely in the foreseeable future.

JEL Classification: J16, J31, J11.
Keywords: gender gap, workforce aging, cohort turnover, wage growth, labor-market entry, entry wages, initial conditions, age pay gap.

*Contact information: Jaime Arellano-Bover, Yale University, IZA, CESifo, 165 Whitney Avenue, New Haven, CT 06511 jaime.arellano-bover@yale.edu; Nicola Bianchi, Kellogg School of Management, Northwestern University, and NBER, 2211 Campus Drive, 60201 Evanston IL, nicola.bianchi@kellogg.northwestern.edu; Salvatore Lattanzio, Bank of Italy, via Nazionale 91, 00184 Rome, salvatore.lattanzio@bancaditalia.it; Matteo Paradisi, EIEF, CEPR, via Sallustiana 62, 00187 Rome, matteo.paradisi@eief.it. We thank Patricia Cortés, Gordon Dahl, Oren Danieli, Nicole Fortin, Fabian Lange, Claudia Olivetti, Michael Powell, Uta Schönberg, Alexandra Spitz-Oener, as well as participants at various seminars and conferences for helpful comments. We thank Sergey Abramenko, Thomas Barden, Carolina Bussotti, Sean Chen, and Chengmou Lei for outstanding research assistance. The realization of this article was possible thanks to the sponsorship of the "VisitINPS Scholars" program. The views expressed in this paper are those of the authors only and should not be attributed to the Bank of Italy nor to the Eurosystem.

1 Introduction

The gender pay gap has been decreasing in many high-income economies since the mid-1970s (Blau and Kahn, 2017). Throughout the same period, younger workers have been faring progressively worse compared to older workers, experiencing a widening age pay gap (Rosolia and Torrini, 2007; Bianchi and Paradisi, 2023), lower likelihood of being promoted to higher-paying jobs (Bianchi et al., 2023), and lower employment (Mohnen, forthcoming). In this paper, we document how the worsening of the labor-market outcomes of younger workers is related to the narrowing gender pay gap observed in the employee population as a whole.

Using both US survey data and Italian administrative records, we show that cross-cohort generational turnover accounts for the entire decrease in the aggregate gender pay gap: newer cohorts with smaller-than-average gender pay gap levels gradually replaced older worker cohorts with larger-than-average gaps, closing the overall gender pay differential. Crucially, we document that cross-cohort gender convergence at labor-market entry was primarily driven by younger men falling closer to younger women in both firms' hierarchies and the market-wide pay distribution. Finally, in both the US and Italy, gender convergence at labor-market entry stopped in the mid-1990s, a period that coincided with a slowdown in the closure of the aggregate gender pay gap (England, Levine, and Mishel, 2020). Since then, the retirement of older worker cohorts with higher-than-average gender pay differentials has accounted for the entire decline in the gender pay gap, a result that has important consequences for projecting the closure of the pay gap. In secondary analyses, we show that survey data from Canada and the UK display similar trends.

We argue that the empirical facts are consistent with an explanation by which an increased supply of older workers crowds out younger workers from higher-ranked and higher-paying jobs. We formalize this logic by extending a model of the labor market with cross-cohort spillovers between younger and older workers (Bianchi and Paradisi, 2023) to include gender groups. Albeit stylized, this model generates several predictions that closely align with our empirical findings.

The aging-driven crowd-out explanation works as follows. Over the past four decades, the supply of older workers in high-income economies has increased as a result of population aging and longer life expectancy. Older workers have been able to hold onto their (often higher-paying and higher-ranked) jobs for longer due to a variety of non-mutually exclusive factors, such as knowledge spillovers (Sandvik et al., 2020), firm-specific human capital (Lazear, 2009), and employment protection laws (Bentolila and Bertola, 1990). While absorbing this positive shock to the employment of older workers, firms, especially those in later stages of their life cycle, have not always been able to expand their ranks at the top and protect promotion paths for younger workers. As previously shown by Bianchi et al. (2023) and Bianchi and Paradisi (2023), younger workers have been progressively crowded out from higher-paying and higher-ranked positions, while older workers have experienced the opposite trend. Crucially, compared to younger women, younger men were substantially more likely to hold higher-paying jobs before workforce aging started. Hence, workforce aging has crowded out younger men from top jobs at a higher rate than younger women, contributing to the shrinkage in gender pay differentials.

Our first empirical result illustrates how the closure of the gender pay gap over the past four decades has stemmed from reductions happening through labor-market entry and exit of worker cohorts, rather than over the life cycle of any given cohort. This finding aligns well with our model's prediction that workforce aging shrinks the gender pay gap among younger workers. Moreover, we show that the recent deceleration in the closing of the total gender pay gap in the economy, which started in the mid-1990s, coincided with a slowdown in the closing of the between-cohort gap.

We then quantify how much of the shrinkage in the gender pay gap stems from cross-cohort convergence between men and women. Specifically, we compute the between-cohort change in the gender pay gap, which excludes any variation in the gap taking place over the cohorts' life cycle. In practice, we assign to all workers in our data the mean earnings that workers of the same gender and birth year earned at 25 years old. The results are robust to alternative choices for the age of labor-market entry. In all countries, more than the entire decline in the aggregate gender pay gap in weekly earnings can be accounted for by a progressive decrease in the between-cohort component. For example, the between-cohort change in the US equals 127 percent of the total decline in the gender pay gap between 1976 and 2019. Our analysis thus reveals that the trend in the aggregate gender pay gap can be entirely explained by the entry of younger cohorts with smaller earnings differentials and the exit of older cohorts with larger differentials.

Next, we show evidence consistent with our model's prediction that workforce aging narrows the gender pay gap across cohorts by worsening the labor outcomes of younger men more than those of younger women. First, we document that the share of higher-paying managerial jobs held by younger men plummeted between the mid-1970s and mid-1990s, while the share of these jobs held by younger women stayed relatively constant throughout. For example, in the US, the share of managerial jobs with pay in the top quartile of the year-specific distribution held by men between 25 and 30 years old dropped from 12 percent in 1976 to 5 percent in 1995, converging to the 2-percent share of these jobs held by younger women.

Second, we examine where men and women at age 25 ranked over time in the overall pay distribution. The narrowing of the gender pay gap at labor-market entry, which lasted until the mid-1990s, was primarily driven by younger men falling closer to younger women in the pay distribution. In the US, the average rank of younger men at age 25 fell from the 50th percentile of the wage distribution in 1976 to the 39th percentile in 1995, while the mean position of women at age 25 remained fairly stable around the 30th percentile during the same period. After the mid-1990s, the positions of younger men and younger women followed the same flat trajectory.

Third, we use the Italian administrative data with employer-employee matched records to show that younger men experienced larger positional losses than younger women within the pay distributions of all types of firms. However, these losses were the largest within higher-paying firms, where older workers have become especially overrepresented over time. Moreover, subsequent cohorts of younger men were progressively less likely to sort into higher-paying firms, eventually converging towards younger women in being underrepresented at higher-paying firms.

We continue to use the Italian administrative data to draw a more direct link between the firm-level exposure to workforce aging and the career outcomes of younger men and women. To this end, we regress the firm-level change in the gender pay gap at ages 25-30 between 1976 and 1986 (the period of maximum convergence between the two genders) on the change in the firm-level share of workers aged 51-60 during the same period, controlling for trends correlated with firm age, firm size, mean firm pay, and province.

Given that the main regressor can be endogenous, we instrument it with the firm-level difference between the share of workers aged 41-50 and the share of workers aged 51-60 in 1976. This instrument leverages cross-firm variation in the exposure to workforce aging stemming from the age distribution of workers over 40 at baseline. The high tenure of these workers indicates that cross-firm differences in this instrumental variable are likely to originate from hiring decisions that firms made decades earlier, rather than recent firms' actions or firm-level shocks. These IV regressions confirm that the gender pay gap among younger workers closed significantly faster within firms that experienced a larger increase in the share of older workers. Within these firms, younger men faced significantly more negative career outcomes and converged more dramatically toward younger women.

Finally, if we focus on the last twenty years of data, we show that the gender pay gap has been shrinking due almost exclusively to the exit of older cohorts from the labor market, as convergence at labor-market entry stopped in the mid-1990s.[1] Therefore, in contrast to forecasts based on trends in the aggregate gender pay gap (World Economic Forum, 2023), which usually predict that the mean earnings of men and women will match in a few decades, we project that the gender pay gap is not slated to disappear in the high-income countries in our sample in the absence of structural breaks.

We consider a variety of alternative explanations for our results, such as time-varying selection into employment, an increase in part-time work, changing trends in educational attainment, variations in the child pay penalty borne by new parents, sectoral shifts, and the decline of manufacturing. Our tests indicate that these channels do not align with the full range of our findings as well as the aging-driven crowd-out channel.

A vast literature documents the existence of different types of gender gaps in the labor market (Altonji and Blank, 1999; Goldin, 2006; Goldin, Katz, and Kuziemko, 2006; Niederle and Vesterlund, 2011; Azmat and Petrongolo, 2014; Fortin, Bell, and Böhm, 2017, Bertrand, 2020; Olivetti, Pan, and Petrongolo, 2024), as well as the recent convergence in the gender pay gap in most high-income economies (Olivetti and Petrongolo, 2016; Blau and Kahn, 2017). Two papers related to our analysis are Goldin (2014) and Blundell, Lopez, and Ziliak (2024), which show that the gender gap in earnings has become progressively smaller for younger cohorts in the United States while increasing within each cohort. Moreover, Fortin (2019) shows similar facts for Canada.

We build upon this evidence by zooming in on the determinants and consequences of cross-

[1] Consistent with prior work (for example, Sloane, Hurst, and Black, 2021), we show that in both Italy and the US, at least among college-educated workers, a large portion of the remaining gender pay gap at labor-market entry by the mid-1990s can be accounted for by differential sorting across higher- and lower-paying majors.

cohort convergence, using data covering a long period and different countries. Our first contribution is to show that cohort-driven effects can fully account for the dynamics of the gender pay gap over the past forty years (i) partially through the *inflows* of newer cohorts with lower gaps for the first two decades, and (ii) exclusively through the *outflows* of older cohorts with higher gaps for the last two decades. Second, we complement prior findings on the progressive improvement in the wage levels and occupational choices of women (Hsieh et al., 2019) by demonstrating that part of the wage compression between younger men and younger women was driven by a relative worsening in the labor-market outcomes of the former group. Third, we show that the trend in the gender pay gap is likely linked to the seemingly unrelated phenomenon of career spillovers between older and younger workers.[2]

Prior work has studied the importance of economic conditions at the time of labor-market entry for the career progression of new entrants. For example, Kahn (2010) and Oreopoulos, von Wachter, and Heisz (2012) have documented that macroeconomic conditions at the time of college graduation affect the careers of new entrants for several years after entry.[3] Bovini, De Philippis, and Rizzica (2024) and Foliano et al. (2024) have instead analyzed the gender pay gap at labor-market entry and its evolution over time. We contribute to this body of work by quantifying the importance of conditions at labor-market entry in shaping the trajectory of the *aggregate* gender pay gap. We find that a greater gender balance in the earnings of new entrants not only was relevant for the gender pay gap of the directly affected cohorts, but was also a major driver of the closure of the economy-wide gender pay gap.

Finally, building on recent work on the worsening outcomes of younger workers (Bentolila et al., 2022; Guvenen et al., 2022; Dabla-Norris, Pizzinelli, and Rappaport, 2023) and on the interconnectedness of the careers of older and younger workers (Bertoni and Brunello, 2021; Boeri, Garibaldi, and Moen, 2022; Ferrari, Kabátek, and Morris, 2023; Guaitoli and Pancrazi, 2023; Bianchi et al., 2023; Bianchi and Paradisi, 2023; Mohnen, forthcoming), we show both empirically and theoretically that aging-driven crowd-out of younger workers from top jobs can disproportionately worsen the labor-market outcomes of younger men, leading to a narrowing of the gender pay gap.

The remainder of this paper is structured as follows. Section 2 discusses the conceptual framework. Section 3 describes our data. Section 4 quantifies the contribution of cross-cohort effects in shrinking the aggregate gender gap. Section 5 zooms in to compare entry wages and job positions of younger men and younger women, both within and between firms. Section 6 explores alternative channels and robustness checks. Section 7 presents direct evidence on the link between workforce aging and the gender gap. Section 8 provides further insights into the sources of cross-cohort convergence and the drivers of the remaining gender pay gap at entry. Section 9 concludes.

[2] Explanations for the closing of the gender gap are typically driven by labor supply, yet our proposed mechanism is driven by labor demand—firms' hiring and promotion frictions (see Black and Spitz-Oener (2010) for another example of a demand-driven mechanism).

[3] Further evidence on the importance of initial conditions includes Arellano-Bover (2022), which documents the effect of initial macroeconomic conditions on long-term skill accumulation, and Arellano-Bover (2024), which shows the relevance of first-employer quality for long-term earnings.

2 A Model of the Gender Gap with Cross-Cohort Spillovers

Bianchi et al. (2023) and Bianchi and Paradisi (2023) illustrate how an increase in the supply of older workers can limit access to higher-paying jobs for younger workers, thus widening the pay gap between the two age groups. Under the same conditions that lead to crowding out of younger workers from top jobs, the stylized framework in this section shows that workforce aging can contribute to shrinking the pay gap between men and women.

Production. An economy with a price-taking representative firm has a fixed supply of l_y younger workers and l_o older workers. Workers differ with respect to their gender so that the labor supply of each age group can be written as the sum of men and women ($l_a = m_a + f_a$). The firm employs these labor inputs to perform a top job t and a bottom job b. Production occurs through the production function $AY\left(L_y, L_o\right)$, where A is a productivity shifter, $Y_{L_a} > 0$, and $Y_{L_a,L_a} < 0$ $\forall a \in \{y, o\}$. Moreover, younger and older workers are complements in production such that $Y_{L_y,L_o} > 0$. The inputs L_y and L_o are efficiency units of younger and older labor, respectively: $L_a = \theta_{a,t}\left(m_{a,t} + f_{a,t}\right) + \theta_{a,b}\left(m_{a,b} + f_{a,b}\right)$, where $\theta_{a,j}$ is the marginal productivity in job $j \in \{t, b\}$ of workers in the age group $a \in \{y, o\}$. We assume that $\theta_{a,t} > \theta_{a,b}$ $\forall a$ to make all workers more productive in the top job.

Cross-cohort spillovers and gender pay gap. In this model, we introduce key features to generate negative career spillovers and a gender pay gap.

We begin by incorporating two components that enable cross-cohort spillovers in employment levels and wages. First, we assume that older workers' wages and job allocations are stickier than those of younger workers. Theoretical and empirical evidence from labor and organizational economics supports this assumption. Older workers are often more sheltered from firm-level shocks due to factors such as backloaded wage schemes (Lazear, 1979; Ke, Li, and Powell, 2018), firm-specific human capital (Lazear, 2009; Gathmann and Schönberg, 2010), knowledge spillovers (Sandvik et al., 2020; Cornelissen, Dustmann, and Schönberg, 2023), and employment protection laws (Bentolila and Bertola, 1990).[4]

In this model, we capture this notion by assuming that the firm inherits a stock of older workers from period -1 and cannot change their job assignment and wages. Specifically, the number of *legacy workers* in job j is $\rho_j l_{o,j}^{-1}$, where ρ_j denotes the retention rate in job j and $l_{o,j}^{-1}$ represents the number of older workers in job j in period -1.

Second, we assume that the firm incurs a quadratic organizational cost proportional to the parameter $\kappa > 0$ to create and maintain K slots at the top ($K = l_{o,t} + m_{y,t} + f_{y,t}$). This cost reflects the constraint the firm faces when expanding higher-ranked positions, even when it is financially capable of paying the associated wages.[5] In fact, top jobs often entail managerial responsibilities, decision-making authority, and complex tasks. Therefore, creating a new top job requires slack

[4] For example, Bianchi and Paradisi (2023) uses Italian administrative data to show that older workers' wages respond significantly less than those of younger workers to firm-level negative value-added shocks.

[5] This idea of constrained slots, especially at the top of firms' hierarchies, aligns with models and findings discussed by Lazear, Shaw, and Stanton (2018), Bianchi et al. (2023), and Bianchi and Paradisi (2023).

in both the firm's organizational capacity (available high-level responsibilities) and its payroll budget.

Next, we introduce a gender wedge in favor of men which leads to a gender pay gap. Specifically, although men and women are perfect substitutes in production, (i) the firm pays a quadratic cost proportional to c_g for employing younger workers of gender g in the top job, and (ii) this cost is higher for younger women $(c_f > c_m)$.[6] These parameters, which make younger women less concentrated in top jobs, can be microfounded as either taste-based or statistical discrimination.[7]

Wage formation and timing. Following Acemoglu and Restrepo (2023), we assume that wages in the top job pay an exogenous rent over wages in the bottom job: $w^g_{a,t} = \mu_a w^g_{a,b}$, where $\mu_a > 1$ is the exogenous rent for age group a and $w^g_{a,j}$ is the wage in job j for workers in age group a and gender group $g \in \{m, f\}$.[8]

The timing is as follows. First, the firm is endowed with older workers from period -1. Then, given a set of wages for younger workers, the firm decides how many younger men and younger women to slot in the top and bottom jobs by equating their marginal revenue products in the two positions to their marginal costs. Based on these decisions, the firm allocates younger workers randomly between the top and bottom jobs until its labor demands in the two positions are satisfied. Finally, the production is realized, and the firm pays all workers.

The firm problem. The firm problem is to choose the number of younger men and younger women (hereafter, shortened to *younger men and women* for ease of exposition) to employ in top and bottom jobs that maximizes its profits, as follows:

$$\max_{m_{y,t}, f_{y,t}, m_{y,b}, f_{y,b}} AY\left(L_y, L_o\right) - \sum_{g\in\{m,f\}} \sum_{a\in\{y,o\}} \sum_{j\in\{t,b\}} \left(w^g_{a,j} g_{a,j}\right) - \frac{\kappa}{2} K^2 - \sum_{g\in\{m,f\}} \left(\frac{c_g}{2} g^2_{y,t}\right).$$

Appendix B discusses the full solution of the firm problem and provides all the proofs of the following results. Moreover, it includes several extensions: an alternative source of the gender pay gap, a different parametrization of the organizational cost of top jobs, an endogenous labor supply without full employment, and no exogenous rents in wages.

In equilibrium, there is a gender gap in employment in top jobs. Specifically, the marginal revenue product of labor of younger men and women in bottom jobs, and hence their wages in both bottom and top jobs, are the same. However, the number of younger women in top jobs is lower than the number of younger men in those jobs because the cost of employing the former is higher than the cost of employing the latter: $c_f > c_m$. In fact, it is optimal for the firm to keep a constant ratio between younger men and women in top jobs below 1: $f_{y,t}/m_{y,t} = c_m/c_f = \delta_f < 1$.

[6] In principle, we could apply these costs also to older workers. However, they would be redundant because the firm does not choose older workers' employment in period 0.

[7] Alternatively, we could rewrite the production function assuming that (i) men and women are imperfect substitutes in the top job, and (ii) women are less productive than men in the top job.

[8] Acemoglu and Restrepo (2023) show that these wedges can be microfounded using either efficiency wages or bilateral wage bargaining.

Comparative statics. In this labor market, cross-cohort spillovers are crucial drivers of the trend in the gender pay gap. By construction, the gender pay gap can change only if the pay gap between younger men and women differs from that found among older workers. To this end, we consider the effects of an increase in the number of legacy older workers in top jobs inherited by the firm in period 0 ($l_{o,t}^{-1}$) on the mean wages of younger men and women.[9] For example, in the United States, the mean age of the population increased from 32.5 years in 1976 to 38.7 years in 2019 (Figure A1, Panel A), while the mean age of managers increased from 40.1 in 1976 to 44.3 in 2019 (Figure A1, Panel B).

The mean wage $\bar{w}_{y,g}$ of younger workers of gender g changes as follows:

$$\frac{\partial \bar{w}_{y,g}}{\partial l_{t,o}^{-1}} = \underbrace{\frac{1}{g_y}\left(\mu_y - 1\right) w_{y,b} \frac{\partial g_{y,t}}{\partial l_{o,t}^{-1}}}_{\text{Career spillovers}} + \underbrace{\left(\frac{1}{g_y}\left(\mu_y g_{t,y} + g_{b,y}\right)\right) \frac{\partial w_{y,b}}{\partial l_{o,t}^{-1}}}_{\text{Wage level}}, \tag{1}$$

for each $g \in \{m, f\}$.

An increase in the number of older workers induces the firm to add slots at the top of its hierarchy due to the complementarity in production between younger and older workers. However, there is a threshold $\bar{\kappa}$ for the cost parameter κ above which this endogenous increase in top jobs K does not compensate for the increased supply of older workers. Hence, when $\kappa > \bar{\kappa}$, the first component of Equation (1) describes negative career spillovers. An increased supply of older workers at the top restricts the career opportunities of younger workers by decreasing their chances of being assigned to higher-paying jobs ($\frac{\partial g_{y,t}}{\partial l_{o,t}^{-1}} < 0$). In contrast, the second component of Equation (1) is positive and refers to changes in the level of wages paid to younger workers in both bottom and top jobs. Having more older workers increases the wages of younger workers (i) due to the complementarity of younger and older workers in the production function (Freeman, 1979; Welch, 1979; Berger, 1985) and (ii) due to the fact that younger workers become more likely to be in bottom jobs and, therefore, their marginal revenue product of labor increases.

Asymmetric impacts for younger men and younger women. The model predicts that these negative career spillovers are larger in magnitude among younger men as long as the latter are more concentrated in top jobs than younger women ($\frac{m_{y,t}}{m_y} > \frac{f_{y,t}}{f_y}$). Therefore, by blocking younger workers from top positions, a larger supply of older workers can narrow the preexisting gender pay gap by compressing the earnings of younger men and women toward the bottom of the distribution.

In this framework, an increase in the retention rate of older workers and a decrease in the rate of economic growth each produce the same negative spillovers on the employment outcomes of younger workers. The data confirm that these two phenomena have coexisted with population aging: workers have progressively experienced a slower GDP growth rate (Figure A1, Panel C) and a lower degree of firm dynamism (Decker et al., 2014; Hummels and Yue, 2024).

[9] The implicit assumption is that firms made past hiring decisions without considering future changes (i) in the relative size of cohorts, (ii) in the length of workers' careers, and (iii) in the economic growth rate. We capture this dynamic consideration within our static framework by studying an unforeseen increase in the number of older workers after period -1.

Extension: heterogeneous firms. We replace the representative firm with N heterogeneous firms to study how the gender pay gap varies within and between different firms. In line with prior research (Antwi and Phillips, 2013; Ruffini, 2022), higher-paying firms (higher A_n) have a higher retention rate of older workers (higher $\rho_{t,n}$). On the labor-supply side, worker i of age group a and gender g derives the following utility when working in job j and firm n: $U_{i,a,j,n} = \log\left(w^g_{a,j,n}\right) + \xi_{i,a,j,n}$, where $\xi_{i,a,j,n}$ represents the idiosyncratic preference of worker i over job j of firm n. We assume that $\xi_{i,a,j,n}$ is unobserved by firms and follows a type-1 extreme distribution with a parameter $\sigma > 0$.

This extension (full details in Appendix B) produces two additional insights. First, an increase in the supply of older workers decreases the probability of younger men holding top jobs more than that of younger women within all firms, regardless of their productivity level. However, these negative career spillovers are larger in magnitude for both younger men and women within higher-productivity and higher-paying firms because the number of older workers increases more in these firms.

Second, given that younger men are more likely to be displaced from the top jobs of higher-paying firms, they are also more likely than younger women to migrate toward the bottom jobs of other firms. Appendix B outlines under what conditions younger men move from higher-paying to lower-paying firms, a result that finds empirical support.

Extension: introducing skills. In this extension (full details in Appendix B), we discuss how our framework can accommodate the slower decline in the gender pay gap that started in the mid-1990s. A key prediction from the baseline model is that younger men are more concentrated in top jobs than younger women and, therefore, suffer more harshly from these positions being progressively occupied by older workers. When younger men and women become equally concentrated in top jobs ($\frac{m_{y,t}}{m_y} \approx \frac{f_{y,t}}{f_y}$), further increases in the supply of older men can still harm the prospects of younger workers but cannot do so differentially across genders. After this point, factors other than the initial job assignment, such as gender imbalances that predate entry into the labor market, become the primary drivers of changes in the gender pay gap.

Here, we assume that each worker enters the labor market with either high (h) or low (l) skills, which represent cross-worker differences in college major choices among other pre-labor-market factors. Moreover, each job is divided into two different tasks, and there is a one-to-one correspondence between skills and tasks. The younger workers' vector of efficiency units of labor is $\mathbf{L_y} = \left(L_{y,h}, L_{y,l}\right)$, and workers with different skills are complements in the production function. The rest of the firm problem is unchanged.

We model a scenario in which younger men and women are concentrated equally in top jobs by assuming that the number of older workers in top jobs is large enough that no higher-ranked positions remain available for younger workers in period 0. We then study what happens to the mean wages of younger workers when the number of older workers in top jobs increases even further. Whether the resulting wage change is larger for men or women crucially depends on their distribution across different skills/tasks. If the complementarity between younger and older workers is

proportional to a task's marginal product, the change in mean wages of younger women is larger than that of younger men if women are overrepresented in high-skill tasks.

Given that prior research has established that men choose higher-return college majors (skill h in the context of our framework) more than women do (for example, see Black et al., 2008; Bertrand, 2020; Huneeus et al., 2021; Sloane, Hurst, and Black, 2021; Bovini, De Philippis, and Rizzica, 2024; Humphries, Joensen, and Veramendi, 2024), an increased supply of older workers is unlikely to further shrink the gender pay gap.

Summary. In the rest of the paper, we document a set of empirical facts that are consistent with several predictions of this stylized framework:

1. Cross-cohort differences in the gender pay gap are the main source of decline in the aggregate gender pay gap (Section 4).[10]

2. The gender pay gap narrows across cohorts because younger men lose more positions in the overall pay distribution and in firms' hierarchies, relative to younger women (Section 5).

3. These cross-cohort spillovers are more negative among younger men than among younger women within all firms. For all younger workers, they are larger in magnitude among higher-paying firms (Section 5).

4. Compared to younger women, younger men leave higher-paying firms at a higher rate (Section 5).

5. When the concentration of younger men and women in top jobs becomes more balanced and predetermined education choices account for a large share of the gender pay gap, the gender pay gap stops shrinking (Section 8).

3 Data

This paper uses a combination of administrative and survey data with 376,814,659 observations from four high-income countries: the United States, Italy, Canada, and the United Kingdom (Table A1). Due to their much larger sample size and informational depth, our main analyses focus on the US and Italy, while leaving data from Canada and the UK primarily to illustrate the generality of the main findings. The Italian and US datasets further allow us to exploit the relative strengths of both administrative population data (employer-employee linkages) and large-scale survey data (detailed information on all persons regardless of labor force participation).

[10] A model inspired by the one in Card and Lemieux (2001), in which older workers are perfect substitutes for younger men and imperfect substitutes for younger women, can produce this result without additional firm-level constraints on the creation of top jobs. However, this model is not compatible with further findings.

3.1 US: Current Population Survey and American Community Survey

Most of our US analyses rely on forty-four years (1976-2019) of repeated cross-sections from the Annual Social and Economic March Supplement of the Current Population Survey (CPS), which we accessed through IPUMS (Flood et al., 2022). We impose similar sample restrictions across all the datasets available to us: we limit our sample to individuals who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks during the past year, and had earned strictly positive earnings.[11]

The CPS data allow us to construct three compensation measures: (i) the annual wage and salary income, (ii) weekly earnings, obtained by dividing the annual wage and salary income by the number of weeks worked during the previous year, and (iii) hourly earnings, obtained by dividing weekly earnings by the usual number of hours worked per week. All compensation measures are expressed in 2015 USD, using the CPI provided by the Bureau of Labor Statistics. Moreover, they are winsorized yearly at the 99.9^{th} percentile from above.

Since the CPS does not include information on college graduates' field of study, we complement some of our analyses with data from the 2009–2019 waves of the American Community Survey (ACS), also accessed through IPUMS (Ruggles et al., 2023).

3.2 Italian Social Security Data

We use confidential administrative data provided by the Italian Social Security Institute (INPS). This dataset comprises forty-four years (1976-2019) of matched employer–employee records for the whole population of private-sector, nonagricultural firms with at least one salaried employee. In each year, we focus on workers who were between 25 and 64 years old, had worked at least 24 weeks, had earned strictly positive earnings, and had not retired by December 31. Unlike the CPS data, the INPS dataset matches workers to firms, a feature that allows us to draw a closer link between firm-level exposure to workforce aging and the trends in the gender pay gap.

The Italian administrative data include two main compensation measures: (i) total annual earnings, which comprise all forms of gross labor compensation, and (ii) full-time-equivalent (FTE) weekly earnings, computed as the ratio between total annual earnings and FTE working weeks. Both variables are expressed in 2015 euros using the CPI provided by the OECD. Moreover, they are winsorized at the 99.9^{th} percentile from above.[12] We complement the administrative data with the Labor Force Survey for the years 2009-2019 to shed light on the changes in college graduates' fields of study.

3.3 Luxembourg Income Study

We leverage survey data from the Luxembourg Income Study (LIS) database for two additional high-income countries that have long time series and sufficiently large sample sizes: Canada (1973-

[11] At baseline, we focus on private-sector workers for comparability with the Italian administrative records. However, we show that our results are robust to including US public-sector workers in the sample.

[12] Annual earnings are also winsorized at €3,000 from below to address a few instances in which yearly compensation is implausibly low.

2019) and the United Kingdom (1976-2019). We can compute weekly earnings only for Canada, while total yearly labor earnings are available in both countries. Both earnings variables are expressed as 2011 purchasing-power-parity US dollars using the conversion tables directly provided by LIS. Moreover, they are winsorized at the 99.9^{th} percentile from above. Whenever possible, we apply the same sample restrictions that we used for the US and Italian data. Specifically, we keep workers who were between 25 and 64 years old, had worked at least 24 weeks (available only in Canada), and had strictly positive labor earnings.

4 The Cross-Cohort Decline in the Gender Pay Gap

4.1 Aggregate and Cohort-Specific Trends in the Gender Pay Gap

Aggregate gap. For at least the last four decades, the United States and Italy have been experiencing a decrease in the gender gap in weekly earnings (Figure 1). Between 1976 and 2019, the gender gap shrank by 0.47 log points or 59 percent (relative to the baseline level of 0.8 log points) in the US and by 0.19 log points or 57 percent (relative to the baseline level of 0.33 log points) in Italy.[13] If we replace yearly earnings for weekly earnings, the decrease in the gender gap remains large, at -59 percent in the US and -26 percent in Italy (Figure A2). The shrinkage in the gender gap extends to the other two countries in our dataset, appearing to be a generalized trend within high-income economies (Figure A3). Specifically, we find that the gender gap decreased by 0.29 log points or 43 percent in Canada between 1973 and 2019 (weekly earnings) and by 0.63 log points or 60 percent in the United Kingdom between 1976 and 2019 (yearly earnings).

Gap between and within cohorts. We now start to highlight the importance of cross-cohort effects in driving the overall negative aggregate trend in the gender pay gap (Figure 2, Panels A and B). Specifically, for various worker birth cohorts, we plot the trend in the gender gap from the year in which all workers in the cohort turned 25 years old to the year in which they turned 50 years old.[14] This graph shows three main results that will inform most of the further analysis in the rest of the paper.

First, consistent with Prediction 1, the aggregate downward trend in the gender gap stems entirely from the fact that younger worker cohorts entered the labor market with smaller initial differences in the log weekly earnings of men and women. In the US, the gender gap at age 25 decreased from 0.55 log points for the cohorts born between 1947 and 1951 to 0.19 log points for the cohorts born between 1967 and 1971. Similarly, in Italy the gender gap at age 25 declined from 0.21 log points for the cohort born in 1951 to 0.09 log points for the cohort born in 1971.[15]

[13] In the United States, the aggregate gender pay gap was largely stable between 1955 and 1975 (Bailey, Helgerman, and Stuart, 2024).

[14] We stop at 50 years old to limit the influence coming from selection into retirement.

[15] In the CPS data, we create groups of five cohorts to increase precision. In this case, we start reporting the cohort-level gaps when all cohorts within a group turn 25 years old. In the remainder, we refer to these groups using their youngest cohort.

Second, the gender gap almost always increases over the life-cycle of each cohort (Bertrand, Goldin, and Katz, 2010).[16] Moreover, this within-cohort increase becomes steeper across subsequent worker cohorts, a trend that works against closing the aggregate gender gap. If we consider the 1951 cohort, the gender gap increased between age 25 and age 30 by 0.07 log points in the US and 0.01 log points in Italy. The same increase for the 1971 cohort was equal to 0.17 log points in the US and 0.03 log points in Italy. These within-cohort dynamics are consistent with many prior findings in the literature. For example, the child penalty incurred by mothers in the labor market (Kleven, Landais, and Søgaard, 2019) and women's propensity to move towards more flexible but lower-paying jobs (Goldin, 2014) are two factors that contributed to increasing the gender pay gap after labor-market entry, especially among more recent worker cohorts.[17]

Third, convergence at labor-market entry across subsequent worker cohorts significantly slowed after 1995 (see also Figure A4). As already discussed, the gender pay gap at age 25 dropped by 0.36 log points in the US and by 0.12 log points in Italy between the 1951 cohort and the 1971 cohort, but then declined by only 0.01 log points in the US and by 0.02 log points in Italy by the 1991 cohort. The end of this form of convergence coincided with a slowdown in the closure of the aggregate gender pay gap.

Similar patterns are present across alternative definitions of earnings, as well as in Canada and the United Kingdom (Figure A5).

4.2 Decomposing the Gender Pay Gap Within and Between Cohorts

We now introduce a more formal test of Prediction 1, which states that most of the decline in the gender pay gap has been taking place across worker cohorts. To this end, we propose a decomposition of the change in the gender pay gap within and between worker cohorts.

First, we write the average log earnings of birth-year cohort c of gender g at time t ($w_{c,g,t}$), as the sum of two terms: $w_{c,g,t} = w^e_{c,g} + \Delta w_{c,g,t}$, where $w^e_{c,g}$ represents cohort c's mean log earnings at the time of entry into the labor market (or alternatively, at an early career stage), and $\Delta w_{c,g,t} = w_{c,g,t} - w^e_{c,g}$ is the growth of the average log earnings of cohort c between year t and labor-market entry (e).[18] Let $w_{g,t}$ represent the average log earnings of all workers of gender g in year t and let $a(c,t)$ denote the age of cohort c in year t. The change between year t and year $t' > t$ in the

[16] The flattening toward the end of each cohort's career is consistent with more negative selection into early retirement among women (Goldin and Mitchell, 2017).

[17] The timing of fertility may have played a role: in the past four decades, the average maternal age at the birth of a child has moved from 26 to 29 in the US and from 28 to 32 in Italy (OECD, 2022).

[18] Our between-within cohorts decomposition does not directly map into "cohort effects" in the canonical model of separate and unrestricted age, time, and cohort effects (see chapter 2.7 in Deaton, 1997). Instead, we define birth cohorts by the interaction of age and time, and our economic mechanism implicitly assumes that there are no fundamental unobserved differences (for example, men's ability) across birth cohorts driving our results. Most of our robustness checks in Section 6 tackle potential failures of this assumption.

average log earnings of gender group g is:

$$w_{g,t'} - w_{g,t} = \underbrace{\sum_{c:a(c,t')\in[25,64]} s_{c,g,t'} \cdot w^e_{c,g} - \sum_{c:a(c,t)\in[25,64]} s_{c,g,t} \cdot w^e_{c,g}}_{\text{Between-Cohort Change}} + \underbrace{\sum_{c:a(c,t')\in[25,64]} s_{c,g,t'} \cdot \Delta w_{c,g,t'} - \sum_{c:a(c,t)\in[25,64]} s_{c,g,t} \cdot \Delta w_{c,g,t}}_{\text{Within-Cohort Change}}, \quad (2)$$

where $s_{c,g,t}$ is the share of workers of gender g from cohort c in year t over the total number of workers of the same gender. This share operates as a weight for a given cohort of active employees (between 25 and 64 years old) in a given year.

The first two terms in Equation (2) quantify the *between-cohort* change between t and t'. This between-cohort component captures variation in the mean log earnings of gender group g that stems from (i) cross-cohort differences in early-career earnings and (ii) changes in the distribution of workers across birth cohorts over time. In contrast, the last two terms of Equation (2) measure the *within-cohort* change between t and t', which isolates variation in the mean log earnings of gender group g that originates from (i) changes in the life-cycle earnings growth for the average cohort active in the labor market between t and t', and (ii) changes in the relative size of cohorts over time. If we subtract Equation (2) for women from Equation (2) for men, we can quantify the contribution of the *between-cohort* and *within-cohort* components to the aggregate trend in the gender pay gap.

To implement this decomposition on the data, we need to know the earnings of each cohort c at the time of labor-market entry e, but this information is available only for Italy, the country for which we have access to panel data. In our baseline results, we compute $w^e_{c,g}$ as the mean earnings of each cohort-gender pair in the year in which the workers in that birth cohort were 25 years old. In Section 6, we show that the results are robust if we make different choices about the entry age.

The data highlight the importance of cross-cohort effects in driving the trend in the overall gender pay gap (Figure 2, Panels C and D). In both the United States and Italy, the between-cohort decline in the gender pay gap is at least as large in magnitude as the total decline. These findings hold for different earnings measures and in other high-income countries (Figure A6). The fact that, in most sample years, the between-cohort component accounts for more than 100 percent of the total decline in the gender pay gap confirms that the gender gap has been increasing over the life cycle of the average worker cohort.

Finally, denote by $\underline{t}$ the first sample year in each dataset; we do not observe the entry earnings of all workers older than $t - \underline{t} + 25$ in year t. In the baseline analysis, we assign the average earnings in year $\underline{t}$ to each cohort-gender pair comprising workers older than 25 in year $\underline{t}$. However, the results are robust to measuring the initial earnings of older cohorts in alternative ways. First, we assess the sensitivity of our results to assigning mean earnings in the first sample year, rather than at age 25, to older cohorts (Table A2, Columns 7 to 8). To this end, we repeat the anal-

ysis starting from 2015, the first year we observe earnings at age 25 for all worker cohorts in our sample. Second, we use data from the US Decennial Census to measure earnings at age 25 before the beginning of our sample in 1976 (Figure A7). In both cases, the between-cohort component continues to account for a majority of the drop in the gender pay gap.

5 The Positions of Younger Workers in the Wage Distribution

Section 5.1 and Section 5.2 below present the empirical evidence in support of Prediction 2, which states that the gender pay gap among younger workers has closed because younger men have fallen toward younger women in firms' hierarchies and in the pay distribution.

5.1 The Probability of Holding Higher-Ranked Positions

We first study trends in the probability of younger men and women holding higher-ranked jobs within firms' hierarchies. We classify higher-ranked positions as all management occupations with associated annual earnings in the top quartile of the year-specific distribution.[19]

In both countries, younger workers have become substantially less likely to hold higher-paying managerial jobs (Figure A8). For example, in the US, the likelihood of a higher-ranked position being filled by a younger worker decreased from 13 percent in 1976 to 7 percent in 2019, even though higher-ranked jobs accounted for a slightly increasing share of all jobs in the economy (from 6.5 percent in 1976 to 7 percent in 2019) during the same period.

Importantly, this crowding out of younger workers from the top of firms' hierarchies has affected younger men more than younger women, which aligns with Prediction 2 (Figure 3, Panels A and B). In the US, the share of higher-ranked jobs filled by men between the ages of 25 and 30 fell from 12 percent in 1976 to 5 percent in 1995, and then to 4 percent in 2019. Meanwhile, the corresponding share for women in the same age group and job category increased from 1 percent in 1976 to 2 percent in 1995, and to 3 percent in 2019.

This finding holds even when the definition of higher-ranked positions is extended to managerial jobs with above-median pay (Figure A9).

5.2 The Pay Rank Gap at Labor-Market Entry

This section shows that the main takeaway of Section 5.1 holds if we investigate changes in the rank of new entrants within the wage distribution that existed at the time of their entry in the labor market, rather than their probability of holding higher-paying managerial jobs.

Let $F_t(w)$ represent the distribution of weekly earnings for all workers in year t.[20] For each

[19] In the CPS data, managerial occupations are identified using 2-digit Standard Occupational Classification (SOC) code 11. In the INPS data, we use the highest ranked position out of the four main job categories in the Italian labor system: in ascending order, these are apprenticeships, blue-collar jobs, white-collar jobs, and managerial jobs. The figures based on Italian data show a spurious trend discontinuity in the mid 1990s because the definition of managerial jobs in the INPS data changes from 1996.

[20] Using the position in the population's pay distribution allows a close mapping between the model and the results. In fact, Prediction 2 concerns the *rank* of younger men and women in the pay distribution, rather than their average wage. In addition, by focusing on ranks, we avoid any potential confounders in the level of wages that stem from

new entrant i with weekly earnings w_{it}^E, we compute where they rank in the overall distribution: $p_{it} \equiv F_t(w_{it}^E)$. We then compute the average rank among all entrants of gender g in year t as follows: $\bar{p}_{gt} = \frac{1}{N_{gt}} \sum_{i \in \{g,t\}} p_{it}$, where N_{gt} is the number of labor market entrants of gender g in year t. As discussed in Section 4.2, we fix the time of labor-market entry in the data at age 25.

Building upon the analysis of the racial wage gap in Bayer and Charles (2018), we also consider an alternative metric focusing on the *median* entrant, rather than the average. Let $F_t(w)$ represent the distribution of weekly earnings for all workers in year t, while $F_{gt}^E(w)$ represents the distribution of weekly earnings at entry among gender g in year t. The gender-specific quantile-q entry pay is given by w_{qgt}^E, defined so that $F_{gt}^E(w_{qgt}^E) = q$. Let P_{qgt} quantify where the qth percentile of weekly earnings at entry of gender g ranks in the overall distribution: $P_{qgt} = F_t(w_{qgt}^E)$ with $q = 0.5$.

Consistent with Prediction 2, the gender pay gap at labor-market entry has shrunk because the mean position of younger men in the wage distribution has declined substantially more than that of younger women (Figure 3, Panels C and D). In 1976, the average rank of younger men at age 25 was equal to the 50th percentile of all weekly earnings in the United States and to the 47th percentile in Italy. By 1995, the position of younger men had fallen to the 39th percentile in the US and to the 36th percentile in Italy. During the same period, the mean position of women at age 25 remained fairly stable around the 30th percentile in both countries.

We reach the same conclusions if we consider the median position, rather than focusing on the mean rank (Figure A10, Panels A-B). Similarly, the results are robust to using (i) survey data from Canada and the United Kingdom (Figure A10, Panels C-D) and (ii) hourly earnings for the US (Figure A11).

One concern is that the positional loss younger men experienced may not reflect a true worsening of their labor-market outcomes if wages at the lower end of the pay distribution have converged to the median. However, prior research shows that low wages have grown more slowly than wages at or above the median in both the Unites States (Autor, Katz, and Kearney, 2008; Danieli, 2024) and Italy (Depalo and Lattanzio, 2024). To further address this concern, we calculate the distance between younger workers' mean log weekly earnings (separately for younger men and women) and the total mean log weekly earnings in the private sector. Even after accounting for changes in the earnings distribution over time, the data indicate that younger men shifted closer to the bottom of the distribution, converging toward younger women (Figure A12).

These results do not necessarily imply younger women have not experienced progress in their conditions at labor-market entry. For example, women may have experienced advancements in their labor outcomes (Hsieh et al., 2019), but more significant improvements among older women may have kept younger women near the bottom of the pay distribution. To explore this possibility, we replicate our baseline analysis by recalculating the ranks of younger men and women within the earnings distribution of men aged over 55, therefore excluding older women from the benchmark group (Figure A13). At least in the United States, younger women gained some ground within this more restricted pay distribution between 1976 and 1995. Nonetheless, the main result

trends in wage inequality, among other factors.

remains: the gender pay gap at labor-market entry closed primarily because younger men fell closer to the bottom of the pay distribution.

5.3 Younger Workers in Lower-Paying and Higher-Paying Firms

This section documents how the gender pay gap among new entrants has changed in lower-paying and higher-paying firms. For this analysis, we are limited to showing evidence from Italy because we must use administrative data matching workers to firms.

We start by dividing workers into one hundred percentiles based on their employer's mean wage in each sample year so that these firm groups have the same number of workers but varying wage levels. As expected, the data indicate that workforce aging is more severe among higher-paying firms, in which monetary returns of postponing retirements are higher (Figure A14). Between 1976 and 1995, the same period in which the careers of younger workers, especially men, substantially slowed down, the share of workers over 50 employed by firms with above-median mean pay increased by 6 percentage points and by firms in the top two deciles of mean pay by 5.2 percentage points.

Next, we analyze the outcomes of younger men and women across these firm groups. To this end, we compute the mean percentile of men and women at age 25 using the wage distribution within each percentile of mean firm pay. The empirical findings support Prediction 3, which states that younger men experience career spillovers that are more negative than younger women do in higher- and lower-paying firms, and that these losses are larger for younger workers of both genders in higher-paying firms, where workforce aging is more extreme. Compared to younger women, younger men have experienced larger declines in their mean percentile within ninety-three out of one hundred firm groups (Figure 4, Panel A). Moreover, the results indicate that the positional losses were more pronounced within higher-paying firms for all younger workers. Between 1976 and 1995, the mean positional change within above-median firm groups was equal to -10 percentiles for men and -6 percentiles for women, while the same change within below-median firm groups was -7 percentiles for men and +0.5 percentiles for women. As already discussed in Section 5.2, the gender pay gap among new entrants stopped closing after the mid-1990s. Consistent with this prior evidence, we find that younger men and women experienced similar rank losses between 1995 and 2019 in all firm groups (Figure 4, Panel B).

Finally, consistent with Prediction 4, the share of younger men in higher-paying firms has declined more than that of younger women (Figure 4, Panel C): between 1976 and 1995, the probability of 25-year-old men working in the top decile of firm groups decreased on average by 6 percentage points (a 62 percent decline from the 1976 level), while the same probability for 25-year-old women fell by only 2 percentage points (a 28 percent decline from the 1976 level). At baseline, younger men were fairly equally distributed across lower-paying and higher-paying firms, while women were overrepresented among lower-paying firms.[21] Over time, the distribution of younger men across different firm groups moved closer to that of younger women. As expected, there are

[21] Our result that younger women were always overrepresented in lower-paying firms is consistent with the findings in Card, Cardoso, and Kline (2016) and Casarico and Lattanzio (2024).

not remarkable differences across the two genders after 1995 (Figure 4, Panel D).

6 Alternative Explanations and Robustness Checks

6.1 Labor Force Participation

In this set of tests, we explore whether our prior results can stem from differential trends in the labor force participation of men and women, especially around labor-market entry. In the US, female labor force participation at age 25 increased by 8 percentage points between 1976 and 1986 and has remained stable ever since (Figure A15).[22]

To account for changes in selection into the labor market, we follow the procedure outlined in Blau et al. (2024). First, we expand the sample to include individuals who worked less than 24 weeks in the year (but at least 100 hours) and had positive labor earnings. Second, we impute the weekly earnings of nonparticipants based on their probability of falling into each decile of the pay distribution. We predict these probabilities via an ordered probit estimated separately by year and gender using the following observable characteristics: years of education, a dummy for college graduation, a dummy for advanced degrees, potential experience (age-years of education-6), potential experience squared, race (white, black, others) and ethnicity (hispanic) dummies, and fixed effects for Census divisions.[23]

All our prior results hold when we replicate them on this selection-adjusted sample (Figure A16). The between-cohort component keeps explaining more than the entire decline in the aggregate gender pay gap. Moreover, the average rank of younger men at age 25 fell from the 57th percentile of the selection-adjusted distribution of weekly earnings in 1976 to the 42nd percentile in 1995, while the average rank of younger women at age 25 remained fairly stable around the 35th percentile during the same period. The relative positions of both genders showed little variation after the mid-1990s. These results alleviate the concern that the increase in the labor-force participation of younger women between 1976 and 1986 had direct negative effects on the career outcomes of younger men, a conclusion that is consistent with the findings in Fukui, Nakamura, and Steinsson (2023). The results are robust also if we impute hourly, rather than weekly, earnings to nonparticipants (Figure A17).

6.2 Sectoral Shifts

We test whether the decline in the gender pay gap across cohorts is related to variation in the sorting of men and women across sectors. To do so, we consider the between-cohort component in Equation (2) and further decompose it between and within sectors. Sectors are defined using

[22] This section focuses on the US for two reasons. First, the INPS data include only employees. Second, data from the Bank of Italy's Survey of Household and Income Wealth indicate that women's labor-force participation was fairly constant in the 1980s and the 1990s and only slightly declined in the 21st century.

[23] See Appendix C for more details.

1-digit NACE Rev. 2 codes in the United States and 2-digit codes in Italy.[24]

In both countries, the data indicate that most of the between-cohort decline in the gender pay gap has taken place within sectors (Figure A18). Specifically, 18 percent of the total between-cohort change between 1976 and 2019 occurred across 1-digit sectors in the US, while none of the decline between 1976 and 2019 took place between 2-digit sectors in Italy. Therefore, the data suggest that a loss of employment in economic sectors where men historically received high wages, such as manufacturing (Charles, Hurst, and Schwartz, 2019), does not appear responsible for a meaningful portion of the cross-cohort decline in the gender pay gap.

Next, we directly address the progressive decline of the manufacturing sector in two ways. First, we replicate our main finding after dropping from the sample all workers who are employed in manufacturing (Figure A19). Second, using data from the first five sample years (1976-1981), we estimate a probit model for an indicator for working outside of manufacturing as a function of several worker-level characteristics.[25] Next, we use the estimated coefficients to predict the probability of working outside of manufacturing for the whole sample. Then, we retain in the sample only individuals with a very high predicted probability (> 0.7), reducing the sample size by 35 percent in the United States and by 57 percent in Italy (Figure A20). All prior results hold after either dropping every manufacturing worker or retaining only workers with a very high predicted probability of working outside of manufacturing.

6.3 Changes in Educational Attainment

We investigate whether an increase in the college graduation rate of women or a decline in that of men can explain the cross-cohort shrinkage of the gender pay gap.[26] Using an approach similar to the one described for studying sectoral shifts, we decompose the total decline in the gender pay gap as well as its between-cohort component within and between two education levels: workers with and without a college degree (Figure A21). In this decomposition, the between-college component quantifies how much of the gender pay gap's decline stems from the fact that women have become more likely to graduate from college over time (or men have become less likely to graduate). In the US, the between-college dimension accounts for only a minor share of the decline in both the total (14 percent) and between-cohort (14 percent) gender gap between 1976 and 2019.

Next, we directly show that the key results in Section 4 and Section 5 hold for individuals with and without a college degree (Figure A22). For example, the finding that younger men have become less likely than younger women to hold higher-paying managerial jobs is robust when analyzing only the jobs held by college graduates or workers without a college education.[27]

[24] Although the information on 2-digit codes is available for the United States, we use 1-digit codes in order to have enough observations in each cohort-year-sector cell.

[25] We use the following regressors: US Census divisions or Italian provinces fixed effects interacted with age, age squared, completed education (less than high school, high school, some college, four years of college or more; US only), gender, race (the same three categories used to address selection into the labor market; US only), ethnicity (US only), and a foreign born dummy (Italy only).

[26] Here, we focus on the US because the INPS data lack information on completed education for most workers in the sample.

[27] In this extension, we define a higher-paying managerial job as one with earnings in the top quartile of the year-and-

These findings are consistent with the hypothesis that workforce aging significantly contributes to the decline in the gender pay gap among younger workers. Older workers with and without a college education have become more numerous, occupying a larger share of the higher-ranked jobs they can achieve within their education level. Even if not all these jobs are positioned at the very top of firms' hierarchies, workforce aging can still impede the career progression of younger workers (especially men) of similar educational backgrounds. In other words, older workers with a college education can crowd out younger men of similar education from top managerial jobs, while older workers without a college education can crowd out younger workers of similar education from middle-management positions.

Finally, we test whether our choice of observing many outcomes at age 25 is problematic. For example, an increasing trend in college graduation rates over time could change the average work experience of 25-year-olds. First, we show that our findings are robust to alternative definitions of younger workers: individuals observed at ages 25 (the baseline), 28, 30, and ages 25 to 30 (Figure A23; Table A2, Columns 3 to 6). Second, the between-cohort component continues to account for the whole decline in the gender pay gap if we exploit the longitudinal dimension of the Italian data to define $w^e_{c,g}$ as the earnings observed during the first year in the labor market for each worker who started working after 1976 (Table A2, Columns 9 to 10). Third, the findings hold when we analyze younger workers by years of potential work experience rather than by age (for the US only; Figure A24 and Figure A25).

6.4 Child Penalty

We test whether a progressive decrease in the child penalty borne by mothers in the labor market could explain the importance of the between-cohort convergence in earnings. Since childbirth tends to happen in a woman's earlier career stages, a progressively lower child penalty could disproportionately benefit women in younger cohorts. We expect that accounting for this factor would reduce the contribution of the between-cohort component, albeit modestly, given that the overall wage effects of the child penalty estimated in the literature are relatively small.

For each year, we estimate how much of the between-cohort change in the gender pay gap is explained by disparities in the negative consequences of having children on the careers of new mothers and fathers. Following Kleven, Landais, and Søgaard (2019) and Casarico and Lattanzio (2023), we first calculate estimates of the child penalty in weekly earnings between mothers and fathers in the United States and between women with and without children in Italy, using a subset of data for which we have fertility information from maternity leave applications. We then create counterfactual weekly earnings by multiplying women's weekly earnings by the product of the child penalty and the fraction of mothers in each year and cohort. This variable is an estimate of women's earnings if we were to eliminate the adverse effects of parenthood on women's careers.[28]

As expected, we find that accounting for the child penalty decreases the contribution of the between-cohort component to the overall decline in the gender gap, but this effect tends to be

education-specific pay distribution.

[28] Appendix D provides more details about this procedure.

small in magnitude (Figure A26). Even after adjusting for the child penalty in the United States and Italy, the main takeaways of the prior two sections still holds: the decline in the gender pay gap stems entirely from differences across worker cohorts, and younger men lost more ground than younger women in the pay distribution.

6.5 Other Trends

Full-time workers. We then move to investigate whether a disproportionate increase in the proportion of full-time female workers could serve as an explanation for the between-cohort convergence in mean earnings. To this end, we replicate all the prior findings, limiting the sample to only full-time workers (Figure A27). In this more restricted sample, the main results hold.

Public employees. For comparability, we omit public-sector employees from our preferred sample as they are present in the CPS data but not in the Italian INPS data. Here, we show that all our prior findings hold if we include public-sector employees in the CPS data. We can also use this larger sample to replicate the procedure outlined in Section 6.1 to account for time-varying selection into both the private and public sector. The data clearly indicate that the results are robust if we (i) include public-sector employees in the sample (Figure A28, Panels A-C) and (ii) expand the sample to public employees while also accounting for time-varying and gender-specific selection into the labor market (Figure A28, Panels D and E).

Residual earnings. Finally, we control for multiple observable characteristics at once. We regress the log of weekly earnings on a dummy for part-time workers, a dummy for temporary workers (only in Italy), a dummy for domestic-born workers (only in Italy), dummies for race (only in the US), a dummy for college graduation (only in the US), and a dummy for workers with children (only in the US). For the US, we have a second version in which we control for even more variables: in addition to the previous regressors, we include a dummy for hispanic ethnicity, dummies for marital status, and fixed effects for Census divisions. We estimate these regressions separately by year and country to allow the coefficients to vary over time and across different datasets. Based on these coefficients, we compute residual earnings and use them to show that the main findings in Section 4 and Section 5 are all robust (Figure A29).

7 Firm-Level Exposure to Workforce Aging

This section draws a direct link between workforce aging and changes in the gender pay gap of younger workers. Using the Italian employer–employee matched data, we exploit variation across firms in the degree of aging they experience. Consider the following firm-level regression equation:

$$\Delta y_f^{(25-30)} = \alpha + \beta \Delta s_f^{(51-60)} + X_f' \gamma + \epsilon_f,$$

where Δ represents changes in firm-level variables between 1976 and 1986, the period of fastest pay convergence between younger men and women. The dependent variable $y_f^{(25-30)}$ measures

either the gender pay gap or the mean outcomes of men and women in firm f, computed for workers aged 25-30. The key independent variable $s_f^{(51-60)}$ denotes the share of firm f's workers who are between 51 and 60 years old. The vector $X_f^{'}$ includes province fixed effects, firm age, firm size, and mean firm pay, all observed at baseline.[29]

The parameter of interest is β, which measures how changes in the share of older workers in firm f affect the evolution of the gender pay gap or mean pay rank of younger employees. However, OLS estimates of β may be biased if unobserved firm-level labor-demand shocks influence both workforce aging and changes in younger workers' outcomes.

7.1 Instrumental Variable: Firms' Age Structure at Baseline

To address this concern, we propose an instrument that exploits variation in the age distribution of firms' workforce at baseline. In particular, we instrument $\Delta s_f^{(51-60)}$ with the ten-year change in the share of firm f's workforce over 50 projected on the basis of the age distribution of workers just younger than 50 in 1976. Denoting the projected $\Delta s_f^{(51-60)}$ by $\widetilde{\Delta} s_f^{(51-60)}$, our instrument is defined as:

$$\widetilde{\Delta} s_f^{(51-60)} = \left(s_{f,1976}^{(41-50)} - s_{f,1976}^{(51-60)} \right),$$

where $s_{f,1976}^{(41-50)}$ is the share of firm f's workers between 41 and 50 years old in 1976.[30] This instrument captures changes in the firm-level share of workers over 50 between 1976 and 1986 driven by the natural aging of employees who were *already* working at firm f in 1976. This strategy mirrors the approach used by Mohnen (forthcoming) to study the effects of differences in retirement rates across US commuting zones on youth employment.

IV Relevance. The variable $\widetilde{\Delta} s_f^{(51-60)}$ has two characteristics that make it a suitable instrument. First, it is highly, but not perfectly, correlated with the endogenous regressor $\Delta s_f^{(51-60)}$ (Table A3, Panel A). A 1-percentage-point increase in the $\widetilde{\Delta} s_f^{(51-60)}$ is associated with a 0.58-percentage-point increase in $\Delta s_f^{(51-60)}$, an effect that is statistically significant at the 1 percent level. The fact that this coefficient is less than one suggests some degree of turnover among older workers. The strength of the first stage is evidenced by Kleibergen-Paap F-statistics that exceed 3,000 (Table A4).

IV Exclusion. Second, the exogeneity of the instrument relies on the assumption that the projected change in workers over 50 in 1976 ($\widetilde{\Delta} s_f^{(51-60)}$) affects changes in younger workers' outcomes between 1976 and 1986 ($\Delta y_f^{(25-30)}$) only through its influence on the actual change in the share of workers over 50 during the same period ($\Delta s_f^{(51-60)}$). Although this exclusion restriction is not

[29] When the dependent variable is the outcomes of younger men and women, rather than the gender pay gap, we estimate stacked regressions in which the outcomes of men and women are treated as separate observations in each year. The modified regression is $\Delta y_{g,f}^{(25-30)} = \alpha + \beta \Delta s_f^{(51-60)} + \gamma \Delta s_f^{(51-60)} \cdot m_g + \delta m_g + X_f^{'} \theta + \epsilon_{g,f}$, where the subscript g denotes gender, and m_g is a dummy variable equal to 1 for men's labor-market outcomes.

[30] The sample for this analysis consists of all firms with more than five total employees in 1976, at least one man and one woman under 30 years old in both 1976 and 1986, and at least one worker over 40 in 1976. These restrictions ensure that we consider non-micro firms employing younger men, younger women, and older workers. To limit noise, we compute firm-level values for year x (for example, 1976) as three-year averages from x to $x+2$.

directly testable, we argue that, unlike $\Delta s_f^{(51-60)}$, differences in $\widetilde{\Delta} s_f^{(51-60)}$ mostly reflect cross-firm variation in hiring decisions made many years before 1976. Consequently, they are unlikely (or at the minimum, much less likely than cross-firm differences in $\Delta s_f^{(51-60)}$) to be directly related to unobserved shocks or firms' decisions affecting the change in the gender pay gap of younger workers between 1976 and 1986.[31] To support this claim, we show that 58.5 percent of workers between 41 and 50 years old in 1986 were hired at their employer ten or more years earlier (Figure A30).[32]

7.2 Main Results

We first show the reduced-form estimates by plotting the relationship between the IV and younger workers' outcomes using binned scatter plots (Figure 5). There is a negative and linear relationship between the projected firm-level workforce aging and the change in the gender pay gap among younger workers (Panel A). The gender pay differential among younger workers narrows more dramatically within firms more exposed to workforce aging because younger men's career outcomes worsen significantly more than those of younger women (Panel D).

We next show OLS and 2SLS regression estimates (Table 1, Panel A). Here, the actual change in the share of older workers is the main regressor of interest, and the outcome variable is the change in young workers' gender pay gap. According to the OLS estimates, a 1-percentage-point increase in the share of workers aged 51-60 between 1976 and 1986 is associated with a 2.9-percentile additional reduction in the gender gap in pay rank among younger coworkers at the firm. This effect is 65 percent larger than the average change in the gender pay gap during this period. The 2SLS estimates corroborate this conclusion, confirming that firms experiencing more pronounced workforce aging saw a sharper convergence in the mean pay rank between younger men and women. Here, a 1-percentage-point rise in the share of older workers is associated with a 4.2-percentile faster closure in the gender pay rank gap. [33]

Consistent with our findings in Section 5, the gender gap in pay rank closes faster in firms with more workforce aging because younger men fall closer to younger women in the pay distribution (Table 2, Panel A). According to the 2SLS regressions, a 1-percentage point increase in the share of workers aged 51-60 correlates with a 5.8-percentile positional loss for younger men and with only a 1.1-percentile positional loss for younger women. The former effect is statistically significant at the 1-percent level, while the latter is insignificant at the 10 percent level.

[31] For example, a negative demand or productivity firm-level shock could relate to (i) ϵ_f if it induces the firm to pay new entrants lower wages, and (ii) $\Delta s_f^{(51-60)}$ if experienced workers are less likely to leave the firm following the negative shock. However, since the IV $\widetilde{\Delta} s_f^{(51-60)}$ is predetermined in 1976, it is unlikely to correlate with such shocks once X_f is controlled for.

[32] We compute the tenure distribution for 1986, the first year with complete data on at least ten years of tenure (due to the dataset's 1976 start date).

[33] The larger magnitude of 2SLS relative to OLS could be driven by a combination of IV correcting for confounders and giving greater weight to "complier" firms—those for whom projected aging better predicts actual aging. These firms could have a more immobile set of experienced workers, more stringent slot constraints, and thus stronger aging-driven career spillovers.

7.3 Robustness and additional results

These findings generally hold if we estimate the gender pay gap as the difference in weekly wages or log weekly wages among younger workers, rather than the difference in their mean pay rank, although the results using log weekly earnings tend to be slightly less precise (Table A4, Panel A).

Moreover, the results hold if we estimate the effects of workforce aging over the entire twenty-year period between 1976 and 1996, during which the pay differentials between younger men and women narrowed substantially (Table 1, Panel B). For these regressions, the instrument for $\Delta s_f^{(51-60)}$ becomes $\widetilde{\Delta} s_f^{(51-60)} = \left(s_{f,1976}^{(31-40)} - s_{f,1976}^{(51-60)} \right)$, where $s_{f,1976}^{(31-40)}$, the share of firm's f workers aged 31-40 in 1976, replaces $s_{f,1976}^{(41-50)}$.[34] This analysis is more taxing on the data, as it requires observing the same firm twenty years apart, rather than only ten. Consequently, the number of firms included in the estimating sample decreases from 25,536 in the baseline analysis to 14,064 in this extension. Despite this reduction, the results remain consistent with the baseline findings. For example, a 1-percentage-point increase in the share of firm f's older workers between 1976 and 1996 is associated with an additional reduction in the pay rank gap between younger men and women of 5.4 percentiles in the OLS estimates and 8.3 percentiles in the 2SLS estimates. The magnitudes of these effects are slightly smaller than those estimated over the 1976-1986 period, relative to the larger mean decline in the gender pay gap between 1976 and 1996.

We can also replicate this analysis over the subsequent twenty-year period, 1996-2016, when convergence in the gender pay gap at labor-market entry largely stalled (Table 1, Panel C).[35] Based on our earlier findings, we expect to observe more muted effects of workforce aging on younger workers' outcomes during this period. Consistent with this expectation, a 1-percentage-point increase in the share of older workers between 1996 and 2016 is associated with an additional decline in the pay rank gap between younger men and women of 0.4 percentiles in the OLS estimates and 1.2 percentiles in the 2SLS estimates, two effects that are statistically insignificant.

8 Two Phases of Convergence in the Gender Pay Gap

Prior sections have shown that the gender pay gap of new entrants into the labor market rapidly shrank until the mid-1990s and then stabilized until 2019. Section 8.1 quantifies the contribution of the convergence in entry earnings to the overall between-cohort change. Section 8.2 then studies the implications of the post-mid-1990s halt in the convergence at labor-market entry for the future of the aggregate gender pay gap. Finally, Section 8.3 connects the stop in the convergence at entry with gender differences in college majors (Prediction 5).

[34] Similar to workers aged 41-50, these slightly younger workers had a high mean tenure in 1986 (Figure A30, Panel C). Specifically, 51 percent of them were hired at their current firm eight or more years earlier.

[35] In this case, the instrument becomes $\widetilde{\Delta} s_f^{(51-60)} = \left(s_{f,1996}^{(31-40)} - s_{f,1996}^{(51-60)} \right)$.

8.1 Convergence from Inflows and Outflows of Worker Cohorts

We quantify to what extent the decline in the between-cohort gender pay gap has stemmed from the fact that men and women have been entering the labor market with progressively more similar mean earnings (*convergence at entry*), or, instead, from the fact that older worker cohorts, who had higher-than-average gender disparity in earnings, have naturally retired over time (*convergence through exit*).

To this end, starting from the between-cohort component in Equation (2), we further neutralize any cross-cohort convergence in the mean earnings of men and women that happened at the time of entry into the labor market. Specifically, considering a baseline year t_b, our new counterfactual measure of pay fixes the earnings at labor-market entry (at age 25) of all cohorts who enter in that year *or later* to the average earnings at entry computed between t_b and the following two years. Hence, this new between-cohort component measures what would have happened if the convergence in the early-career earnings of men and women had stopped at t_b and if only the natural turnover of older cohorts had affected the overall level of the gender pay gap.

Under this counterfactual scenario, in which we freeze convergence at entry in the first sample year, the data indicate that both sources of convergence (at entry and through exit) are important drivers of the decline in the gender gap that took place between 1976 and 2019 (Figure 6, Panels A and B). In 2019, convergence at entry accounted for 36 percent of the total between-cohort shrinkage in the gender gap in the US and 43 percent in Italy.

Consistent with our prior findings, we observe that the importance of the convergence at labor-market entry wanes as t_b increases, while the opposite is true for convergence through labor-market exit (Figure 6, Panel C). The latter begins to consistently explain at least 100 percent of the total decline in the between-cohort gap when the benchmark year t_b is equal to 2001 in the US and 2003 in Italy. In other words, the entire decline in the gender pay gap between the early 2000s and 2019 originated from the fact that older cohorts, who had higher levels of the gender pay gap, progressively retired, therefore reducing the mean pay differential between men and women who were still active in the labor market.[36]

8.2 Consequences for Future Convergence

The ongoing importance of convergence through exit carries relevant implications for the future of the gender pay gap. Given that retirees with higher-than-average gaps are currently the only source of convergence, the gender pay gap is *not* projected to close. In the absence of further changes in the labor market, it can, at best, converge to the level of the gender pay gap currently observed among younger worker cohorts at the time of labor-market entry.

Strikingly, if we were to predict future convergence by looking at the recent trends in the *aggregate* gender pay gap computed on the whole workforce, we would reach vastly different conclusions. Since the aggregate gender pay gap has been decreasing, any future projection would predict full convergence in a few decades in most high-income economies. For example, in its

[36] These findings also hold in the data from Canada and the United Kingdom (Figure A31).

2023 Global Gender Gap Report, the World Economic Forum predicts that Europe will reach gender parity in 67 years, while North America will get there in 95 years (World Economic Forum, 2023).

To better appreciate the differences between the two predictions, we estimate the rate of earnings convergence at early career stages across subsequent worker cohorts and over time. Denote with g_s^e the gender pay gap at age 25 of the cohort who entered the labor market in year s. Assuming that the convergence rate is linear, we model the gender pay gap at entry in year t as follows:

$$g_s^e = \alpha_t - \beta_t (s - \underline{s}_t), \tag{3}$$

where $\underline{s}_t$ is the entry year of the cohort with the maximum age used to estimate convergence in year t (age 45 in the baseline analysis), $s \in [\underline{s}_t, t]$, α_t is the gender gap at age 25 for the cohort $\underline{s}_t$, and the coefficient β_t measures the rate of convergence in the gender pay gap at age 25 observed between year $\underline{s}_t$ and t. Starting from a given year t, if the convergence continued at the same rate in the following cohorts and the demographic composition of men and women remained the same, the gender pay gap would close for the first time for the cohort who entered the labor market in year $s^* = \frac{\alpha_t}{\beta_t} + \underline{s}_t$.[37] In addition to estimating Equation (3), we can estimate the linear trend in the aggregate gender pay gap for each year t using the prior twenty years of data. We can then use the estimated coefficients from these regressions to predict the first year in which the total gap would close if its linear path continued without modifications.

As expected, the between-cohort gender pay gap at age 25 is not bound to converge. We first show the estimated linear function in Equation (3) in the US and Italy for two years: the year 2000 and the last sample year (Figure 7, Panels A and D). Both countries' convergence rate β_t dramatically decreased from 2000 to 2019. For example, in the United States, β_t declined from 0.008 in 2000 to 0.0003 in 2019. In addition to highlighting two different years, we show the evolution of β_t across all years in the sample (Figure 7, Panels B and E): at least from 1995, the convergence rate rapidly decreased until it reached zero in the second half of the 2010s.

Next, we show how the projected year of closure of the gender pay gap at age 25 changed over time (Figure 7, Panels C and F). Specifically, we plot the predicted s^*, or year of entry into the labor market for the first cohort with zero gap at age 25, for all years t in our sample. The first year with no gender gap at age 25 has been following the opposite trend of the convergence rate β_t. In 1995, the projected first worker cohort without a positive gender pay gap at age 25 was slated to enter the market in 2022 in the United States and 2028 in Italy. By 2019, the entry year of this same cohort had slipped after year 2300 in the US, while no cohort was projected to have a zero gap in Italy. In contrast, simply extrapolating the current negative trend in the total gender gap suggests that convergence should be attainable in 2073 in the US and 2062 in Italy.[38]

[37] If older cohorts have larger gender pay gaps than younger cohorts, s^* underestimates the year in which the gender pay gap would close for the labor market as a whole.

[38] Data from Canada and the United Kingdom allow us to reach the same conclusions (Figure A32).

8.3 Why Did Convergence at Labor-Market Entry Stop?

This section examines why the convergence at entry stopped after the mid-1990s. Our stylized framework predicts that when younger men and women are fairly equally concentrated in higher-paying jobs, a larger supply of older workers cannot shrink the gender pay gap by pushing more younger men toward lower-ranked positions. Consistent with the theory, we have shown that the differences between younger men and women in the share of top jobs (Section 5.1) and in mean pay rank (Section 5.2) were small by 1995 and remained fairly constant afterwards.

After the mid-1990s, the effects of a further increase in the number of older workers on the gender pay gap depend on the sources of the remaining earning differential between younger men and women. If (i) the outstanding gap depends primarily on education outcomes that predate their labor-market entry, such as their choice of a college major (Black et al., 2008; Bertrand, 2020; Huneeus et al., 2021; Sloane, Hurst, and Black, 2021; Bovini, De Philippis, and Rizzica, 2024; Humphries, Joensen, and Veramendi, 2024) and (ii) older workers are not closer complements of younger workers in low-skill jobs, Prediction 5 states that further workforce aging will not close the gender pay gap.

This section documents that gender differences in college majors have accounted for a substantial share of the gender pay gap at labor-market entry since the mid-1990s, a finding that is in line with the model's prediction. For this analysis, we focus on college graduates and assess the portion of the remaining entry gender pay gap explained by their college major choices. We use the American Community Survey for the United States and the Italian Quarterly Labor Force Survey for Italy. Building upon Bertrand (2018), we start from the population of full-time native-born male employees and compute average residual weekly earnings for each college major from regressions that net out year fixed effects and a quadratic polynomial of age. We then quantify the share of individuals in each cohort and gender group that graduated in a specific major. By interacting average residual weekly earnings with the shares of graduates in each college major, we compute the major-predicted average weekly earnings of each cohort-gender combination at labor-market entry.[39]

In both countries, the gender pay gap at entry predicted by younger workers' major choices slightly declined until the mid 1980s (Figure 8).[40] After this period, it has remained remarkably stable for nearly three decades. Notably, the major-predicted gap has been stable since the convergence at entry stopped, constituting approximately 63 percent of the entry gap for college graduates in the United States and 51 percent in Italy.[41] These two large shares are likely to be a lower

[39] Given that we observe the combination of wages and college majors only in more recent years, the major-predicted earnings for older cohorts implicitly assume that the relative wages between majors have been stable over time. In support of this assumption, we show that college major choices have been stable after 1995 in both countries (Figure A33). In addition, our results show a stable major-predicted gender pay gap, which is unlikely to be driven by changes in average wages that perfectly offset each other.

[40] This effect is quantitatively small relative to the observed convergence at entry in average wages (for the US, -0.04 log points compared to -0.25 log points).

[41] Our results based on Italian survey data are slightly smaller than estimates based on Italian administrative data (60 percent from Bovini, De Philippis, and Rizzica (2024)).

bound of the importance of pre-labor-market choices for the gender pay gap at entry, given that we are considering only a single predetermined factor (major choice).

9 Conclusion

In this paper, we examine how workforce aging and resulting negative spillovers on the careers of younger workers contribute to the narrowing of the gender pay gap.

Over the past four decades, older workers have become more numerous and have extended their participation in the labor force, increasing their likelihood of holding higher-ranked and higher-paying positions for longer periods. At the same time, many firms, facing this positive employment shock and dwindling growth prospects, have struggled to expand the number of available slots at the top of their organization, limiting the career progression of younger workers. Our main hypothesis is that younger men have been more exposed to these negative consequences because they were more likely than younger women to hold higher-ranked positions that have been increasingly occupied by older workers.

We report several findings that support this crowd-out mechanism. First, the reduction in the gender pay gap originates primarily from cohort turnover, with newer worker cohorts entering the labor market with smaller gender pay gaps than those of older cohorts. Second, younger men have experienced greater career stagnation and positional losses within firms' hierarchies, leading to a relative convergence with younger women. Third, these effects are especially pronounced in firms that have faced more binding constraints in expanding their top positions, such as those experiencing a larger increase in their supply of older workers. Fourth, in the absence of structural breaks, the labor markets in high-income economies are unlikely to reach full gender pay convergence, as the ongoing reduction in the gap stems mainly from the exit of older workers with large pay disparities.

Overall, this paper highlights the critical role of workforce aging in shaping trends in the gender pay gap. Our results suggest that future efforts to reduce the gender pay gap should view new cohorts' early years in the labor market as a critical period to influence gender pay differentials for each cohort's entire life cycle. Moreover, they should consider the importance of protecting younger workers' job opportunities and career progression to more effectively shield them from the potentially negative consequences of low-growth forecasts, declining dynamism, and an aging workforce.

References

Acemoglu, Daron and Pascual Restrepo. 2023. "Automation and Rent Dissipation: Implications for Inequality, Productivity, and Welfare." working paper, MIT.

Altonji, Joseph G. and Rebecca M. Blank. 1999. "Race and Gender in the Labor Market." *Handbook of Labor Economics*, , ed. David Card and Orley Ashenfelter Vol. 3, Chapter 48, 3143–3259. Elsevier B.V.

Antwi, James and David C. Phillips. 2013. "Wages and health worker retention: Evidence from public sector wage reforms in Ghana." *Journal of Development Economics*, 102: 101–115.

Arellano-Bover, Jaime. 2022. "The effect of labor market conditions at entry on workers' long-term skills." *Review of Economics and Statistics*, 104(5): 1028–1045.
Arellano-Bover, Jaime. 2024. "Career consequences of firm heterogeneity for young workers: First job and firm size." *Journal of Labor Economics*, 42(2): 549–589.
Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics*, 90(2): 300–323.
Azmat, Ghazala and Barbara Petrongolo. 2014. "Gender and the Labor Market: What Have We Learned From Field and Lab Experiments?" *Labour Economics*, 30: 32–40.
Bailey, Martha J., Thomas Helgerman, and Bryan A. Stuart. 2024. "How the 1963 Equal Pay Act and 1964 Civil Rights Act Shaped the Gender Gap in Pay." *The Quarterly Journal of Economics*, 139(3): 1827–1878.
Bayer, Patrick and Kerwin Kofi Charles. 2018. "Divergent paths: A new perspective on earnings differences between black and white men since 1940." *Quarterly Journal of Economics*, 133(3): 1459–1501.
Bentolila, Samuel and Giuseppe Bertola. 1990. "Firing costs and labour demand: How bad is eurosclerosis?" *Review of Economic Studies*, 57(3): 381–402.
Bentolila, Samuel, Florentino Felgueroso, Marcel Jansen, and Juan F. Jimeno. 2022. "Lost in Recessions: Youth Employment and Earnings in Spain." *SERIEs*, 13(1-2): 11–49.
Berger, Mark C. 1985. "The Effect of Cohort Size on Earnings Growth: A Reexamination of the Evidence." *Journal of Political Economy*, 93(3): 561–573.
Bertoni, Marco and Giorgio Brunello. 2021. "Does a Higher Retirement Age Reduce Youth Employment." *Economic Policy*, 36(106): 325–372.
Bertrand, Marianne. 2018. "Coase Lecture–The Glass Ceiling." *Economica*, 85(338): 205–231.
Bertrand, Marianne. 2020. "Gender in the Twenty-First Century." *AEA Papers and Proceedings*, 110: 1–24.
Bertrand, Marianne, Claudia Goldin, and Lawrence F. Katz. 2010. "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors." *American Economic Journal: Applied Economics*, 2(3): 228–255.
Bianchi, Marco, Björn R. Gudmundsson, and Gylfi Zoega. 2001. "Iceland's natural experiment in supply-side economics." *American Economic Review*, 91(5): 1564–1579.
Bianchi, Nicola and Matteo Paradisi. 2023. "Countries for Old Men: An Analysis of the Age Wage Gap." working paper, Northwestern University.
Bianchi, Nicola, Giulia Bovini, Jin Li, Matteo Paradisi, and Michael Powell. 2023. "Career Spillovers in Internal Labor Markets." *Review of Economic Studies*, 90(4): 1800–1831.
Black, Dan A., Amelia M. Haviland, Seth G. Sanders, and Lowell J. Taylor. 2008. "Gender Wage Disparities Among the Highly Educated." *Journal of Human Resources*, 43(3): 630–659.
Black, Sandra E. and Alexandra Spitz-Oener. 2010. "Explaining Women's Success: Technological Change and the Skill Content of Women's Work." *The Review of Economics and Statistics*, 92(1): 187–194.
Blau, Francine D. and Lawrence M. Kahn. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations." *Journal of Economic Literature*, 55(3): 789–865.
Blau, Francine D., Lawrence M. Kahn, Nikolai Boboshko, and Matthew Comey. 2024. "The Impact of Selection into the Labor Force on the Gender Wage Gap." *Journal of Labor Economics*, 42(4): 1093–1133.
Blundell, Richard, Amanda Gosling, Hidehiko Ichimura, and Costas Meghir. 2007. "Changes in the Distribution of Male and Female Wages Accounting for Employment Composition Using Bounds." *Econometrica*, 75(2): 323–363.
Blundell, Richard, Hugo Lopez, and James P. Ziliak. 2024. "Labour Market Inequality and the Changing Life Cycle Profile of Male and Female Wages." *Working Paper, Institute for Fiscal Studies*.
Boeri, Tito, Pietro Garibaldi, and Espen R. Moen. 2022. "In Medio Stat Victus: Labor Demand Effects of an Increase in the Retirement Age." *Journal of Population Economics*, 35(2): 519–556.
Bovini, Gulia, Marta De Philippis, and Lucia Rizzica. 2024. "The Origins of the Gender Pay Gap: Education and Job Characteristics." *Temi di discussione, Rome: Bank of Italy*.
Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *Quarterly Journal of Economics*, 131(2): 633–686.
Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and labor market inequality: Evidence and some theory." *Journal of Labor Economics*, 36(S1): S13–S70.
Card, David and Thomas Lemieux. 2001. "Can falling supply explain the rising return to college for

younger men? A cohort-based analysis." *Quarterly Journal of Economics*, 116(2): 705–746.
Casarico, Alessandra and Salvatore Lattanzio. 2023. "Behind the child penalty: understanding what contributes to the labour market costs of motherhood." *Journal of Population Economics*, 36(3): 1489–1511.
Casarico, Alessandra and Salvatore Lattanzio. 2024. "What Firms Do: Gender Inequality in Linked Employer-Employee Data." *Journal of Labor Economics*, 42(2): 325–355.
Charles, Kerwin Kofi, Erik Hurst, and Mariel Schwartz. 2019. "The Transformation of Manufacturing and the Decline in US Employment." *NBER Macroeconomics Annual*, 33(1): 307–372.
Cornelissen, Thomas, Christian Dustmann, and Uta Schönberg. 2023. "Knowledge Spillovers, Competition, and Individual Careers." *Working paper, UCL*.
Cross-National Data Center. 2023. "Luxembourg Income Study (LIS) Database."
Dabla-Norris, Era, Carlo Pizzinelli, and Jay Rappaport. 2023. "Job Polarization and the Declining Wages of Young Female Workers in the United Kingdom." *Oxford Bulletin of Economics and Statistics*.
Danieli, Oren. 2024. "Skill-Replacing Technology and Bottom-Half Inequality." working paper, Tel Aviv University.
Deaton, Angus. 1997. *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy.* Baltimore and London:Johns Hopkins University Press for the World Bank.
Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2014. "The Role of Entrepreneurship in US Job Creation and Economic Dynamism." *Journal of Economic Perspectives*, 28(3): 3–24.
Depalo, Domenico and Salvatore Lattanzio. 2024. "The Increase in Earnings Inequality in Italy: The Role and Persistence of Atypical Contracts." *Review of Income and Wealth*, forthcoming.
Eissa, Nada and Hilary Williamson Hoynes. 2004. "Taxes and the labor market participation of married couples: the earned income tax credit." *Journal of Public Economics*, 88(9-10): 1931–1958.
England, Paula, Andrew Levine, and Emma Mishel. 2020. "Progress toward gender equality in the United States has slowed or stalled." *Proceedings of the National Academy of Sciences*, 117(13): 6990–6997.
Ferrari, Irene, Jan Kabátek, and Todd Morris. 2023. "Longer careers: A barrier to hiring and coworker advancement." IZA Discussion Paper No. 16098.
Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, and Michael Westberry. 2022. "Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]."
Foliano, Francesca, Alex Bryson, Heather Joshi, Bożena Wielgoszewska, and David Wilkinson. 2024. "Gender Wage Gap Among Young Adults: A Comparison Across British Cohorts." *Labour Economics*, 91: 102614.
Fortin, Nicole M. 2019. "Increasing Earnings Inequality and the Gender Pay Gap in Canada." *The Canadian Journal of Economics*, 52(2): 407–440.
Fortin, Nicole M., Brian Bell, and Michael Böhm. 2017. "Top Earnings Inequality and the Gender Pay Gap: Canada, Sweden, and the United Kingdom." *Labour Economics*, 47: 107–123.
Freeman, Richard B. 1979. "The Effect of Demographic Factors on Age-Earnings Profiles." *The Journal of Human Resources*, 14(3): 289.
Fukui, Masao, Emi Nakamura, and Jon Steinsson. 2023. "Women, Wealth Effects, and Slow Recoveries." *American Economic Journal: Macroeconomics*, 15(1): 269–313.
Gathmann, Christina and Uta Schönberg. 2010. "How general is human capital? A task-based approach." *Journal of Labor Economics*, 28(1): 1–49.
Goldin, Claudia. 2006. "The Quiet Revolution That Transformed Women's Employment, Education, and Family." *American Economic Review P&P*, 96(2): 1–21.
Goldin, Claudia. 2014. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review*, 104(4): 1091.
Goldin, Claudia and Joshua Mitchell. 2017. "The New Life Cycle of Women's Employment: Disappearing Humps, Sagging Middles, Expanding Tops." *Journal of Economic Perspectives*, 31(1): 161–182.
Goldin, Claudia, Lawrence F. Katz, and Ilyana Kuziemko. 2006. "The Homecoming of American College Women: The Reversal of the College Gender Gap." *Journal of Economic Perspectives*, 20(4): 133–156.
Guaitoli, Gabriele and Roberto Pancrazi. 2023. "Global Trends in Income Distribution Across Age Groups?" working paper, University of Warwick.
Guvenen, Fatih, Greg Kaplan, Jae Song, and Justin Weidner. 2022. "Lifetime Earnings in the United States over Six Decades." *American Economic Journal: Applied Economics*, 14(4): 446–79.
Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica*, 47(1): 153–161.

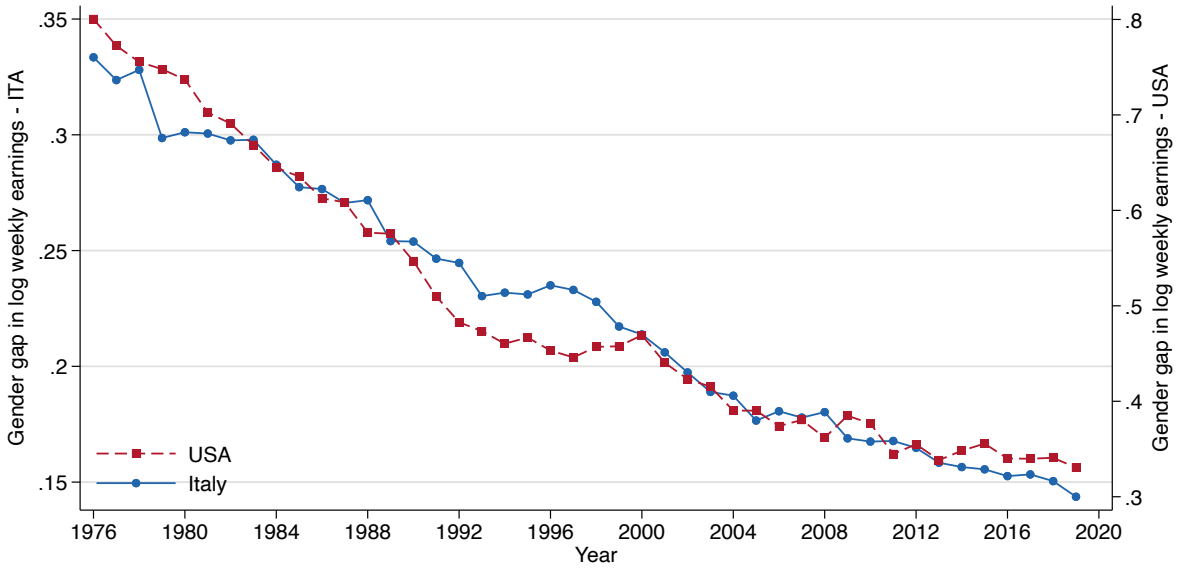
Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow. 2019. "The Allocation of Talent and U.S. Economic Growth." *Econometrica*, 87(5): 1439–1474.
Hummels, David and Kan Yue. 2024. "On the Ubiquity of Declining Business Dynamism." *NBER Working Paper*, 32637.
Humphries, John Eric, Juanna Schrøter Joensen, and Gregory F. Veramendi. 2024. "The Gender Wage Gap: Skills, Sorting, and Returns." *AEA Papers and Proceedings*, 114: 259–264.
Huneeus, Federico, Conrad Miller, Christopher Neilson, and Seth Zimmerman. 2021. "Firm Sorting, College Major, and the Gender Earnings Gap." *Industrial Relations Section Working Paper Series*, 648.
Istituto Nazionale della Previdenza Sociale. 2021. "Database UNIEMENS."
Istituto Nazionale di Statistica. 2023. "Rilevazione Sulle Forze di Lavoro."
Kahn, Lisa B. 2010. "The Long-Term Labor Market Consequences of Graduating From College in a Bad Economy." *Labour Economics*, 17(2): 303–316.
Ke, Rongzhu, Jin Li, and Michael Powell. 2018. "Managing Careers in Organizations." *Journal of Labor Economics*, 36(1): 197–252.
Kleven, Henrik. 2023. "The Geography of Child Penalties and Gender Norms: A Pseudo-Event Study Approach." *NBER Working Paper*, 30176.
Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard. 2019. "Children and Gender Inequality: Evidence from Denmark." *American Economic Journal: Applied Economics*, 11(4): 181–209.
Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler. 2022. "Imperfect competition, compensating differentials, and rent sharing in the US labor market." *American Economic Review*, 112(1): 169–212.
Lazear, Edward P. 1979. "Why Is There Mandatory Retirement?" *Journal of Political Economy*, 87(6): 1261–1284.
Lazear, Edward P. 2009. "Firm-Specific Human Capital: A Skill-Weights Approach Edward P." *Journal of Political Economy*, 117(5): 914–940.
Lazear, Edward P., Kathryn L. Shaw, and Christopher T. Stanton. 2018. "Who Gets Hired? The Importance of Competition Among Applicants." *Journal of Labor Economics*, 36(S1): S133–S181.
Manoli, Dayanand and Andrea Weber. 2011. "Nonparametric evidence on the effects of retirement benefits on labor force participation decisions." *SSRN Working Paper*.
Mohnen, Paul. forthcoming. "The Impact of the Retirement Slowdown on the U.S. Youth Labor Market." *Journal of Labor Economics*.
Mulligan, Casey B. and Yona Rubinstein. 2008. "Selection, Investment, and Women's Relative Wages Over Time." *Quarterly Journal of Economics*, 123(3): 1061–1110.
Niederle, Muriel and Lise Vesterlund. 2011. "Gender and competition." *Annual Review of Economics*, 3(1): 601–630.
OECD. 2022. "Age of mothers at childbirth and age-specific fertility." *OECD Strategic Foresight 2.3*.
Olivetti, Claudia and Barbara Petrongolo. 2008. "Unequal Pay or Unequal Employment? A Cross-Country Analysis of Gender Gaps." *Journal of Labor Economics*, 26(4): 621–654.
Olivetti, Claudia and Barbara Petrongolo. 2016. "The Evolution of Gender Gaps in Industrialized Countries." *Annual Review of Economics*, 8(1): 405–434.
Olivetti, Claudia, Jessica Pan, and Barbara Petrongolo. 2024. "The Evolution of Gender in the Labor Market." *NBER Working Paper*, 33153.
Oreopoulos, Philip, Till von Wachter, and Andrew Heisz. 2012. "The Short- and Long-Term Career Effects of Graduating in a Recession." *American Economic Journal: Applied Economics*, 4(1): 1–29.
Rosolia, Alfonso and Roberto Torrini. 2007. "The generation gap: relative earnings of young and old workers in Italy." Bank of Italy, Economic Research and International Relations Area Temi di discussione (Economic working papers) 639.
Ruffini, Krista. 2022. "Worker earnings, service quality, and firm profitability: Evidence from nursing homes and minimum wage reforms." *Review of Economics and Statistics*, 1–46.
Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Renae Rogers, and Megan Schouweiler. 2023. "IPUMS USA: Version 14.0 [dataset]."
Sandvik, Jason J., Richard E. Saouma, Nathan T. Seegert, and Christopher T. Stanton. 2020. "Workplace Knowledge Flows." *The Quarterly Journal of Economics*, 135(3): 1635–1680.
Sloane, Carolyn M., Erik G. Hurst, and Dan A. Black. 2021. "College Majors, Occupations, and the Gender Wage Gap." *Journal of Economic Perspectives*, 35(4): 223–248.

U.S. Census Bureau. 2023. "American Community Survey."
Welch, Finis. 1979. "Effects of Cohort Size on Earnings: The Baby Boom Babies' Financial Bust." *Journal of Political Economy*, 87(5): S65–S97.
World Economic Forum. 2023. "Global Gender Gap Report."

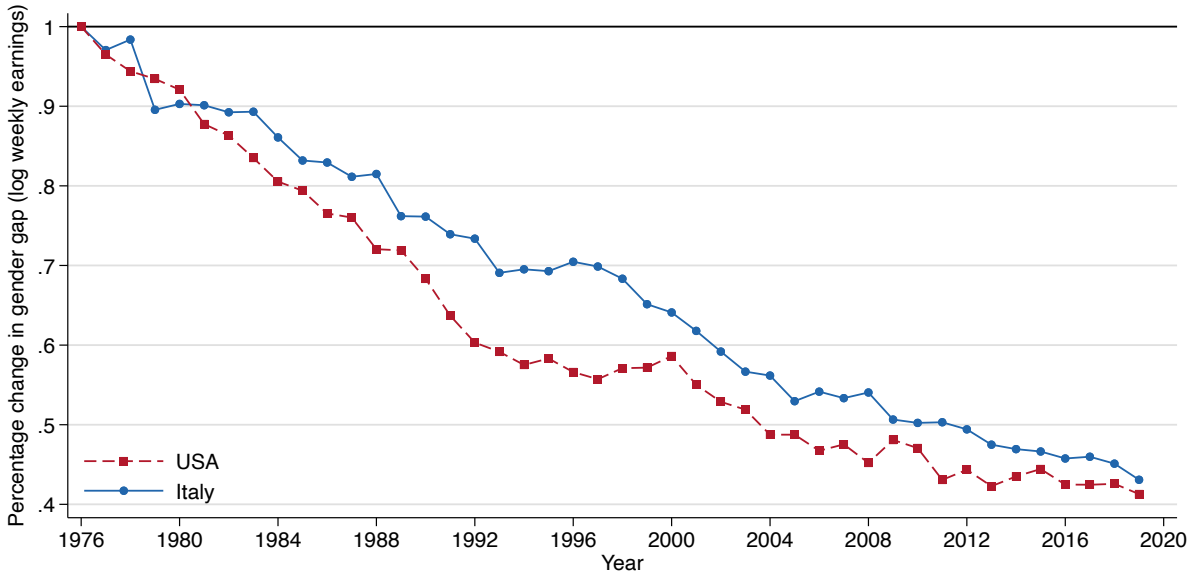
Figures and Tables

Figure 1: Gender Gap in Weekly Earnings

Panel A: Raw gap

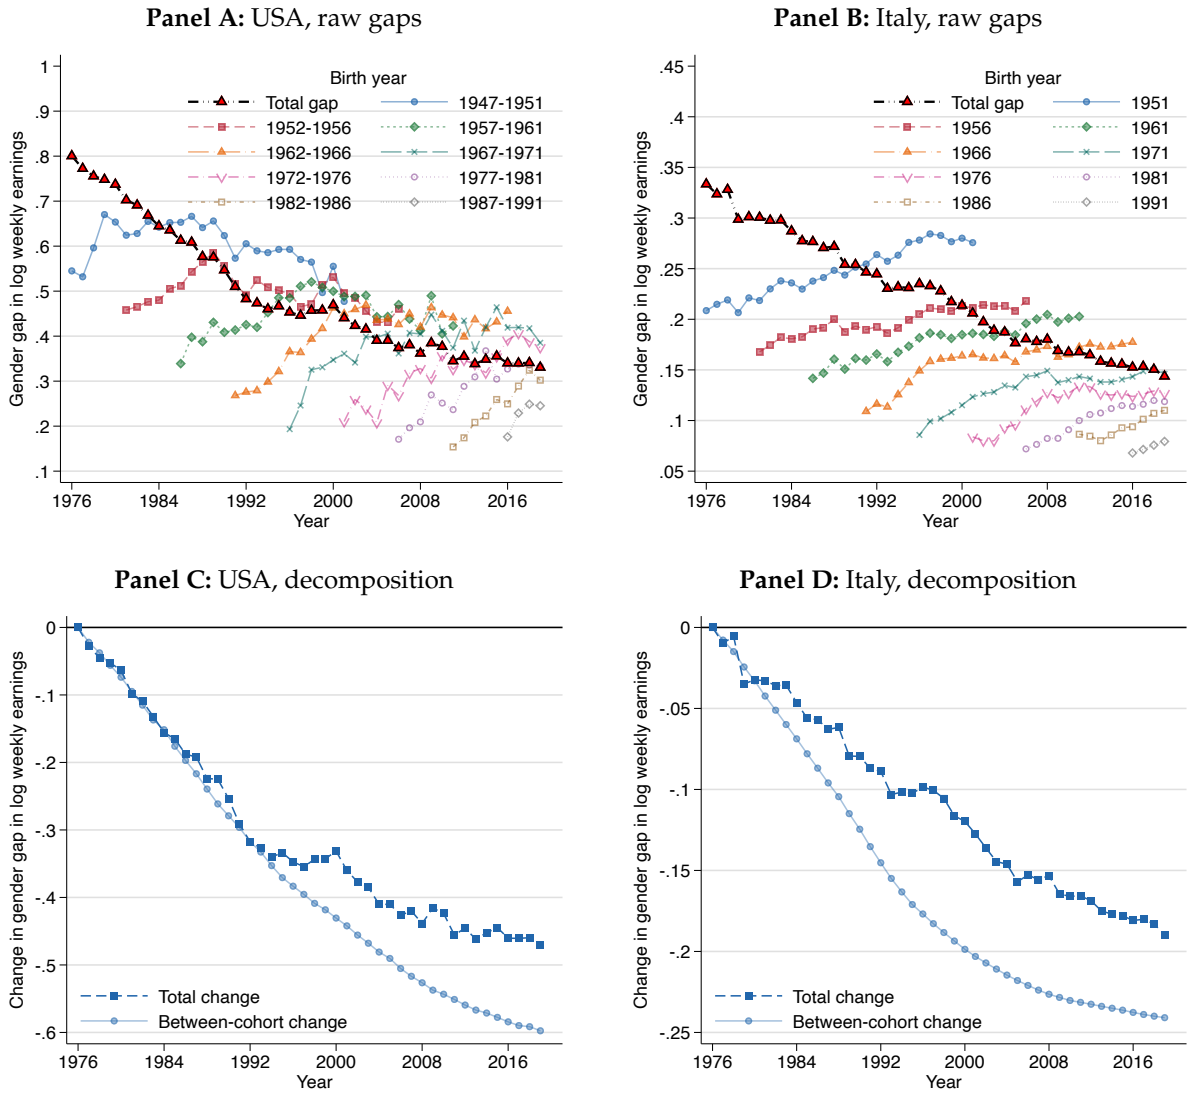


Panel B: Percentage deviation from first year



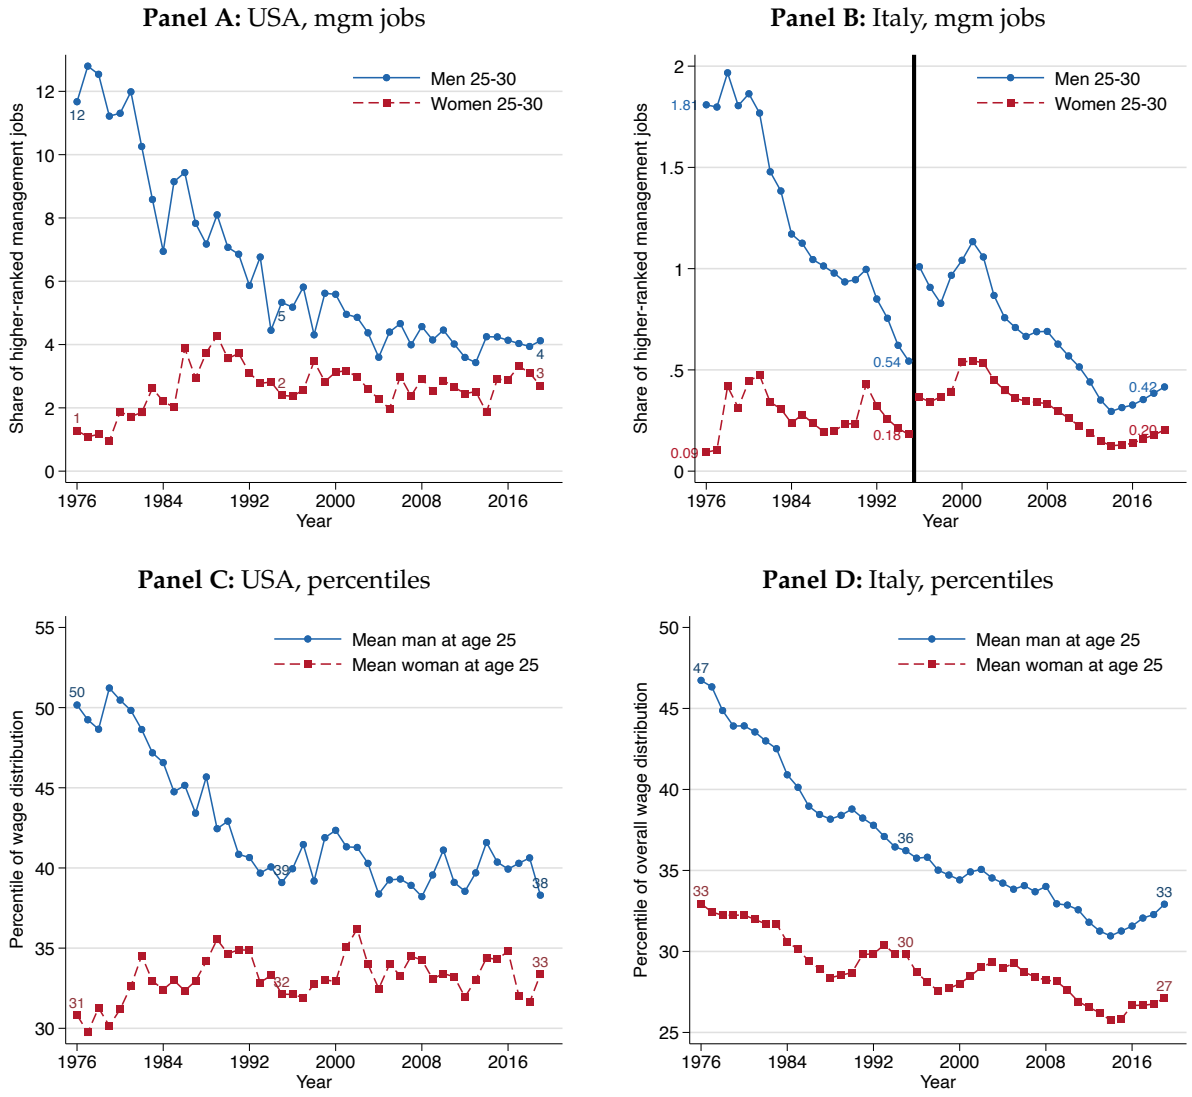
Notes: Panel A plots the trend in the raw mean gender pay gap (log weekly earnings of men - log weekly earnings of women) in Italy and the United States. Panel B shows the percentage deviation from the first year (1976). In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings, and (only in Italy) had not retired by December 31.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure 2: Gender Pay Gap Between and Within Cohorts



Notes: Panels A and B show the trend in the mean gender pay gap (log weekly earnings) across different birth cohorts in the United States and Italy, respectively. The red triangles depict the trend in the mean gender pay gap across all cohorts active in the labor market in each year. This analysis includes only workers aged 50 or younger to limit the influence of cross-cohort changes in selection into retirement. Panels C and D show the change in the total gender pay gap and its between-cohort component in the United States and Italy, respectively, for log weekly earnings. To compute the between-cohort component, we assign to each cohort (defined as a combination of birth year and gender) its mean log weekly earnings at labor-market entry in each year (Equation (2)). In the baseline analysis, entry into the labor market corresponds with the year in which workers in each cohort were 25 years old.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

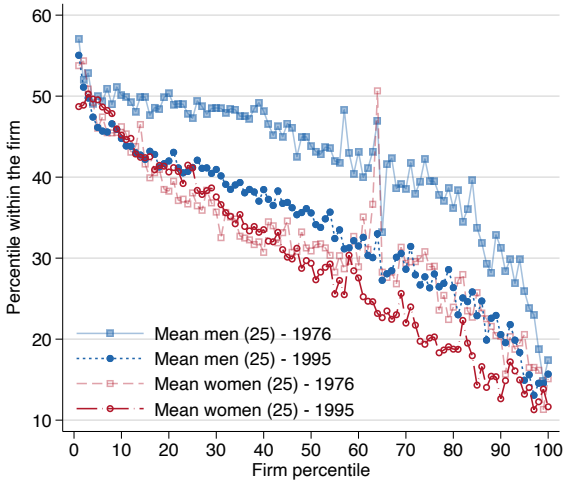
Figure 3: Younger Workers' Positions in Firms' Hierarchies and Pay Distribution



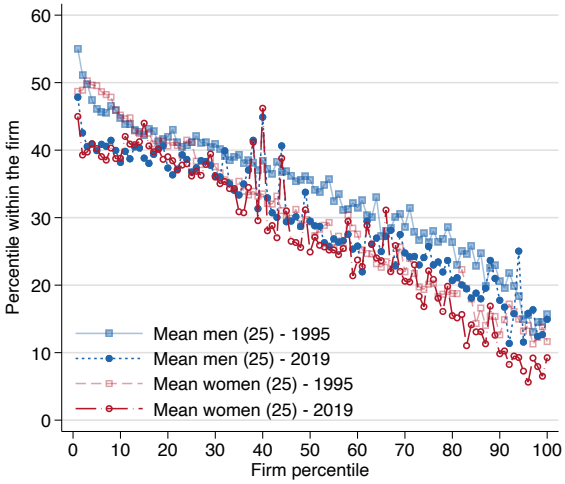
Notes: Panels A and B show the share of high-ranked managerial jobs held by men and women between 25 and 30 years old in the United States and Italy, respectively. We define as higher-ranked managerial jobs all managerial occupations with annual earnings in the top quartile of the year-specific distribution of annual earnings. In the CPS data, managerial occupations are identified using 2-digit Standard Occupational Classification (SOC) code 11. In the INPS data, we use the highest ranked position out of the four main job categories in the Italian labor system: in ascending order, these are apprenticeships, blue-collar jobs, white-collar jobs, and managerial jobs. The figures based on Italian data show a spurious trend discontinuity in the mid 1990s because the definition of managerial jobs in the INPS data changes from 1996. Panels C and D show the average earning percentile of men and women at 25 years old in the United States and Italy, respectively. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings, and (only in Italy) had not retired by December 31.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure 4: Distribution Within and Between Firms for Workers Age 25—Italy

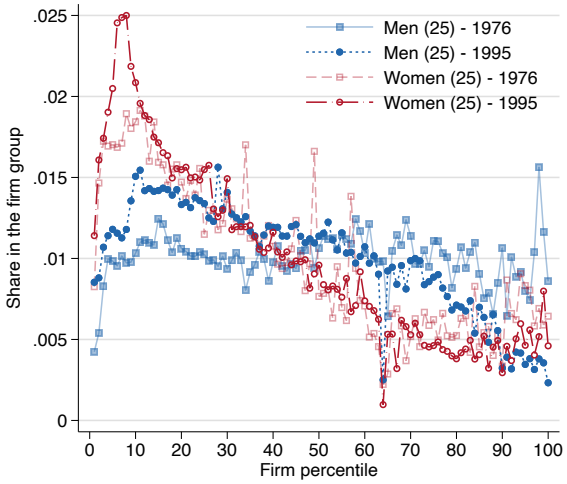
Panel A: Rank within firm, 1976 and 1995



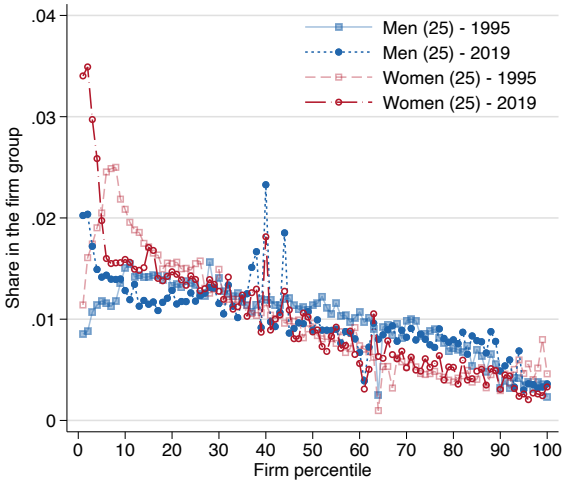
Panel B: Rank within firm, 1995 and 2019



Panel C: Share in firm percentile, 1976 and 1995

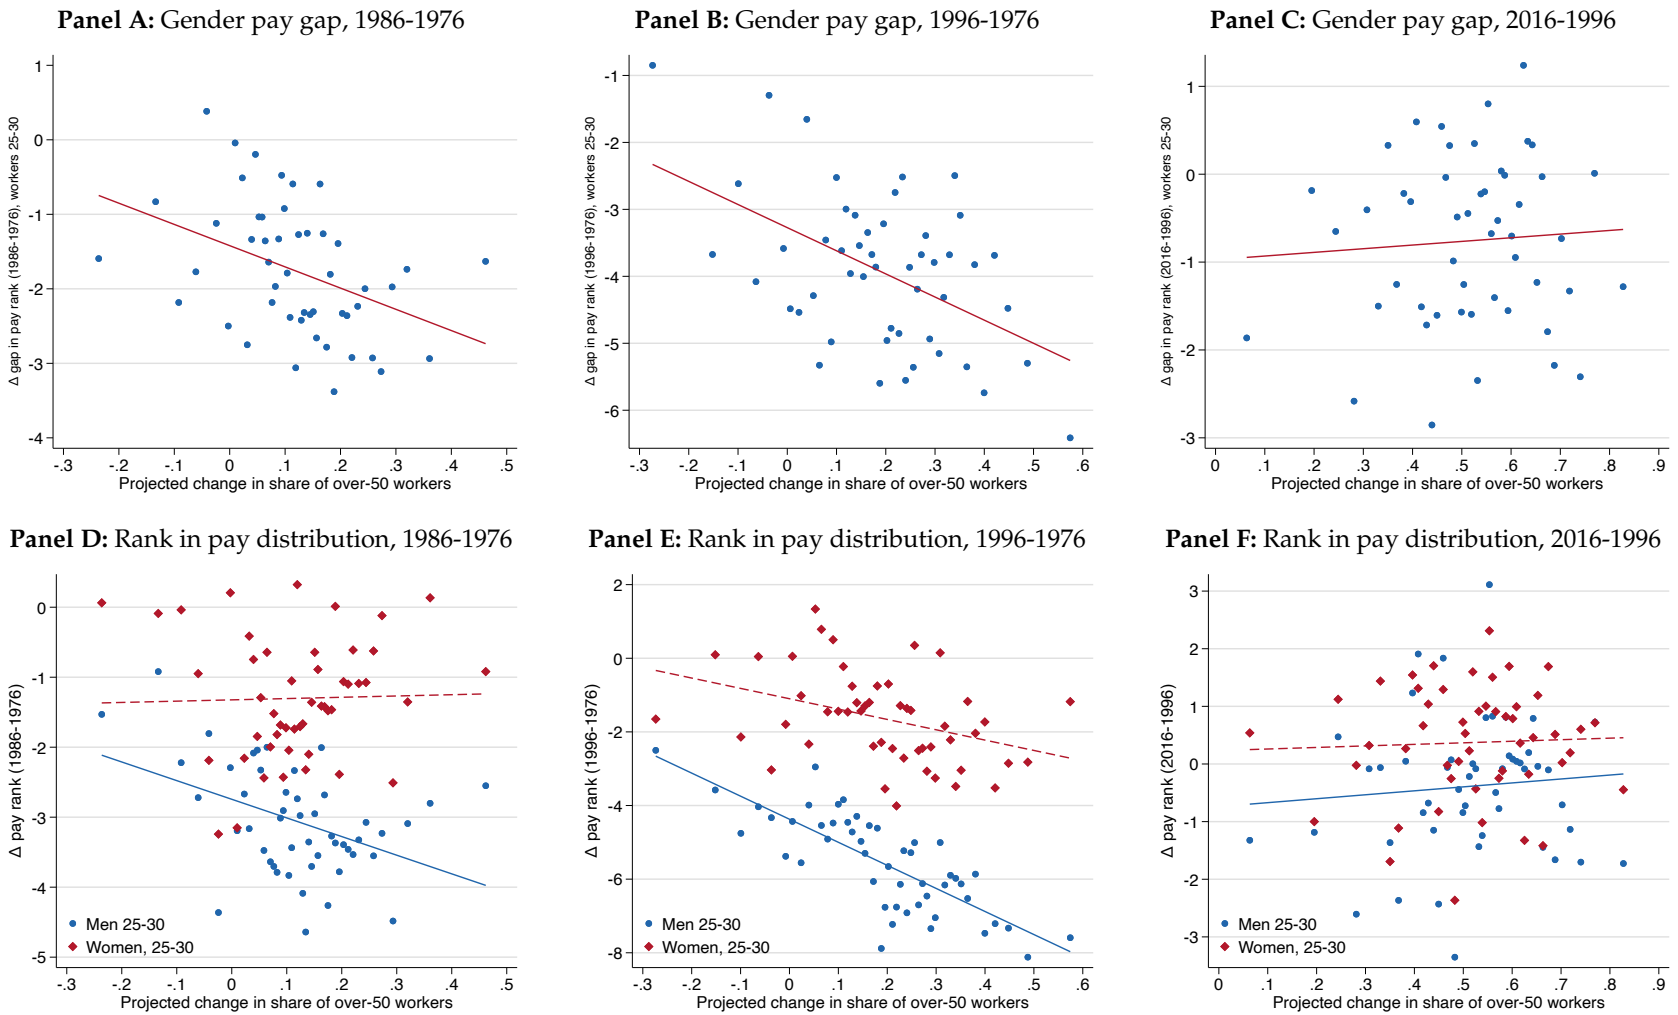


Panel D: Share in firm percentile, 1995 and 2019



Notes: Panels A and B show the average percentile, in the distribution of weekly earnings, of men and women at age 25 across percentiles of firm mean pay in 1976, 1995, and 2019. Panels C and D show the share of men and women at age 25 across percentiles of firm average pay in 1976, 1995, and 2019. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings, and had not retired by December 31.
Source: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

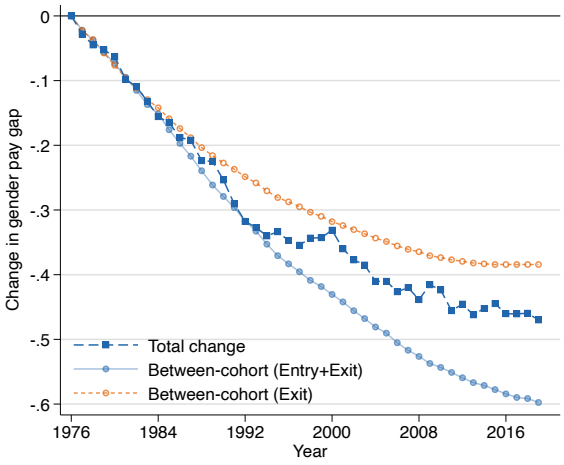
Figure 5: Firm-Level Workforce Aging and the Gender Pay Gap—Italy



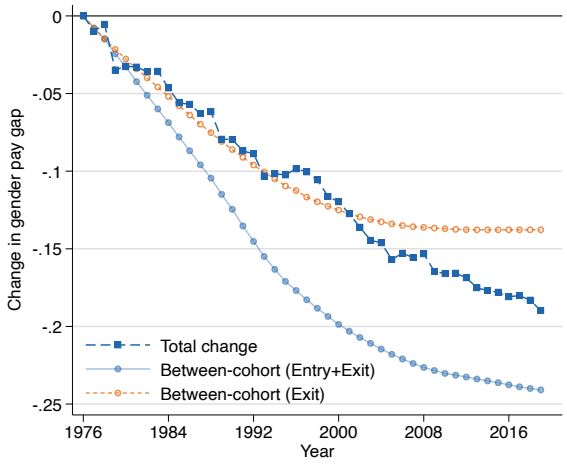
Notes: These figures show binned scatterplots that correlate firm-level changes in the gender pay gap of younger workers to firm-level projected changes in the share of workers over 50. Panel A: firm-level change in the gender gap in the pay rank among younger workers (25-30) from 1976 to 1986 on the y-axis and firm-level difference between the share of workers aged 41-50 and the share of workers aged 51-60 in 1976 on the x-axis. Panels B and C replicate the same analysis on the periods 1976-1996 and 1996-2016, respectively. Given that they cover 20-year periods, the projected change in the share of over-50 workers becomes the firm-level difference between the share of workers aged 31-40 and the share of workers aged 51-60 at baseline. Panel D: firm-level change in the pay rank (mean perc. in firm pay distribution) of younger men and women (25-30 years old) from 1976 to 1986 on the y-axis and the same variable used in Panel A on the x-axis. Panels E and F replicate the same analysis on the periods 1976-1996 and 1996-2016, respectively (changing the projected change in the share of over-50 workers accordingly). Firm-level values for year x are computed as three-year averages over x and $x+2$. In each panel, the sample of firms include all firms that had more than five total employees at baseline, at least one man and one woman under 30 years old in both the initial and final period, and at least one worker over 40 in the initial period. *Source:* UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

Figure 6: Between-Cohort Change Without Convergence in Entry Earnings

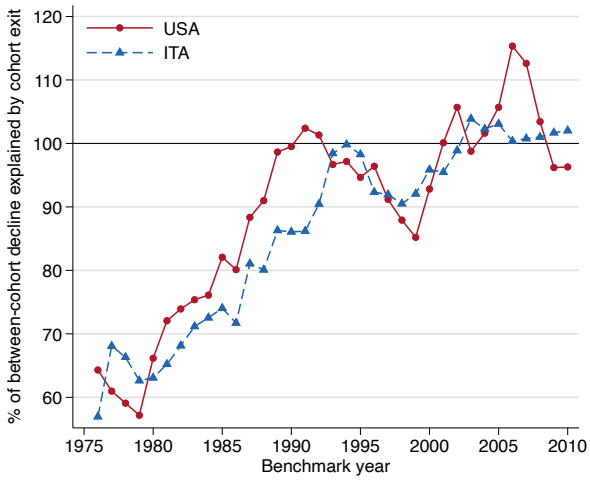
Panel A: United States



Panel B: Italy



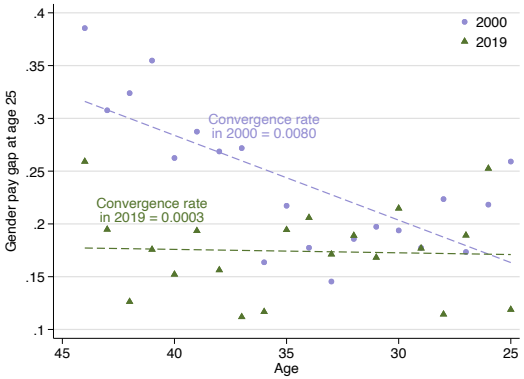
Panel C: Share of convergence at labor-market exit over total between-cohort convergence



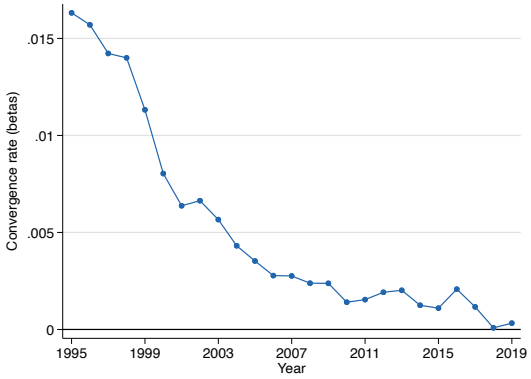
Notes: Panels A and B compute for the United States and Italy, respectively, what would have happened to the gender pay gap if the cross-cohort convergence in the earnings of men and women at age 25 had stopped in benchmark year $t_b = 1976$. In this counterfactual exercise, the weekly earnings of cohorts who entered the labor market after 1976 are set equal to the average weekly earnings of cohorts who entered the labor market between 1976 and the following two years. Panel C shows the ratio between the change (last year - first year) in the wage gap predicted by this new counterfactual scenario and the change in the gender gap predicted by the total between-cohort component from Equation (2) when the benchmark year t_b moves between 1976 and 2010. The ratio is such that 100 implies that the decline in the gender pay gap predicted by the new counterfactual scenario accounts for the entire between-cohort change. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings, and (only in Italy) had not retired by December 31.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure 7: Projected Convergence in the Gender Pay Gap

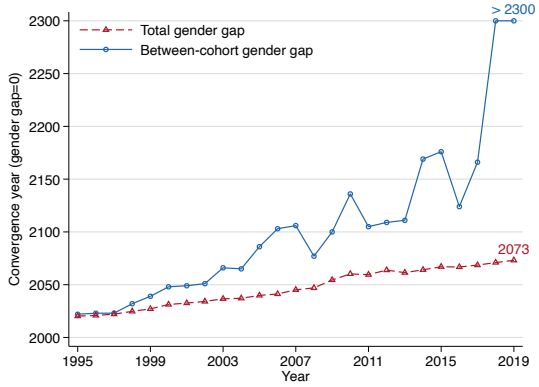
Panel A: Between-cohort trends (USA, log weekly earnings)



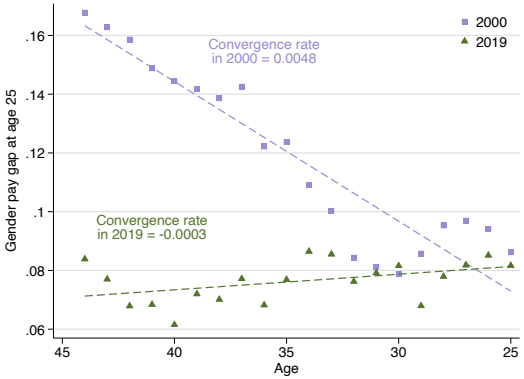
Panel B: Between-cohort convergence rate (USA, log weekly earnings)



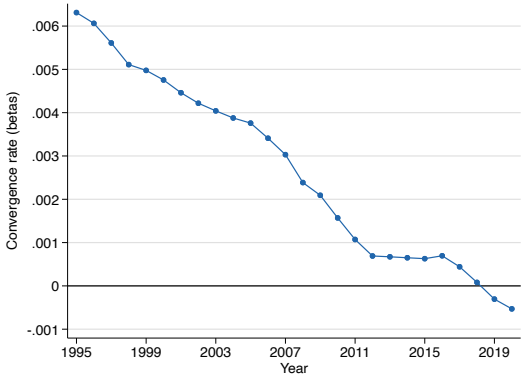
Panel C: Between-cohort convergence year (USA, log weekly earnings)



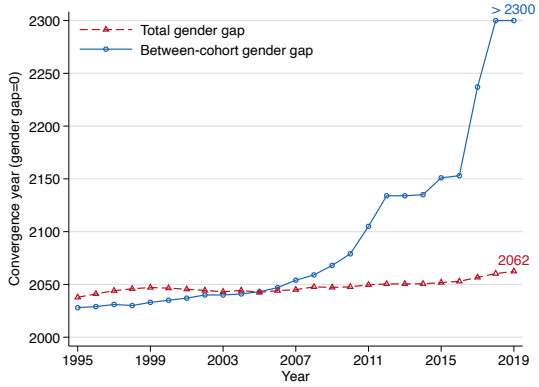
Panel D: Between-cohort trends (Italy, log weekly earnings)



Panel E: Between-cohort convergence rate (Italy, log weekly earnings)



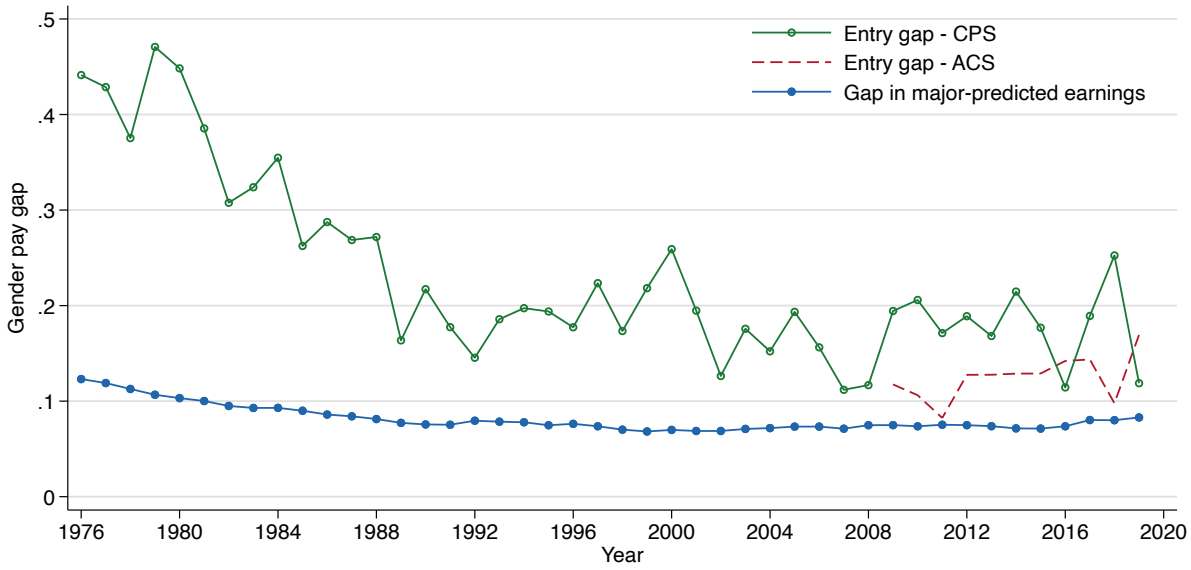
Panel F: Between-cohort convergence year (Italy, log weekly earnings)



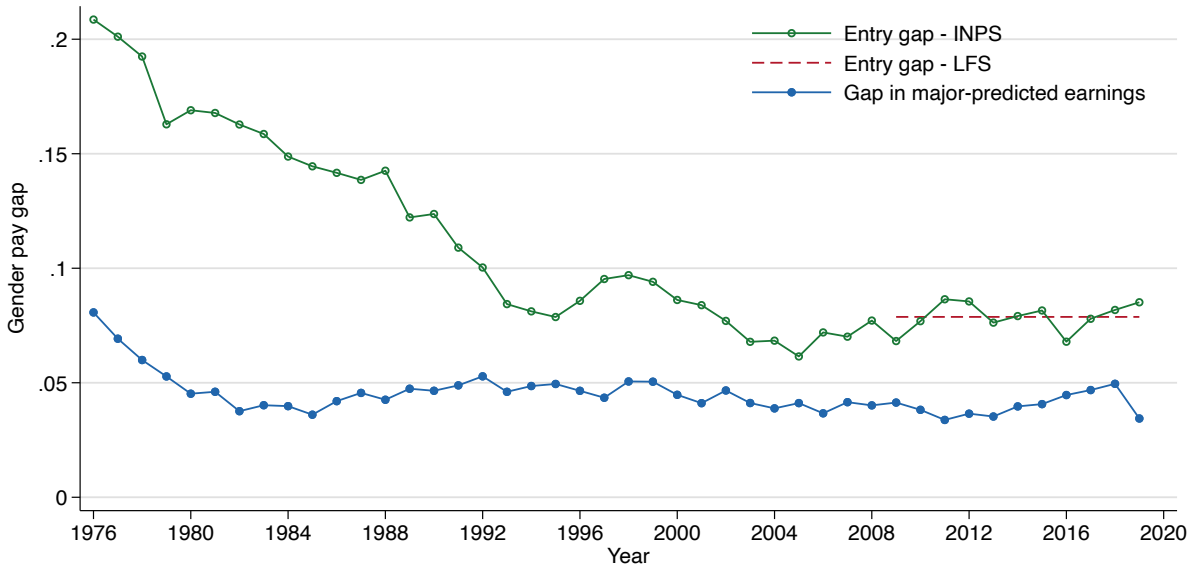
Notes: In each year t, we compute the gender gap at an early career stage for workers in age group a using their earnings at age 25. Then, we estimate the linear relationship between the mean gender gap at labor-market entry and age (Equation (3)). Panels A and D show the best fit line in 2000 and 2019 for the United States and Italy, respectively. Panels B and E show the coefficients of age (β_t) for each year t. Panels C and F show the first year of convergence in the gender gap (first cohort with gap at most equal to zero at age 25) predicted by Equation (3) for each year between 1976 and 2019. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings, and (only in Italy) had not retired by December 31.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States:* Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure 8: Major-Predicted Gender Pay Gap

Panel A: USA



Panel B: Italy



Notes: The figure shows the gender gap in major-predicted residual weekly earnings computed from the American Community Survey in the United States (Panel A) and from the Labor Force Survey in Italy (Panel B). Predicted residual weekly earnings are computed by averaging across college majors the residual weekly earnings of native-born male employees working full-time in the private sector. Residual weekly earnings are obtained from OLS log wage regressions that control for a quadratic polynomial in age and time fixed effects over the period 2009-2019 in both countries. Then, we multiply these major-specific residuals by cohort-gender shares in each major. There are 32 college majors in the data for Italy and 176 for the United States. The gender gap at entry (age 25) is measured with both CPS and ACS (for college graduates) data in the United States and with both INPS and LFS (for college graduates) data in Italy. In the last case, we report the average between 2009 and 2019 because yearly averages are too noisy.
Sources for Italy: Quarterly Labour Force Survey, Istituto Nazionale di Statistica (ISTAT); UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Sources for the United States*: Integrated Public Use Microdata Series, American Community Survey; Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Table 1: Firm-Level Workforce Aging and the Gender Pay Gap

	Gap in pay rank between men and women at age 25-30					
	OLS		Reduced form		2SLS	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 1976-1986						
△ share older workers	-2.76**	-2.93***	-2.70***	-2.44**	-4.66***	-4.18**
	(1.08)	(1.09)	(0.95)	(0.98)	(1.64)	(1.69)
KP F-stat					7,331	11,166
Mean dep. var.	-1.77	-1.77	-1.77	-1.77	-1.77	-1.77
SD dep. var.	19.94	19.94	19.94	19.94	19.94	19.94
Obs.	25,536	25,536	25,536	25,536	25,536	25,536
Panel B: 1976-1996						
△ share older workers	-5.14***	-5.37***	-3.54***	-3.77***	-7.85***	-8.28***
	(1.78)	(1.80)	(1.28)	(1.25)	(2.80)	(2.72)
KP F-stat					4,597	4,392
Mean dep. var.	-3.92	-3.92	-3.92	-3.92	-3.92	-3.92
SD dep. var.	20.41	20.41	20.41	20.41	20.41	20.41
Obs.	14,064	14,064	14,064	14,064	14,064	14,064
Panel C: 1996-2016						
△ share older workers	-0.37	-0.35	-0.47	-0.46	-1.21	-1.16
	(1.42)	(1.39)	(1.25)	(1.32)	(3.24)	(3.33)
KP F-stat					2,245	3,397
Mean dep. var.	-0.77	-0.77	-0.77	-0.77	-0.77	-0.77
SD dep. var.	20.19	20.19	20.19	20.19	20.19	20.19
Obs.	14,709	14,709	14,709	14,709	14,709	14,709
Controls	No	Yes	No	Yes	No	Yes

Notes: This table shows results of OLS and IV firm-level regressions in which changes in the gender pay gap among younger workers (difference in mean pay rank between men and women between 25 and 30 years old) are regressed on changes in exposure to workforce aging. The "OLS" columns regress Δ gender gap in pay rank on Δ share of workers 51-60, both observed at the firm level. The "Reduced form" columns regress Δ gender pay rank gap on the difference between the share of workers aged 41-50 (Panel A) or 31-40 (Panels B and C) and the share of workers 51-60 at baseline. The "2SLS" columns regress Δ gender pay rank gap on Δ share of workers 51-60, instrumenting the latter with the difference between the share of workers aged 41-50 (Panel A) or 31-40 (Panels B and C) and the share of workers 51-60 at baseline. Controls include firm age, firm size, mean firm pay, and province fixed effects, all observed at baseline. Firm-level values for year x are computed as three-year averages over x and $x + 2$. In each panel, the sample of firms include all firms that had more than five total employees at baseline, at least one man and one woman under 30 years old in both the initial and final period, as well as at least one worker over 40 in the initial period.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

Table 2: Firm-Level Workforce Aging and Younger Workers' Rank

	Men, 25-30				Women, 25-30			
	Reduced form		2SLS		Reduced form		2SLS	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: 1976-1986								
△ share older workers	-2.64***	-3.34***	-4.57***	-5.75***	0.05	-0.65	0.09	-1.09
	(0.80)	(0.70)	(1.38)	(1.20)	(0.85)	(0.79)	(1.47)	(1.35)
KP F-stat			3,665	3,665			3,665	3,665
Mean dep. var.	-2.16	-2.16	-2.16	-2.16	-2.16	-2.16	-2.16	-2.16
SD dep. var.	16.51	16.51	16.51	16.51	16.51	16.51	16.51	16.51
Obs.	51,072	51,072	51,072	51,072	51,072	51,072	51,072	51,072
Panel B: 1976-1996								
△ share older workers	-6.15***	-1.74**	-13.60***	-3.86**	-2.60*	1.80*	-5.76*	3.98*
	(1.12)	(0.81)	(2.57)	(1.76)	(1.50)	(1.08)	(3.35)	(2.39)
KP F-stat			2,298	2,298			2,298	2,298
Mean dep. var.	-3.61	-3.61	-3.61	-3.61	-3.61	-3.61	-3.61	-3.61
SD dep. var.	18.30	18.30	18.30	18.30	18.30	18.30	18.30	18.30
Obs.	28,128	28,128	28,128	28,128	28,128	28,128	28,128	28,128
Panel C: 1996-2016								
△ share older workers	-1.77	-1.52	-4.59	-3.86	-1.30	-1.06	-3.38	-2.65
	(1.09)	(0.99)	(2.88)	(2.55)	(1.21)	(1.06)	(3.18)	(2.71)
KP F-stat			1,122	1,122			1,122	1,122
Mean dep. var.	-0.02	-0.02	-0.02	-0.02	-0.02	-0.02	-0.02	-0.02
SD dep. var.	17.92	17.92	17.92	17.92	17.92	17.92	17.92	17.92
Obs.	29,418	29,418	29,418	29,418	29,418	29,418	29,418	29,418
Controls	No	Yes	No	Yes	No	Yes	No	Yes

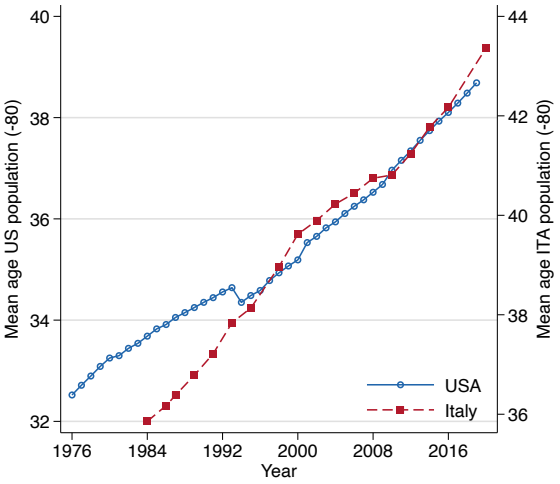
Notes: This table shows results of OLS and IV firm-level regressions in which changes in the mean percentile of younger men and women (aged 25-30) in their firms' pay distributions are regressed on changes in firms' exposure to workforce aging. The "Reduced form" columns regress Δ mean pay rank of younger men and women on firm-level differences between the share of workers aged 41-50 (Panel A) or 31-40 (Panels B and C) and the share of workers 51-60 at baseline. The "2SLS" columns regress firm-level Δ mean pay rank of younger men and women on firm-level Δ share of workers 51-60, instrumenting the latter with differences between the share of workers aged 41-50 (Panel A) or 31-40 (Panels B and C) and the share of workers 51-60 at baseline. Controls include firm age, firm size, mean firm pay, and province fixed effects, all observed at baseline. Firm-level values for year x are computed as three-year averages over x and $x + 2$. In each panel, the sample of firms include all firms that had more than five total employees at baseline, at least one man and one woman under 30 years old in both the initial and final period, as well as at least one worker over 40 in the initial period.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

Online Appendix

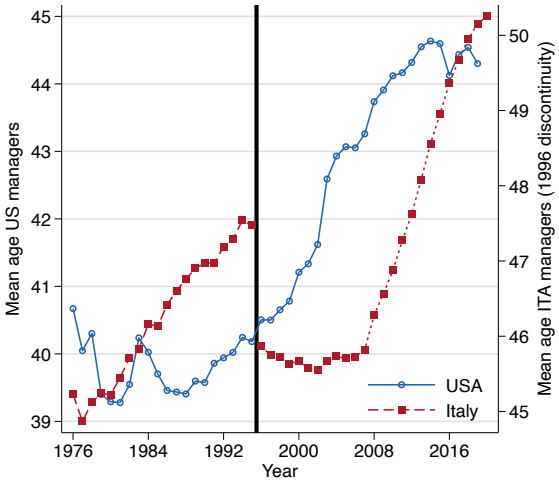
A Additional Figures and Tables

Figure A1: Workforce Aging and GDP Growth

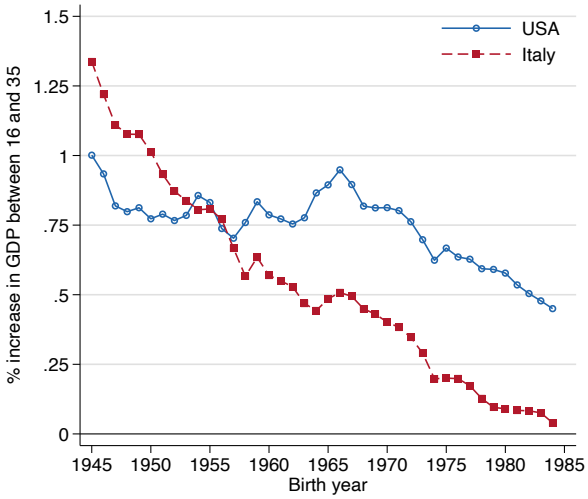
Panel A: Population



Panel B: Managers



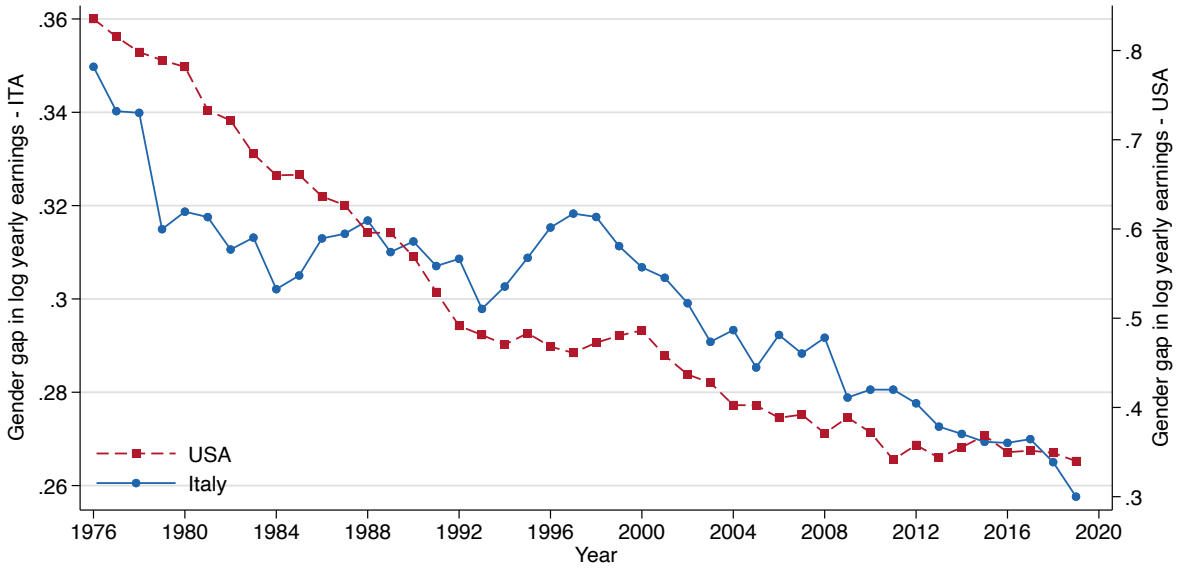
Panel C: GDP growth



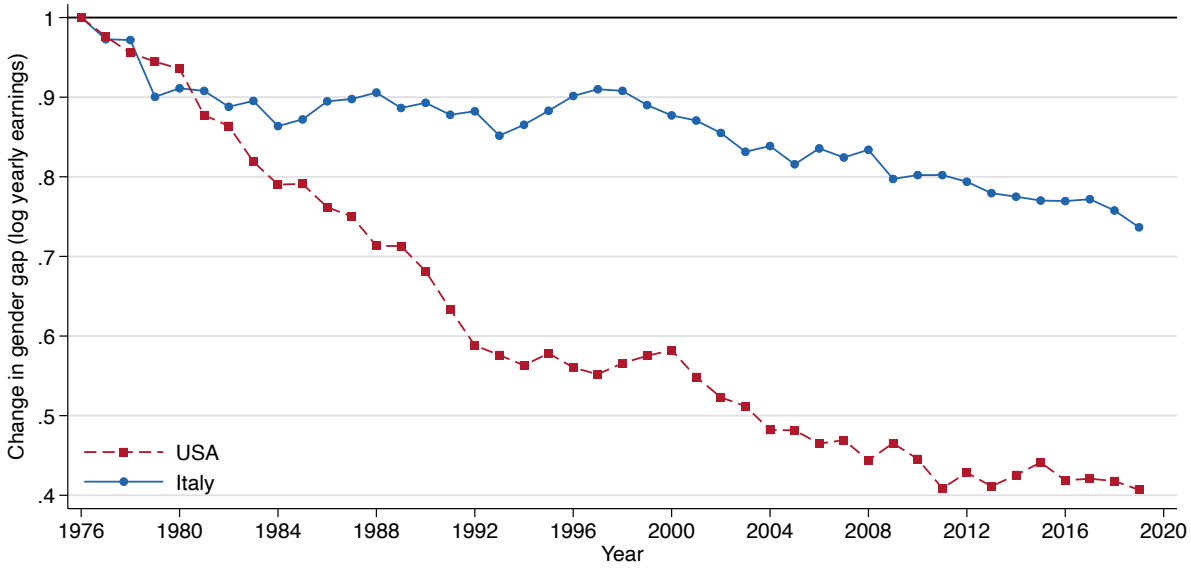
Notes: Panels A shows the mean age of the population. Panel B shows the mean age of managers in the private sector. Managerial positions are identified using (i) 2-digit Standard Occupational Classification (SOC) code 11 in the United States and (ii) the highest ranked position in the Italian labor system in Italy. The figures based on Italian data show a spurious trend discontinuity in the mid 1990s because the definition of managerial jobs in the INPS data changes from 1996. All age variables are winsorized at age 80. Panel C computes the cumulative percentage change in GDP (in 2010 USD) over the first years in the labor market for individuals born in different years. For example, the data point for the birth year "1945" computes the percentage growth in GDP between 1961 (when individuals born in 1945 were 16 years old) and 1980 (when individuals born in 1945 were 35 years old). Sources for Italy: Survey of Household Income and Wealth, Bank of Italy (Panel A). UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS) (Panel B). Source for the United States: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS. Source for GDP data: World Development Indicators by the World Bank, last accessed on 04/21/2023 at https://databank.worldbank.org/reports.aspx?source=2&series=NY.GDP.MKTP.CD&country=.

Figure A2: Gender Gap in Yearly Labor Earnings

Panel A: Raw gap



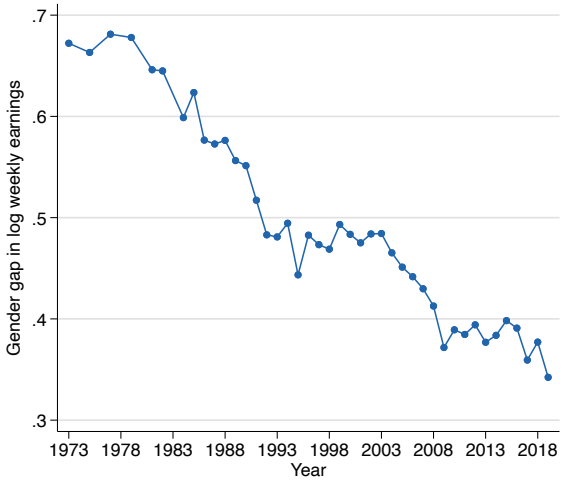
Panel B: Percentage deviation from first year



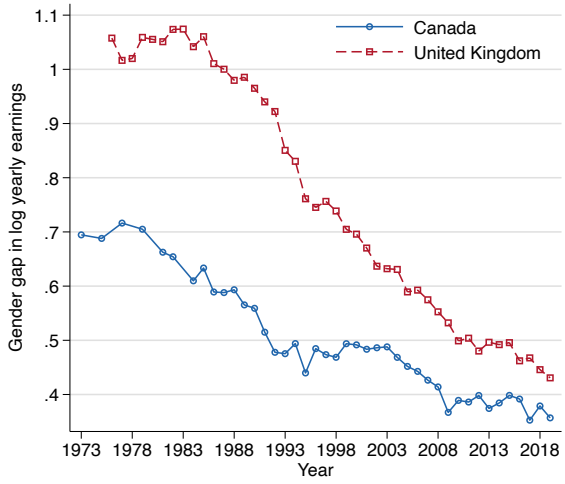
Notes: Panel A plots the trend in the raw mean gender gap (log yearly earnings of men - log yearly earnings of women) in Italy and the United States. Panel B shows the percentage deviation from the first year (1976). In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings, and (only in Italy) had not retired by December 31.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A3: Gender Pay Gap in Other High-Income Countries

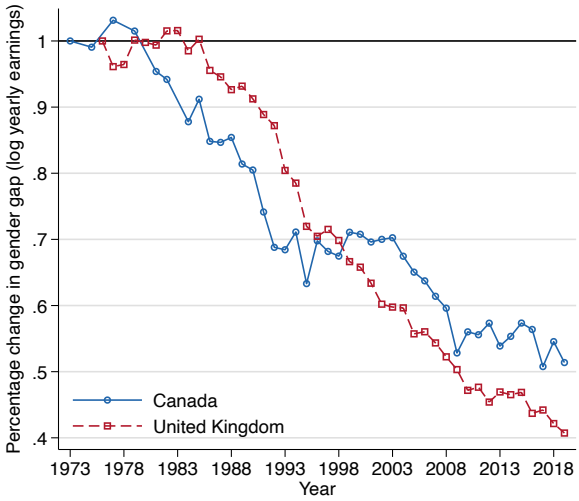
Panel A: Log weekly earnings (Canada)



Panel B: Log yearly earnings

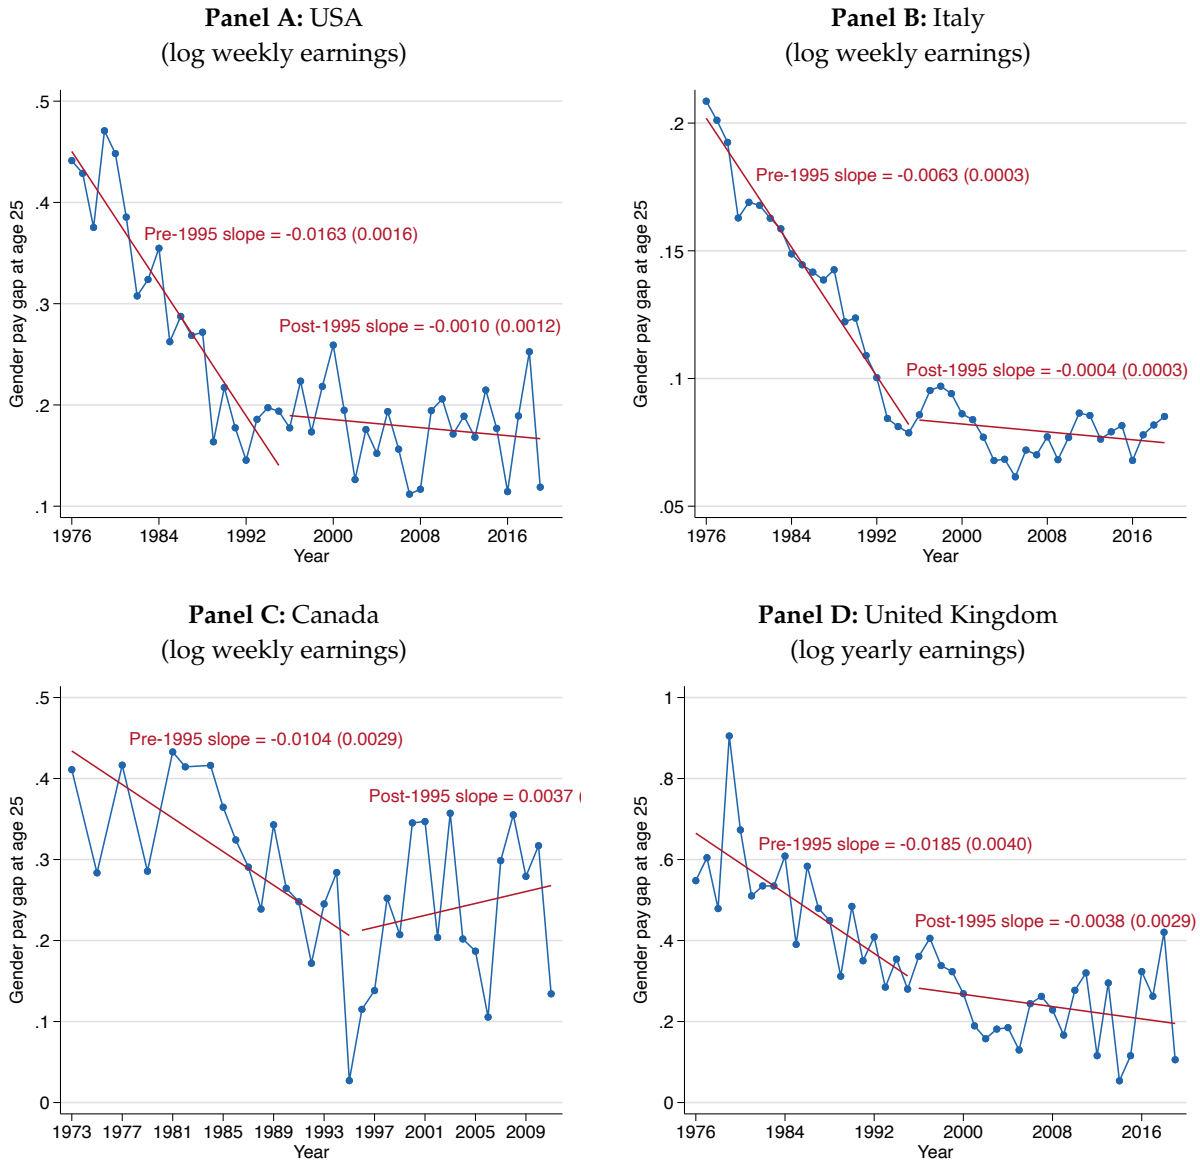


Panel C: Percentage deviation from first year (log yearly earnings)



Notes: Panel A plots the trend in the raw mean gender gap in log weekly earnings for Canada. Panel B plots the trend in the raw mean gender gap in log yearly earnings, which are available in Canada and the United Kingdom. Panel C shows the percentage deviation from the first in the gender gap in log yearly earnings. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks (only in Canada), and had earned strictly positive earnings.
Source for LIS data: Luxembourg Income Study (LIS) Database, last accessed on 06/01/2023 at https://www.lisdatacenter.org/.

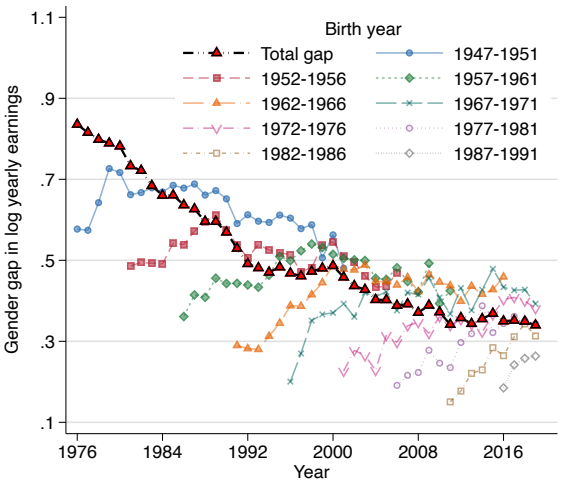
Figure A4: Gender Pay Gap at Age 25



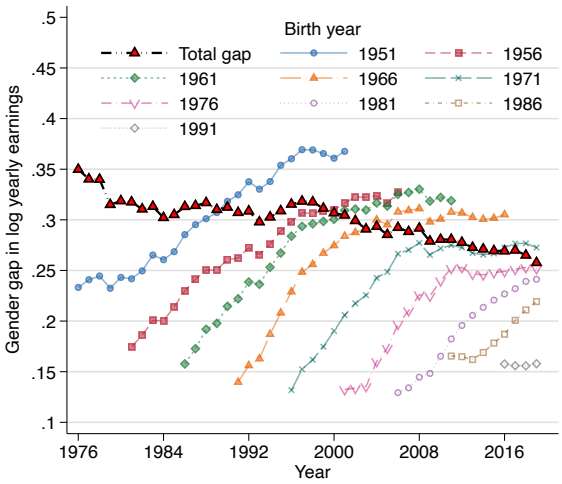
Notes: Panels A to D show the trend in the mean gender gap in log earnings at age 25 in the United States, Italy, Canada, and the United Kingdom, respectively. In each year, the data encompass information about all workers who were 25 years old, had worked in the private sector for at least 24 weeks (not available in the UK), had earned strictly positive earnings, and (only in Italy) had not retired by December 31.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS. *Source for LIS data:* Luxembourg Income Study (LIS) Database, last accessed on 06/01/2023 at https://www.lisdatacenter.org/.

Figure A5: Raw Cross-Cohort Trends in Gender Pay Gap, More Countries and Earnings Measures

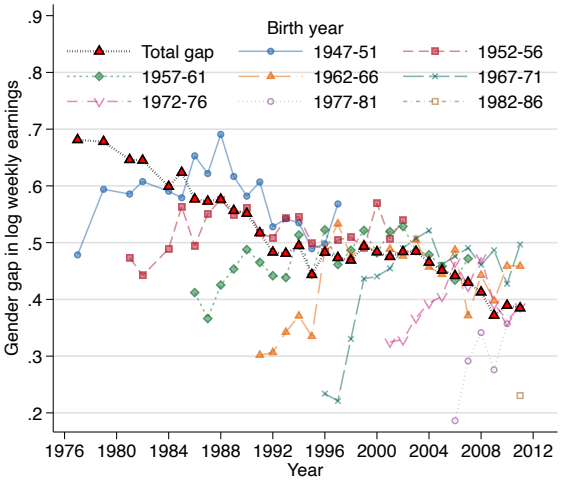
Panel A: USA
(log yearly earnings)



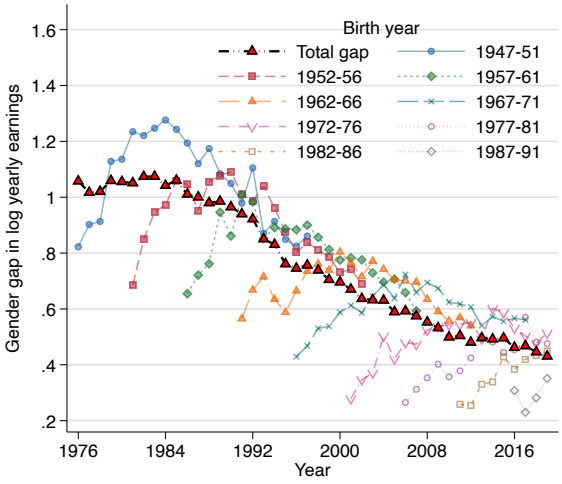
Panel B: Italy
(log yearly earnings)



Panel C: Canada
(log weekly earnings)

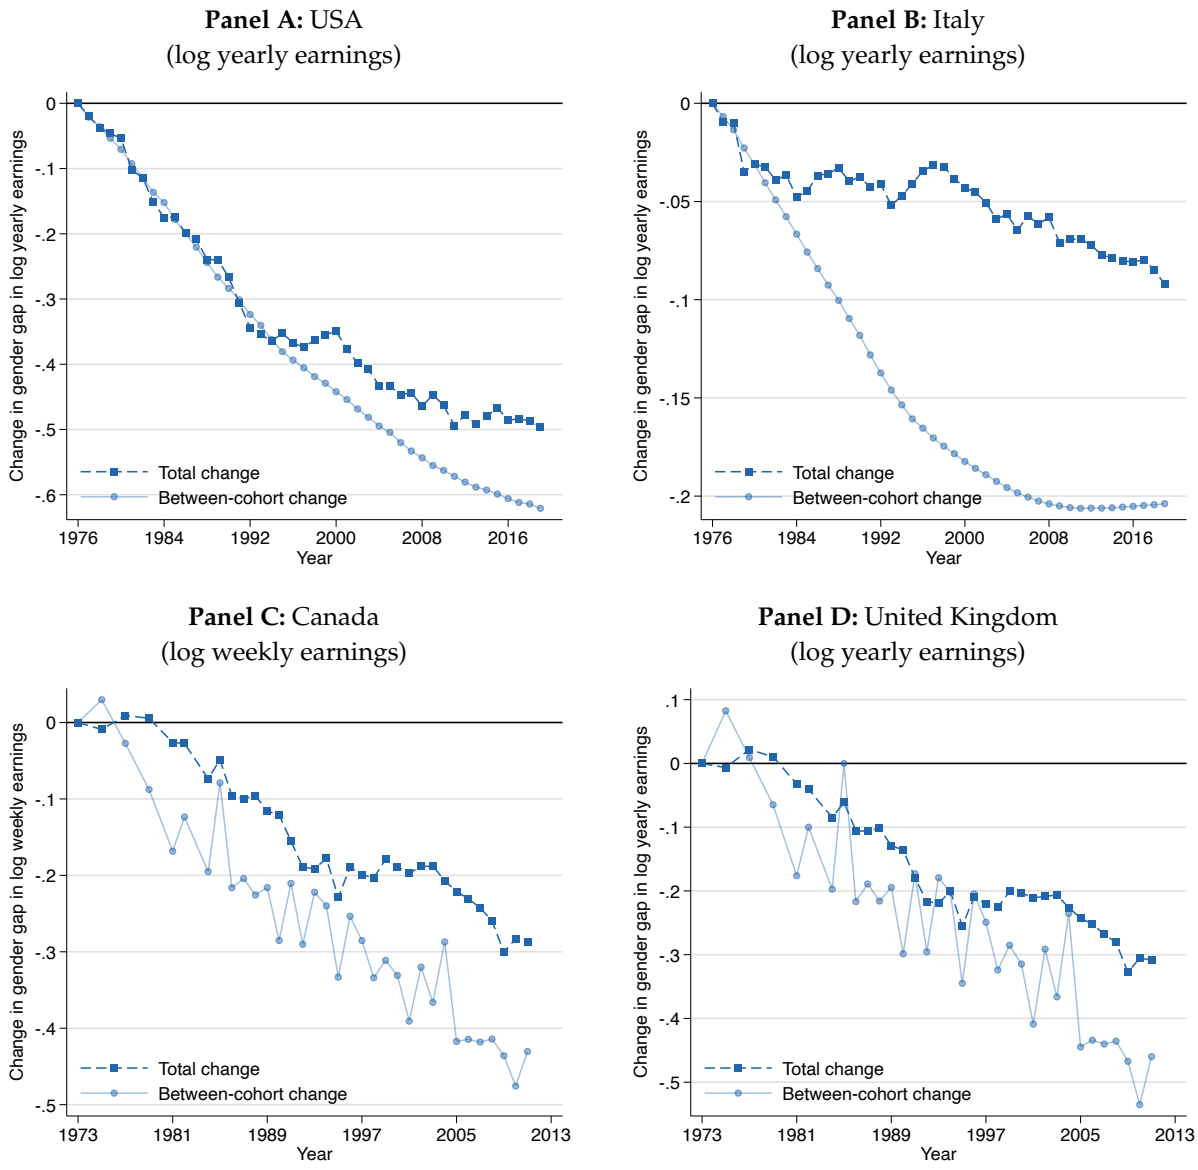


Panel D: United Kingdom
(log yearly earnings)



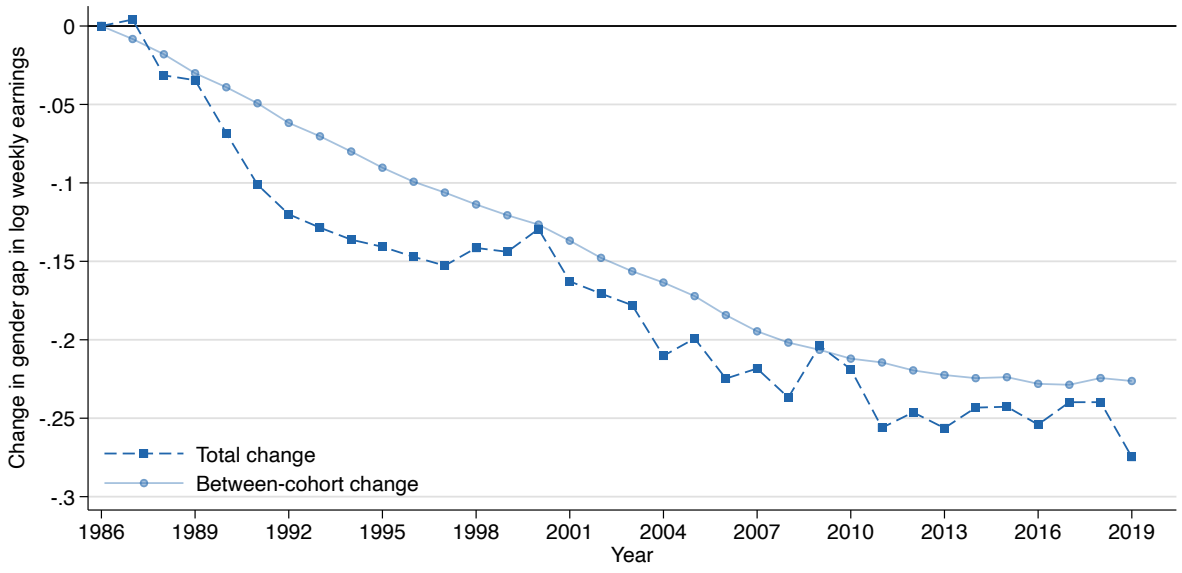
Notes: Panels A, B, and D show the trend in the mean gender gap in log yearly earnings across different birth cohorts in the United States, Italy, and the United Kingdom, respectively. Panel C shows the cross-cohort trend in the mean gender gap in log weekly earnings in Canada. The red triangles depict the trend in the mean gender pay gap across all cohorts active in the labor market in each year. This analysis includes only workers aged 50 or younger to limit the influence of cross-cohort changes in the selection into retirement. In each year, the data encompass information about all workers who were between 25 and 50 years old, had worked in the private sector for at least 24 weeks (not available in the UK), had earned strictly positive earnings, and (only in Italy) had not retired by December 31.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS. *Source for LIS data:* Luxembourg Income Study (LIS) Database, last accessed on 06/01/2023 at https://www.lisdatacenter.org/.

Figure A6: Between-Cohort Decline in Gender Pay Gap, More Countries and Earnings Measures



Notes: Panels A, B and D show the change in the total gender gap and its between-cohort component in the United States, Italy, and the United Kingdom, respectively for log yearly earnings. Panel C shows the change in the total gender gap and its between-cohort component in Canada for log weekly earnings. To compute the between-cohort component, we assign to each cohort (defined as a combination of birth year and gender) its mean log (yearly or weekly) earnings at labor-market entry in each year (Equation (2)). In the baseline analysis, entry in the labor market corresponds with the year in which workers in each cohort were 25 years old. We assign to cohorts who were older than 25 in the first sample year their mean weekly earnings in the first sample year. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks (not available in the UK), had earned strictly positive earnings.
Source for LIS data: Luxembourg Income Study (LIS) Database, last accessed on 06/01/2023 at https://www.lisdatacenter.org/.

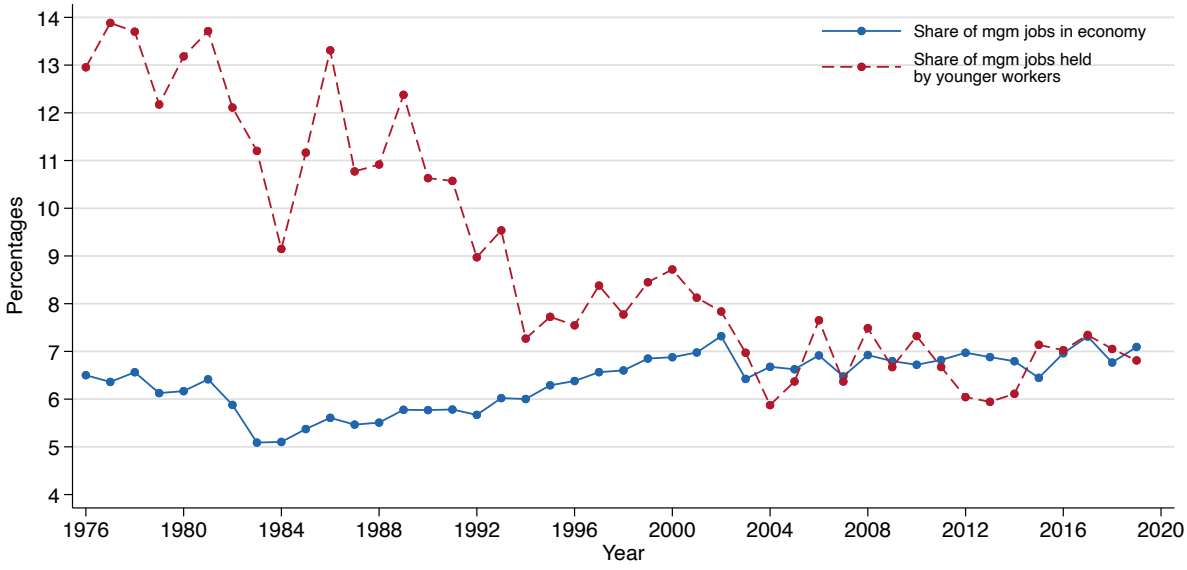
Figure A7: Between-Cohort Decline in Gender Pay Gap, Augmenting CPS with US Decennial Censuses



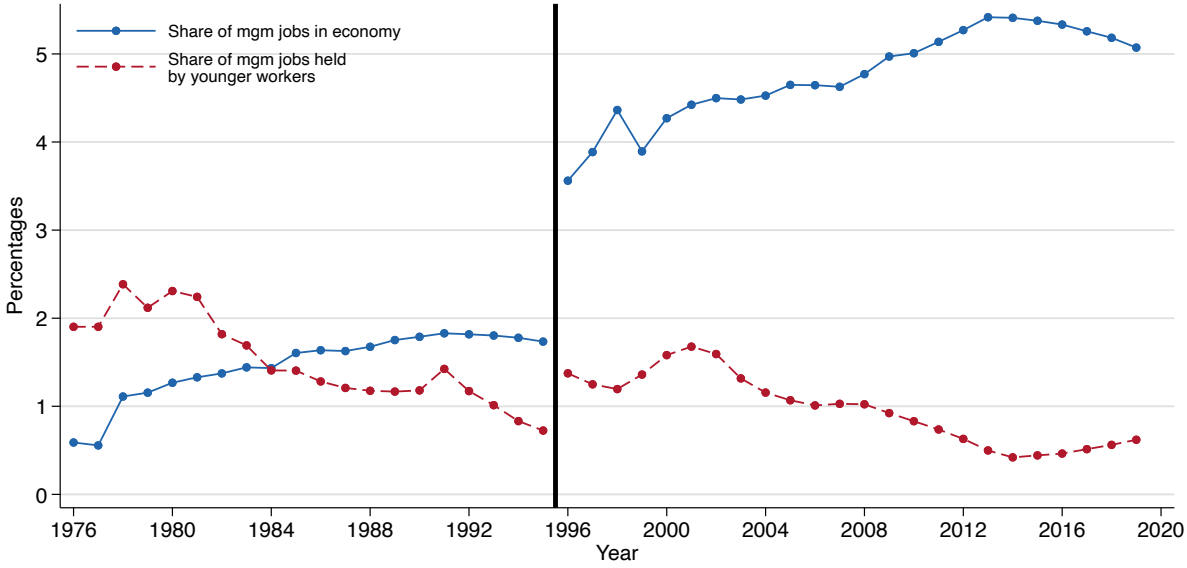
Notes: Change in the total gender pay gap and its between-cohort component in the United States, for log weekly earnings, 1986–2019, workers ages 25–50. We add Decennial Censuses data to augment the CPS data starting in 1976. First, we compute average log weekly earnings at age 25, separately by men and women, by year (1940, 1950, 1960, 1970, 1976–2019). We then linearly interpolate the male and female age-25 series for the years with missing data, effectively assigning gender-specific average earnings at age 25 for all cohorts from the 1915 birth cohort. onwards We then use these averages to compute the decomposition in Equation (2). We restrict attention to ages 25-50 and years 1986-2019 to limit the influence of interpolated averages in the decomposition.
Source for the United States. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS. IPUMS USA: Version 15.0 [dataset]. Minneapolis, MN: IPUMS, 2024.

Figure A8: Share of Higher-Ranked Managerial Positions in Economy

Panel A: USA

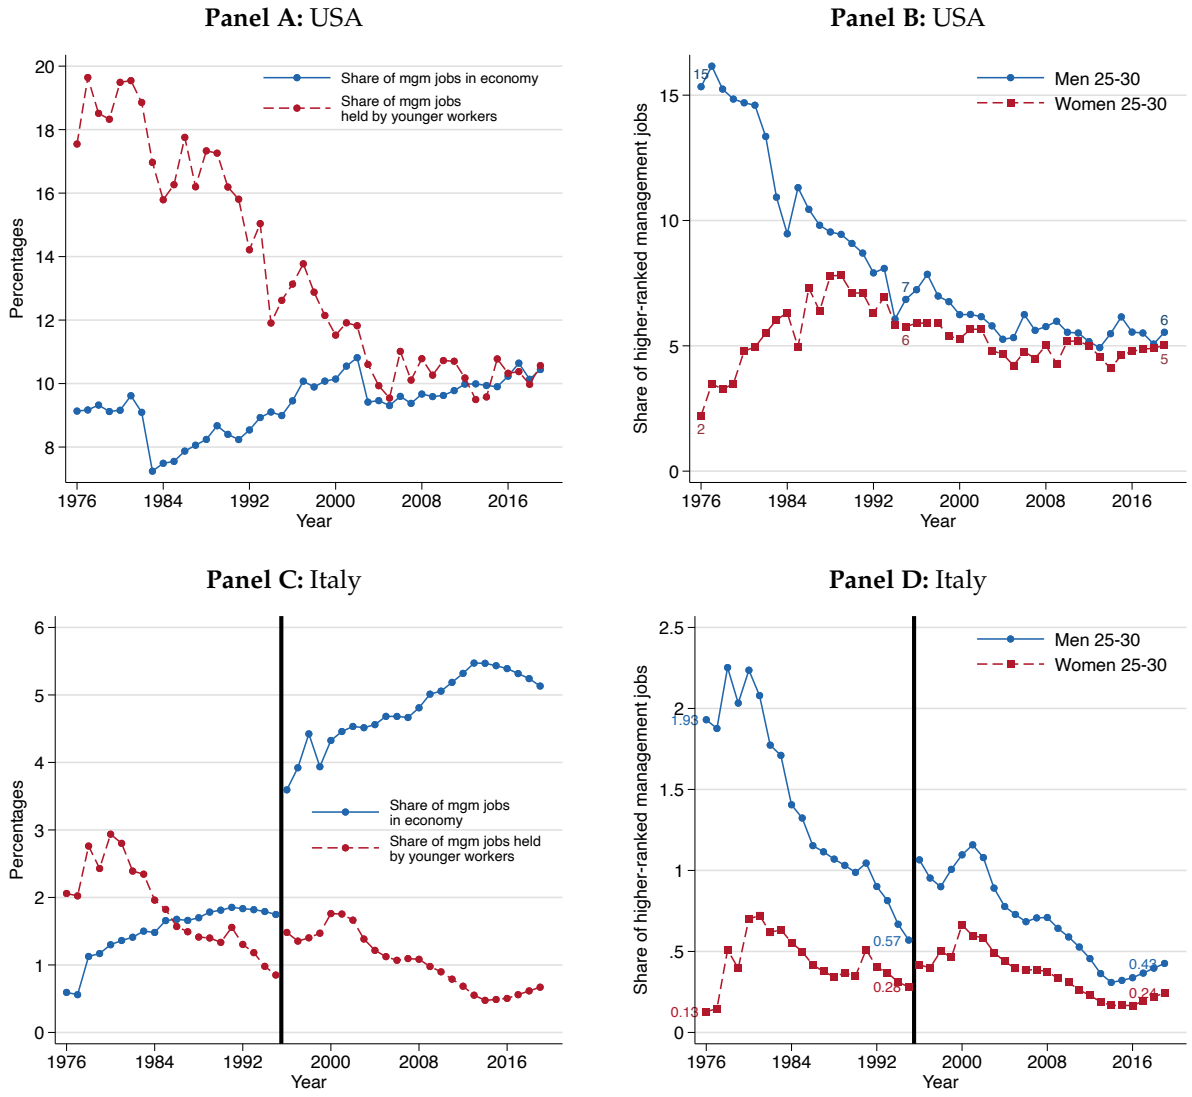


Panel B: Italy



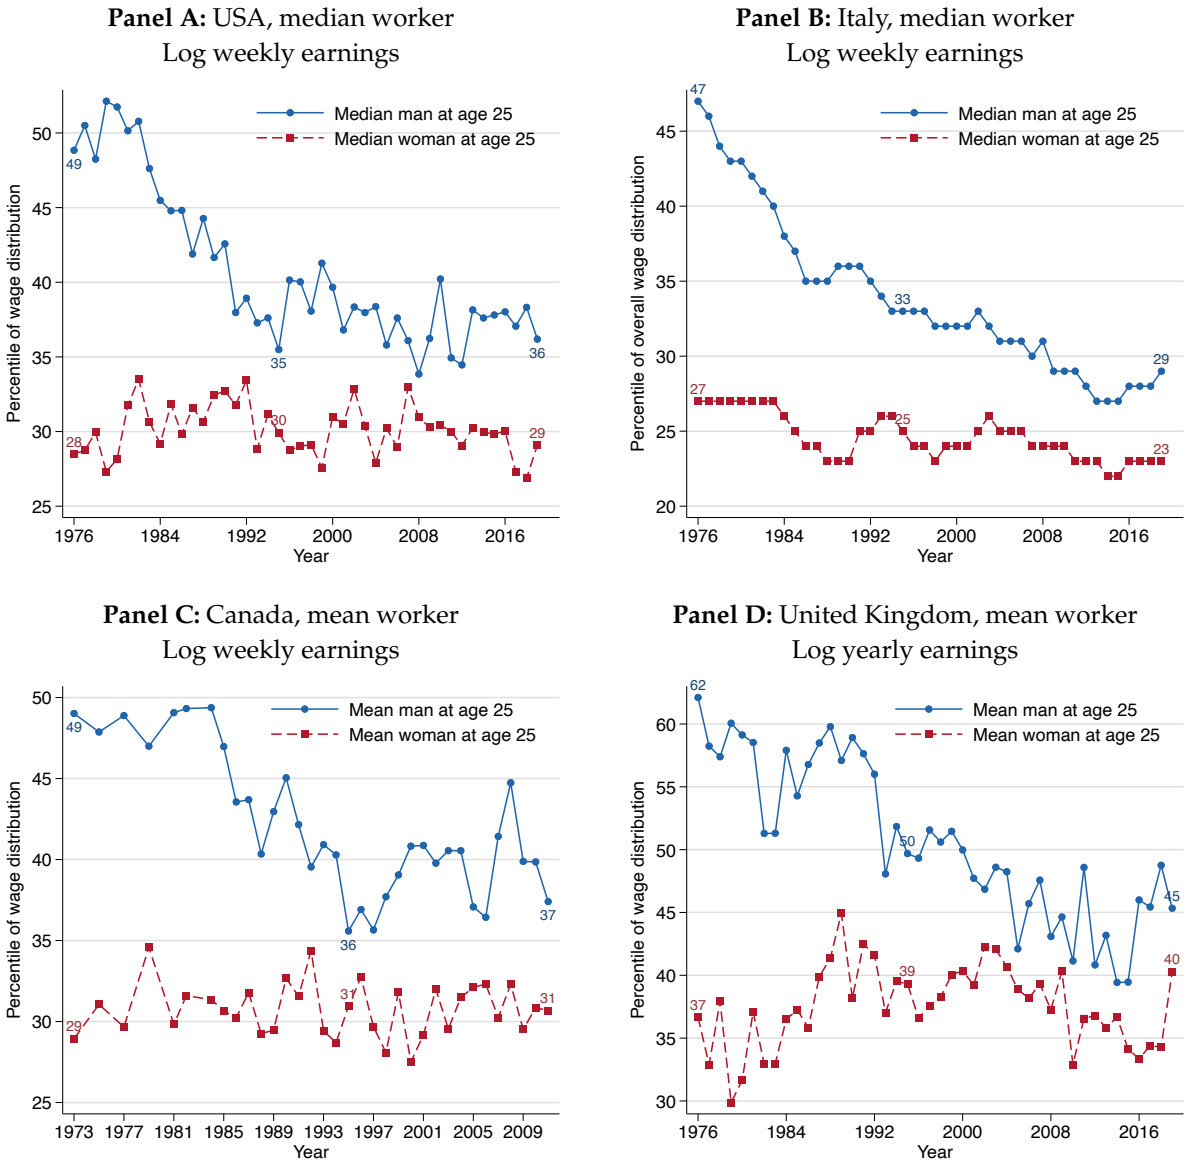
Notes: Panels A and B show the share of high-ranked managerial jobs out of all jobs in the economy and the share of high-ranked managerial jobs held by workers between 25 and 30 years old in the United States and Italy, respectively. We define as higher-ranked managerial jobs all managerial occupations with annual earnings in the top quartile of the year-specific distribution of annual earnings. In the CPS data, managerial occupations are identified using 2-digit Standard Occupational Classification (SOC) code 11. In the INPS data, we use the highest ranked position out of the four main job categories in the Italian labor system: in ascending order, these are apprenticeships, blue-collar jobs, white-collar jobs, and managerial jobs. The figures based on Italian data show a spurious trend discontinuity in the mid 1990s because the definition of managerial jobs in the INPS data changes from 1996.
Source for the United States. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A9: Managerial Positions With Above-Median Pay



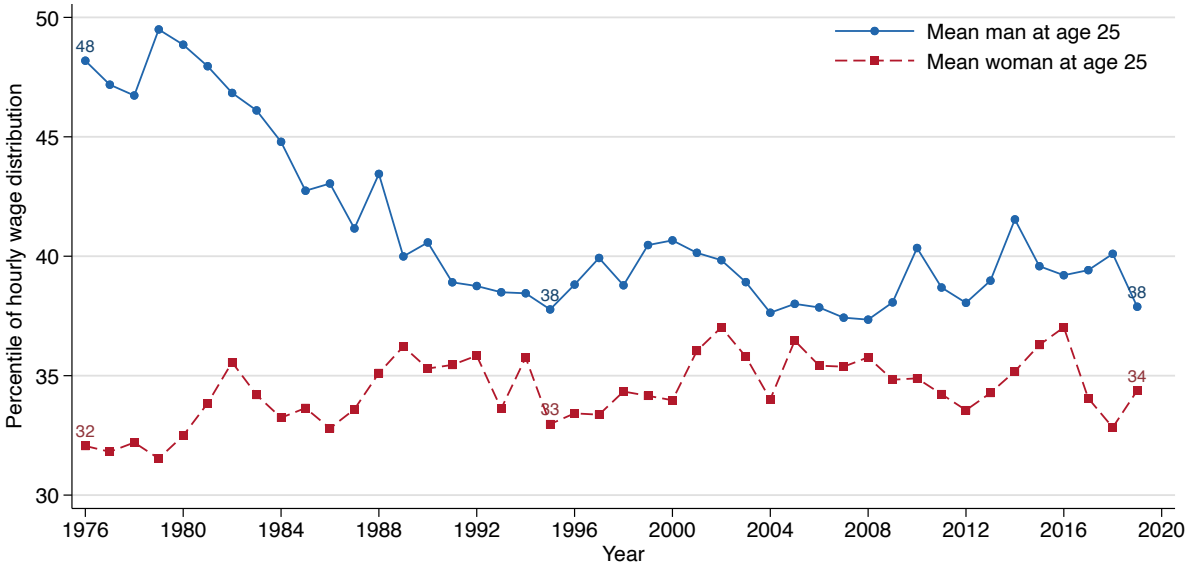
Notes: We define as higher-ranked managerial jobs all managerial occupations with annual earnings above the median of the year-specific distribution of annual earnings. In the CPS data, managerial occupations are identified using 2-digit Standard Occupational Classification (SOC) code 11. In the INPS data, we use the highest ranked position out of the four main job categories in the Italian labor system: in ascending order, these are apprenticeships, blue-collar jobs, white-collar jobs, and managerial jobs. The figures based on Italian data show a spurious trend discontinuity in the mid 1990s because the definition of managerial jobs in the INPS data changes from 1996. Panels A and C show the share of high-ranked managerial jobs out of all jobs in the economy and the share of high-ranked managerial jobs held by workers between 25 and 30 years old in the United States and Italy, respectively. Panels B and D show the share of high-ranked managerial jobs held by men and women between 25 and 30 years old in the United States and Italy, respectively.
Source for the United States. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A10: Rank at Labor-Market Entry; Median and More Countries



Notes: Panels A and B show the median earning percentile of men and women at age 25 in the United States and Italy, respectively. Panels C and D show the average earning percentile of men and women at age 25 in Canada and the United Kingdom, respectively. In each year, the data encompass information about all workers who were 25 years old, had worked in the private sector for at least 24 weeks (not available in the UK), had earned strictly positive earnings, and (only in Italy) had not retired by December 31.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS. *Source for LIS data:* Luxembourg Income Study (LIS) Database, last accessed on 06/01/2023 at https://www.lisdatacenter.org/.

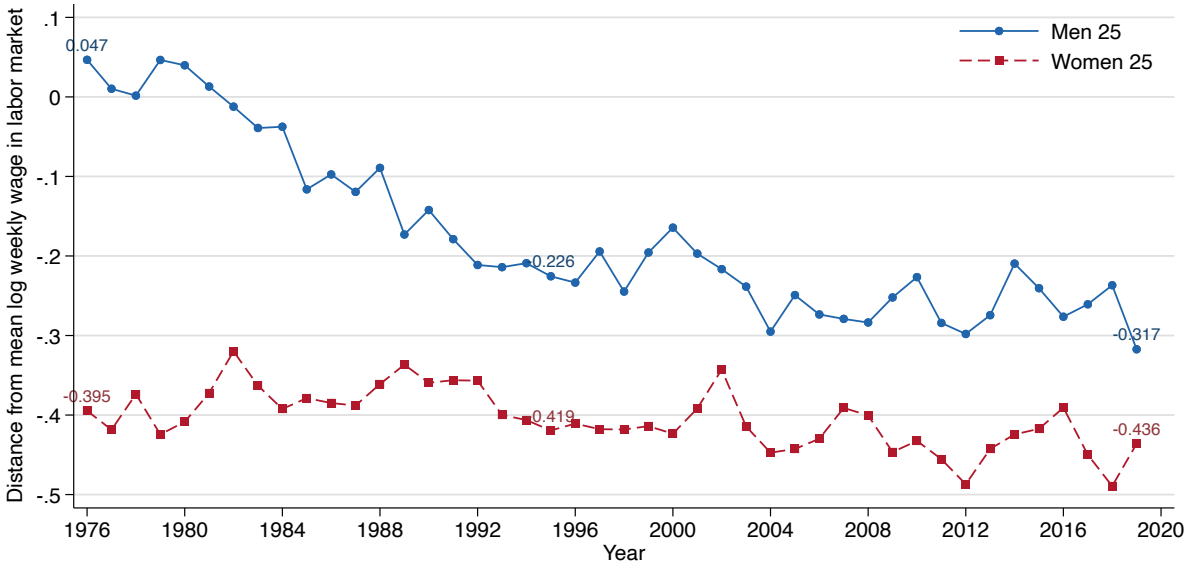
Figure A11: Rank in Distribution of Hourly Earnings (US)



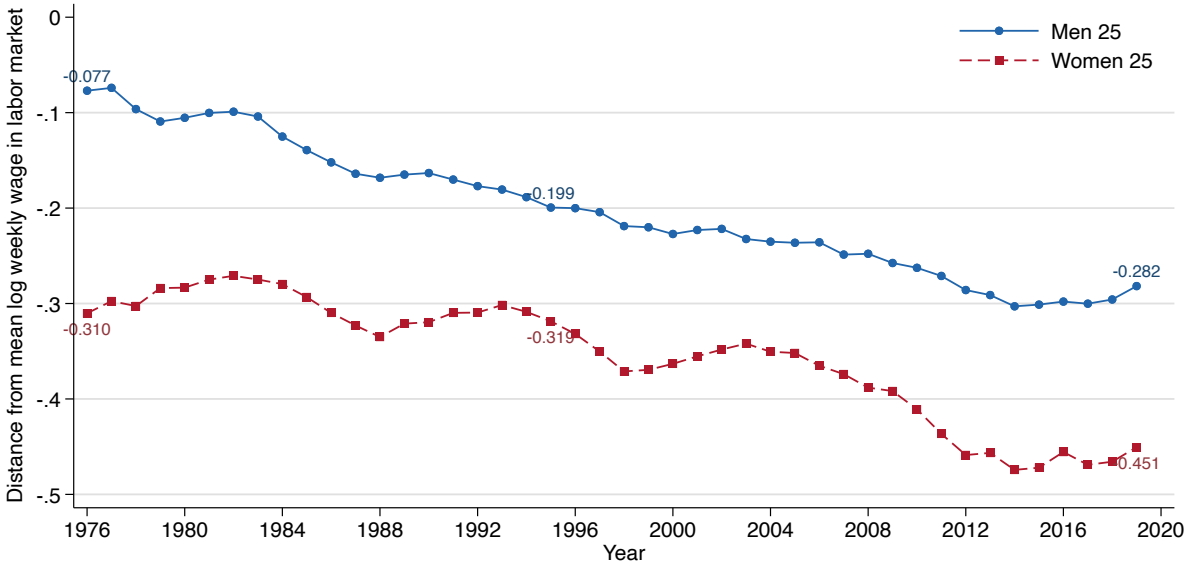
Notes: The figure shows the average percentile of men and women at age 25 in the distribution of hourly earnings in the United States. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings.
Source for the United States. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A12: Younger Workers' Distance from Mean Earnings

Panel A: USA



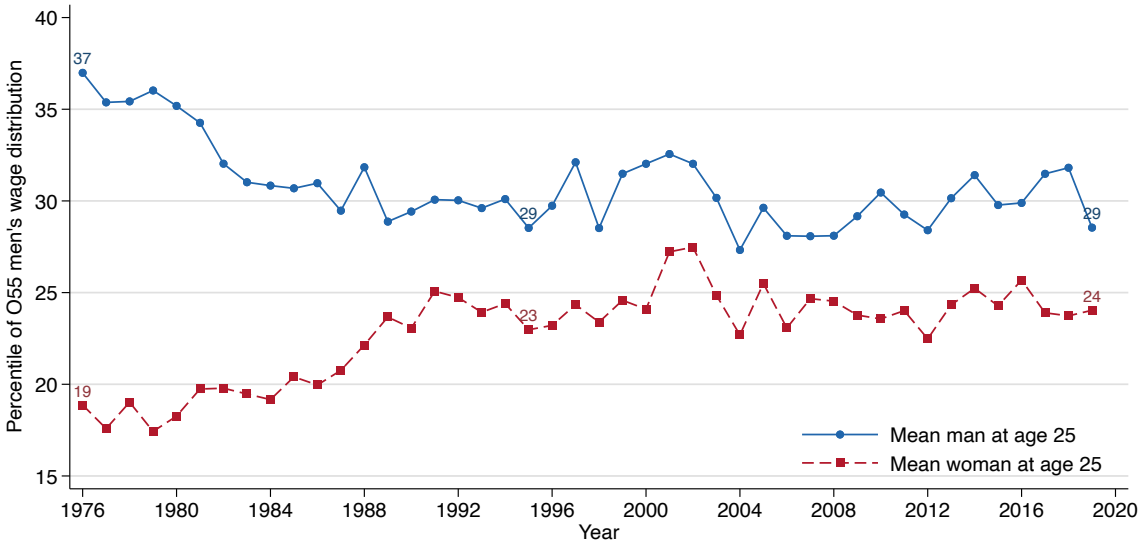
Panel B: Italy



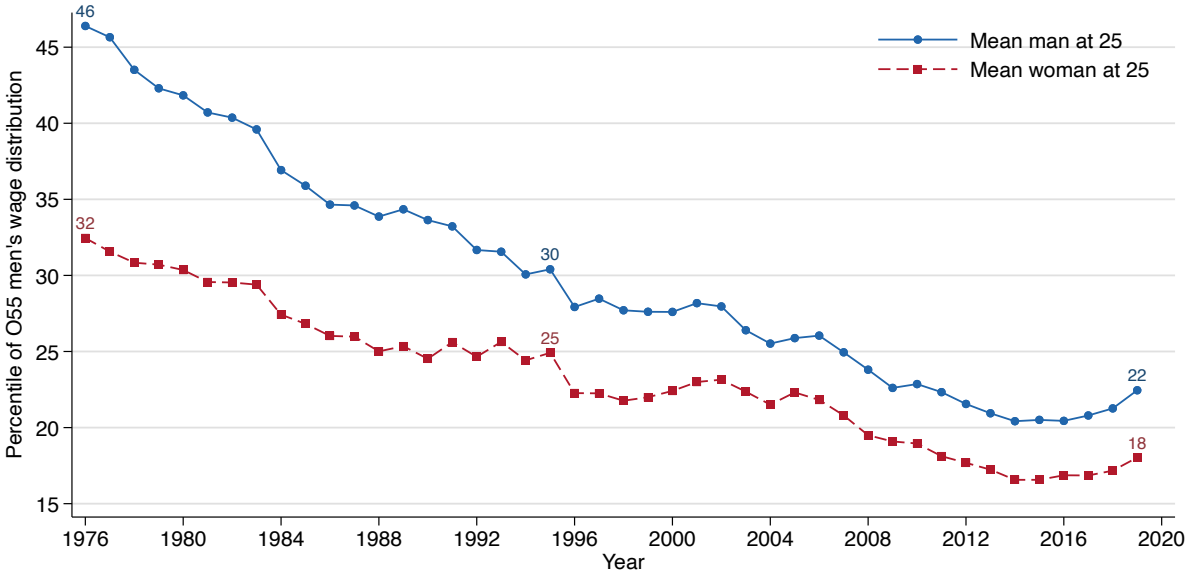
Notes: Panels A and B show the difference between the mean log weekly earnings of men and women at age 25 and the mean log weekly earnings computed using all private-sector workers in the labor market in the United States and Italy, respectively. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings, and (only in Italy) had not retired by December 31.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A13: Rank in Pay Distribution of Over-55 Men

Panel A: USA

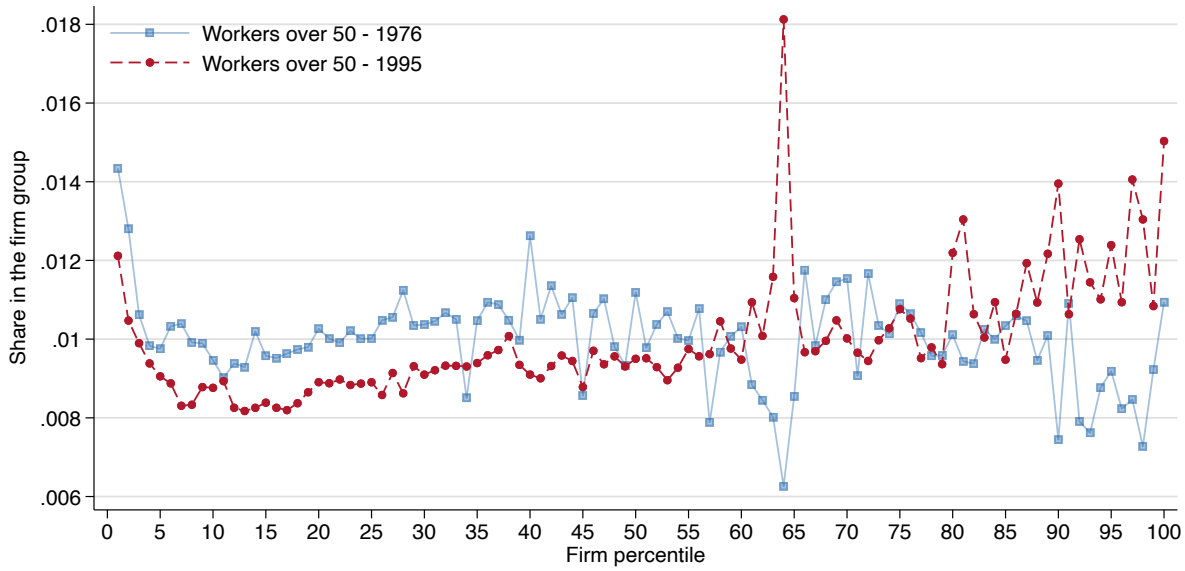


Panel B: Italy



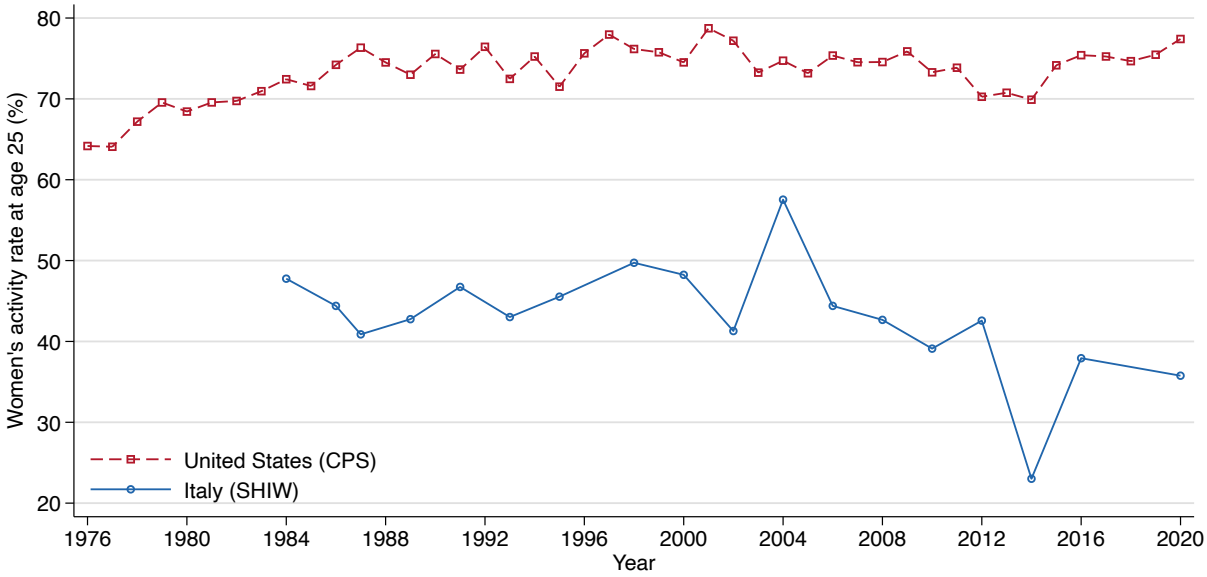
Notes: Panels A and B show the average earning percentile of men and women at 25 years old in the distribution of weekly earnings of over-55 male workers in the United States and Italy, respectively.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A14: Share of Workers Over 50 in Lower-Paying and Higher-Paying Firms



Notes: Figure shows the distribution of workers over 50 across percentiles of firms' mean pay in 1976 and 1995 in Italy. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings, and had not retired by December 31. *Source:* UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

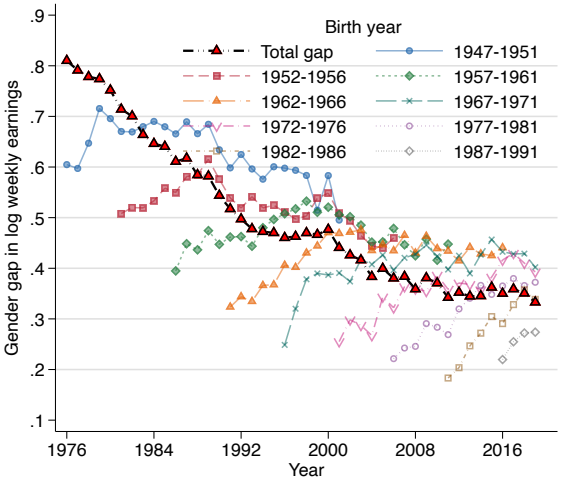
Figure A15: Participation Rates at Age 25



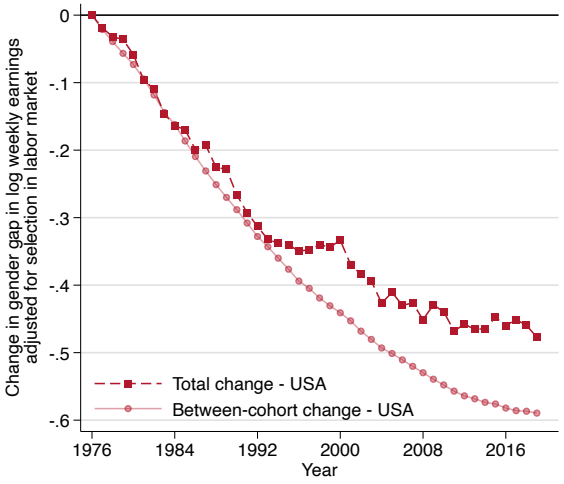
Notes: The figure shows participation rates of women at age 25 in the United States and Italy. *Source for Italy*: Survey of Household Income and Wealth, Bank of Italy. *Source for the United States*: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A16: Results Adjusted for Selection into Labor Market, Weekly Earnings

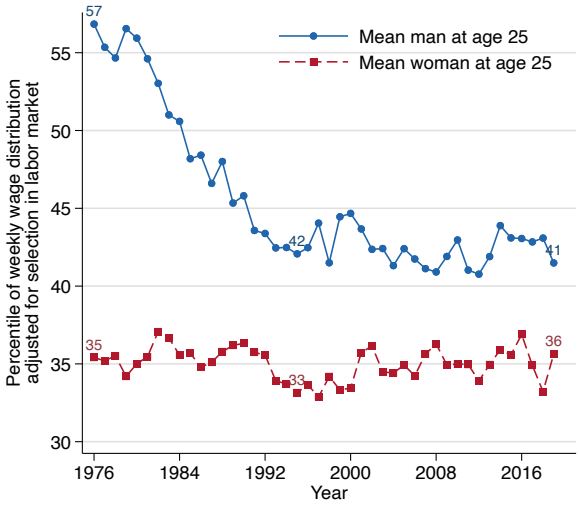
Panel A: Raw gaps



Panel B: Between-cohort change



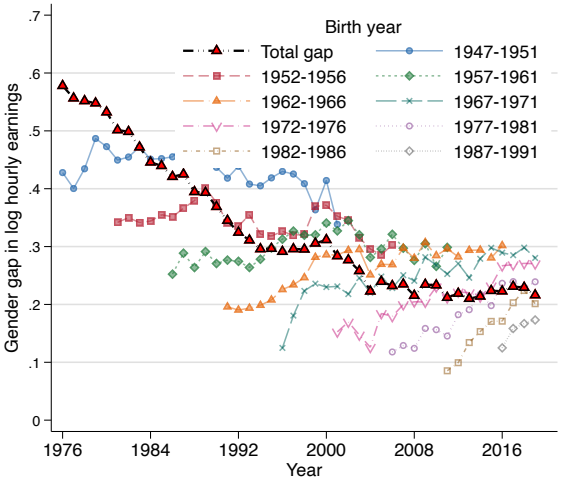
Panel C: Rank in pay distribution



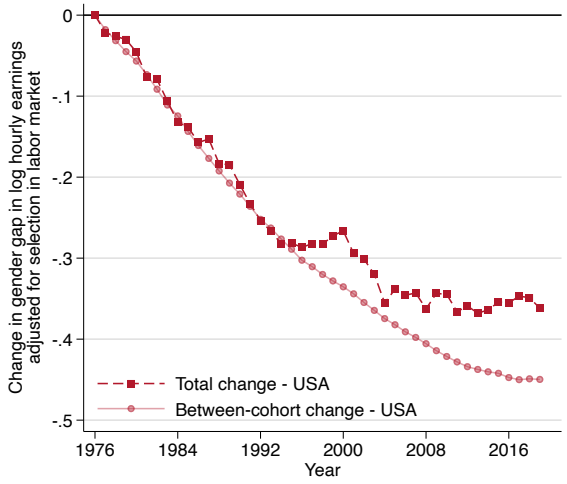
Notes: These figures replicate the main findings of the paper using a larger sample in which nonparticipants in the labor market are included with imputed weekly earnings. To account for changes in selection into the labor market, we follow the procedure outlined in Blau et al. (2024). First, we expand the sample to include individuals who worked less than 24 weeks in the year (but at least 100 hours) and had positive labor earnings. Second, we impute the weekly earnings of nonparticipants based on their probability of falling into each decile of the pay distribution. We predict these probabilities via an ordered probit estimated separately by year and gender using the following observable characteristics: years of education, a dummy for college graduation, a dummy for advanced degrees, potential experience (age-years of education-6), potential experience squared, race (white, black, others) and ethnicity (hispanic) dummies, and fixed effects for Census divisions. Each nonparticipant is then included in the sample ten times, one per decile, and assigned the decile-specific median weekly earnings, while the predicted probabilities of falling in each decile are used to adjust the sampling weights. This procedure is repeated until the earning distributions at the start and end of the imputation process become sufficiently similar.
Source for the United States. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A17: Results Adjusted for Selection into Labor Market, Hourly Earnings

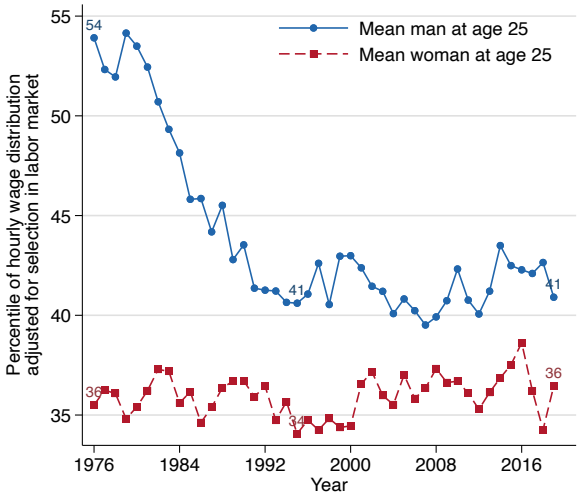
Panel A: Raw gaps



Panel B: Between-cohort change



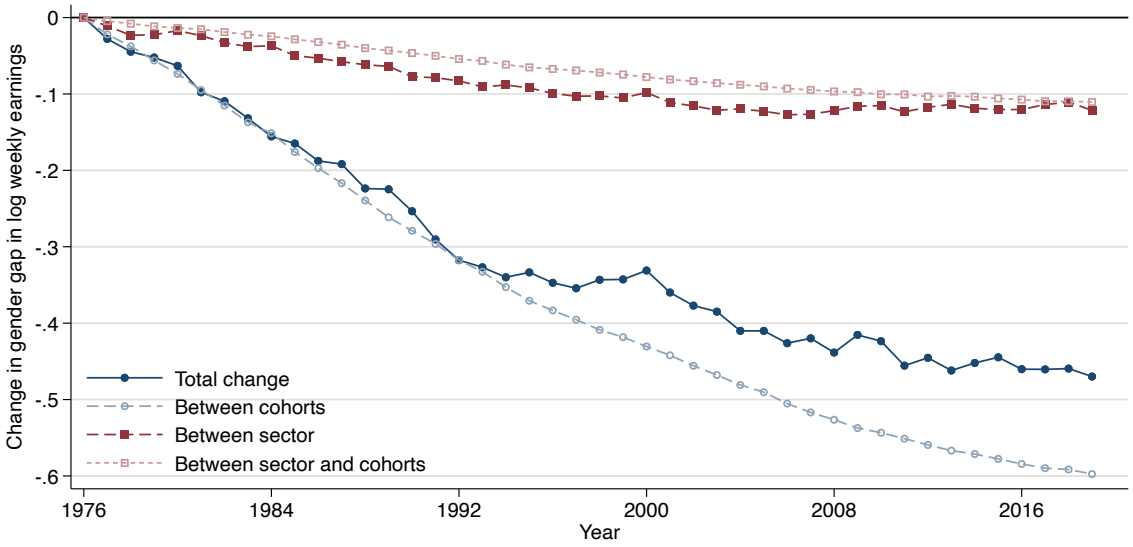
Panel C: Rank in pay distribution



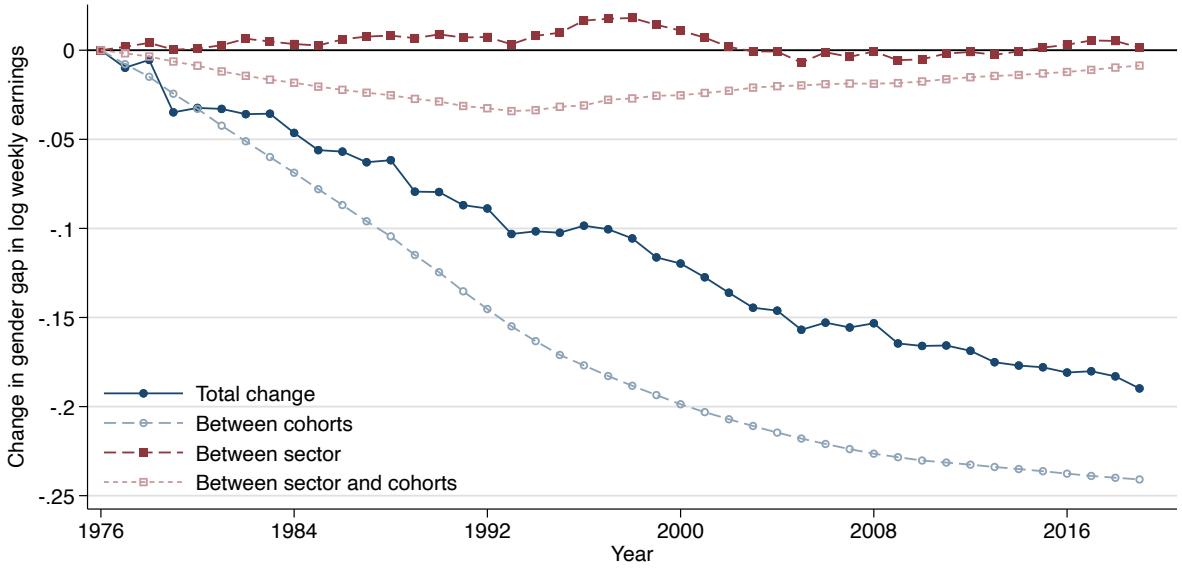
Notes: These figures replicate the main findings of the paper using a larger sample in which nonparticipants in the labor market are included with imputed hourly earnings. To account for changes in selection into the labor market, we follow the procedure outlined in Blau et al. (2024). First, we expand the sample to include individuals who worked less than 24 weeks in the year (but at least 100 hours) and had positive labor earnings. Second, we impute the hourly earnings of nonparticipants based on their probability of falling into each decile of the pay distribution. We predict these probabilities via an ordered probit estimated separately by year and gender using the following observable characteristics: years of education, a dummy for college graduation, a dummy for advanced degrees, potential experience (age-years of education-6), potential experience squared, race (white, black, others) and ethnicity (hispanic) dummies, and fixed effects for Census divisions. Each nonparticipant is then included in the sample ten times, one per decile, and assigned the decile-specific median hourly earnings, while the predicted probabilities of falling in each decile are used to adjust the sampling weights. This procedure is repeated until the earning distributions at the start and end of the imputation process become sufficiently similar.
Source for the United States. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A18: Between-Cohort Decline in Gender Pay Gap, Between and Within Sectors

Panel A: USA

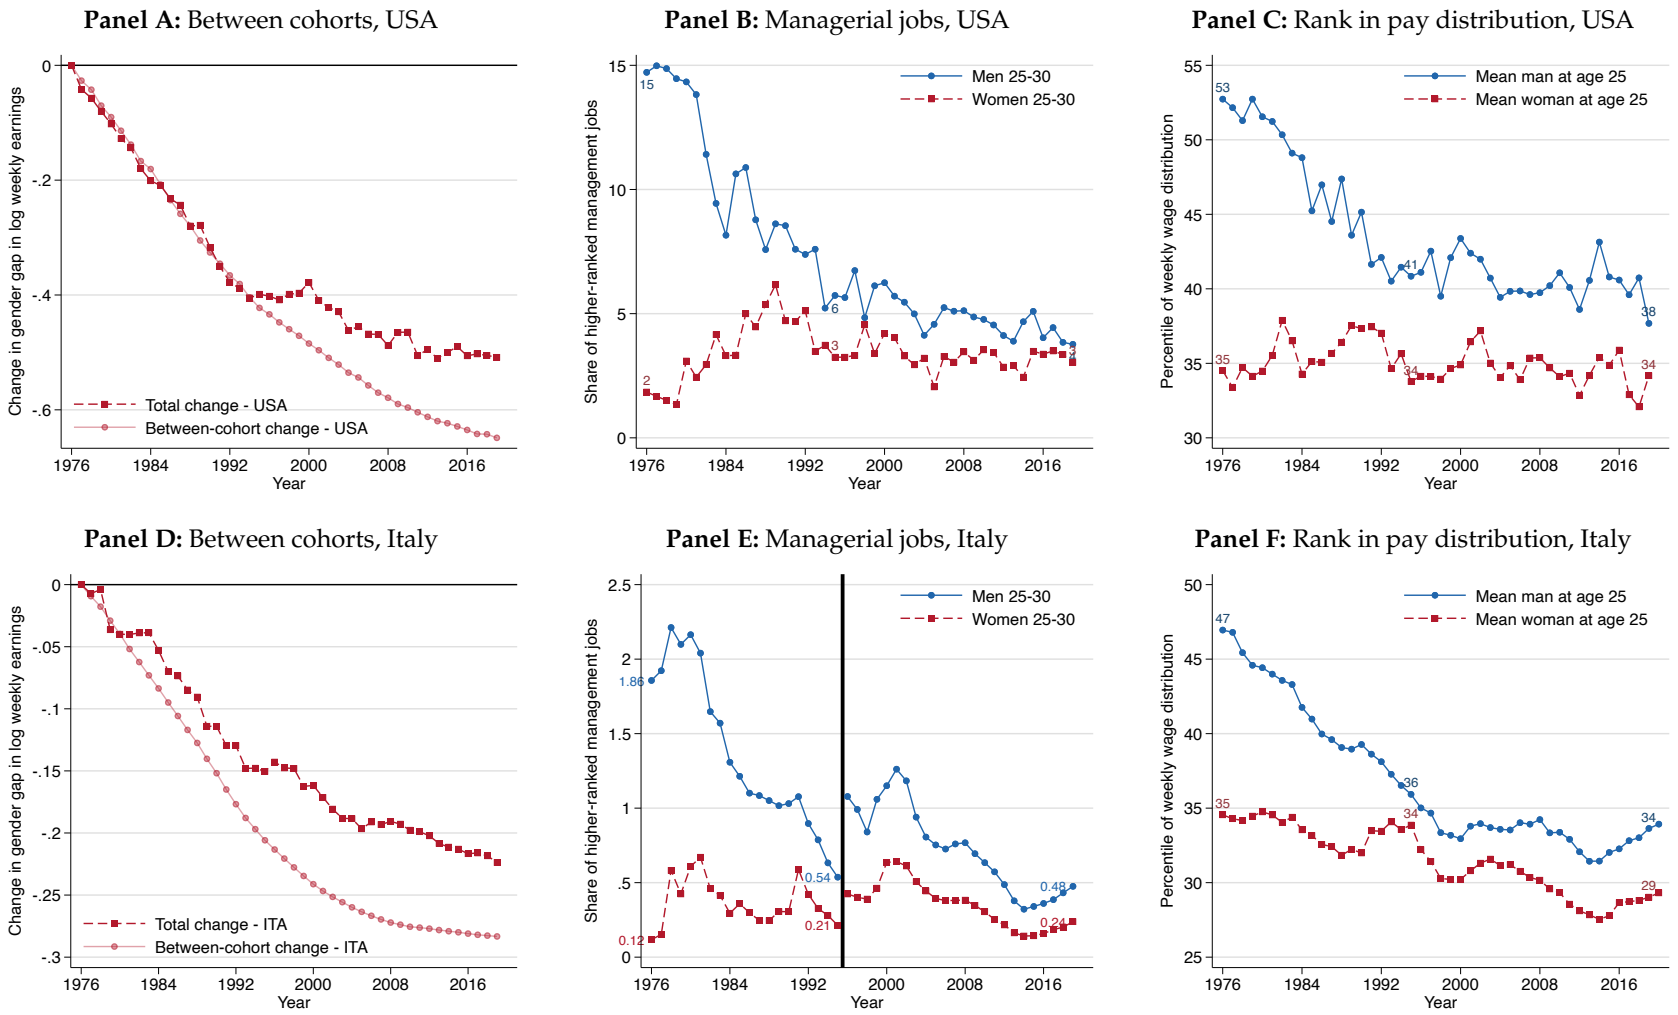


Panel B: Italy



Notes: Panels A and B decompose both the total change in gender pay gap and just its between-cohort component between and within sector (1- and 2-digit NACE Rev. 2 in the United States and Italy, respectively).
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

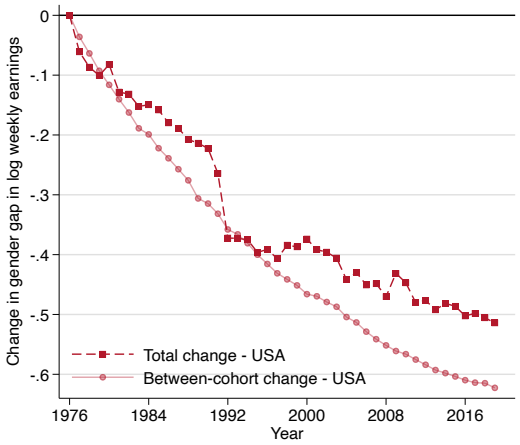
Figure A19: Dropping Workers in Manufacturing



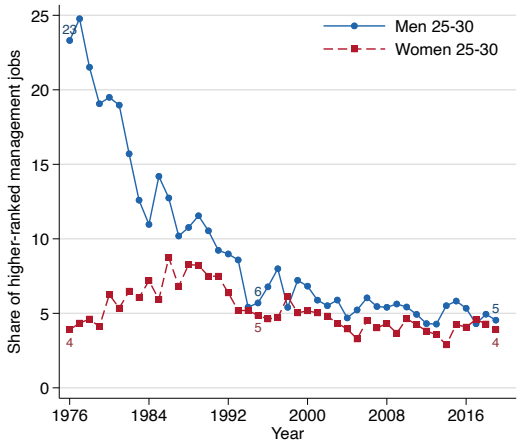
Notes: These figures replicate the main findings after dropping from the sample all workers in manufacturing. Panels A and D decompose the decline in the gender pay gap into a between-cohort and within-cohort component. Panels B and E show the probability of younger men and women holding higher-paying (top quartile of pay distribution) managerial jobs. Panels C and F show the average percentile of younger men and women in the pay distribution. *Source for Italy*: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A20: Workers With High Probability of Working Outside Manufacturing

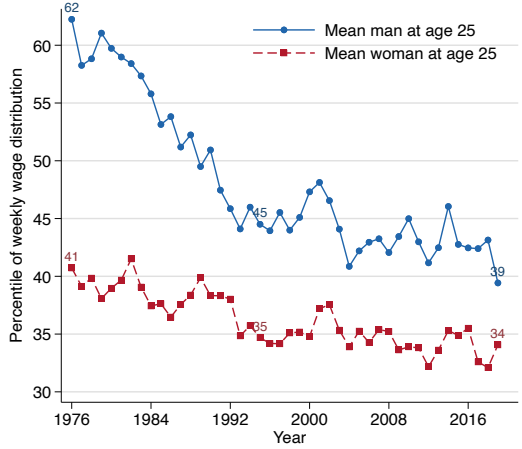
Panel A: Between cohorts, USA



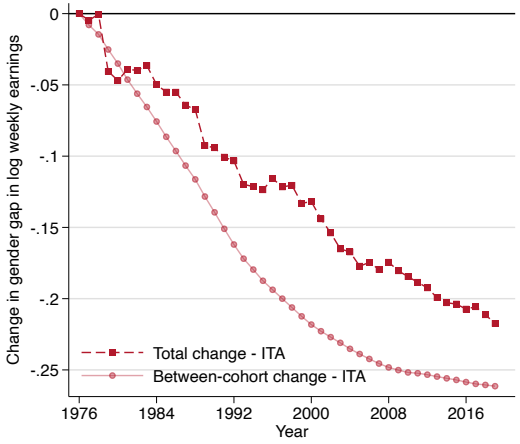
Panel B: Managerial jobs, USA



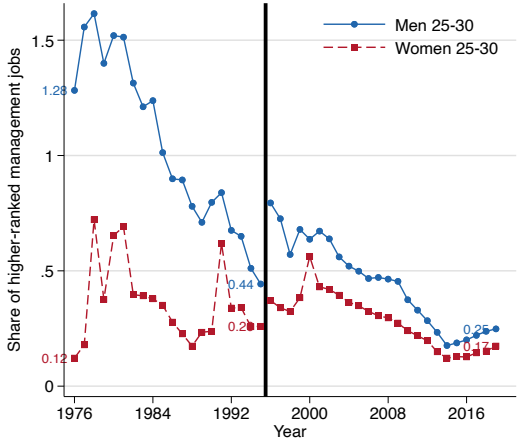
Panel C: Rank in pay distribution, USA



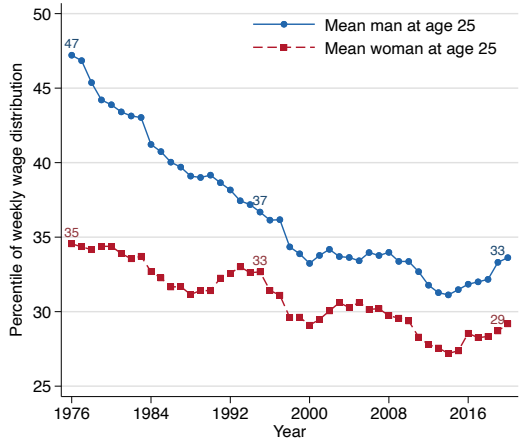
Panel D: Between cohorts, Italy



Panel E: Managerial jobs, Italy

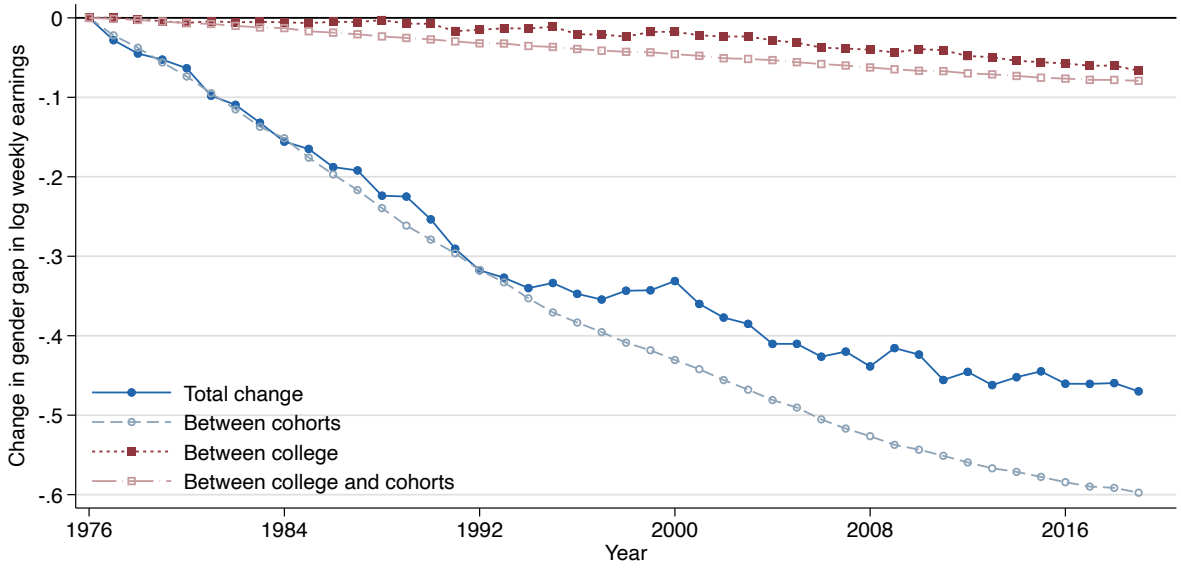


Panel F: Rank in pay distribution, Italy



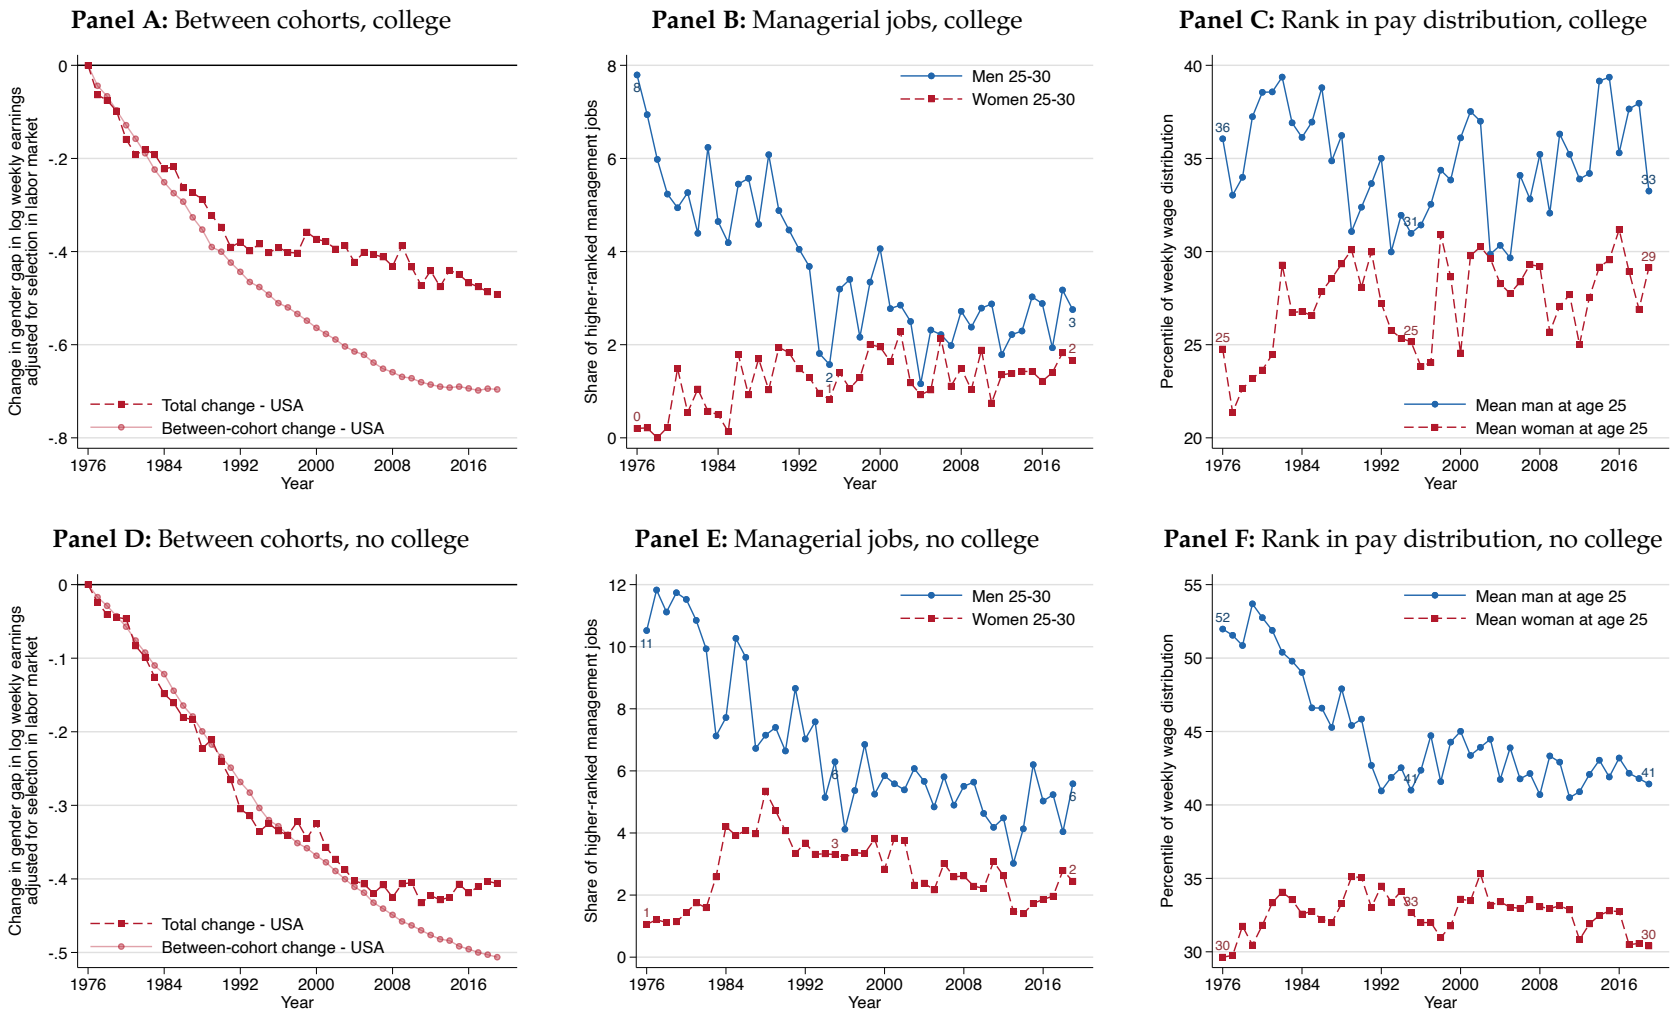
Notes: These figures replicate the main findings after keeping in the sample only workers with high predicted probability of working outside manufacturing. We predict the probability of working outside of manufacturing using the first five sample years (1976-1981) and the following observable characteristics: Census divisions (US) or province (IT) FEs interacted with age, age squared, completed education (less than high school, high school, some college, four years of college or more; US), gender, race (the same three categories used to address selection into the labor market; US), ethnicity (US), and a foreign-born dummy (IT). Next, we use the coefficients to predict the probability of working outside of manufacturing for the whole sample. Then, we retain in the sample only individuals with predicted probability$>$ 0.7. *Source for Italy*: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A21: Between-Cohort Decline in Gender Pay Gap, Between and Within College Education



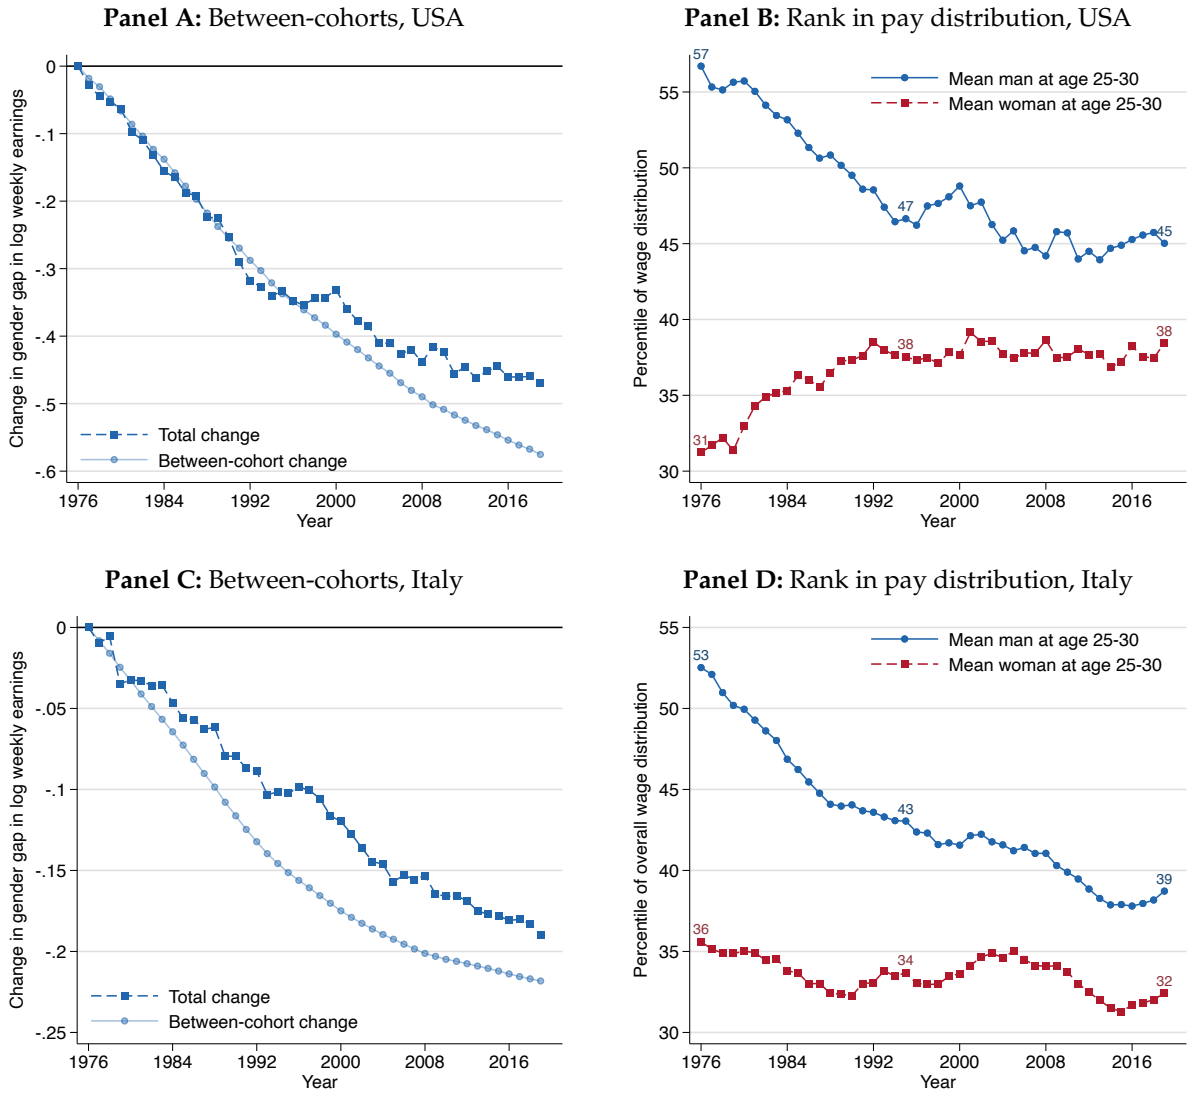
Notes: The figure decomposes both the total change in gender pay gap and just its between-cohort component between and within college graduation in the United States.
Source for the United States: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A22: Results for Workers With and Without College Education



Notes: These figures replicate the main findings separately for workers with (Panels A-C) and without (Panels D-F) a college education in the United States.
Source for the United States: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

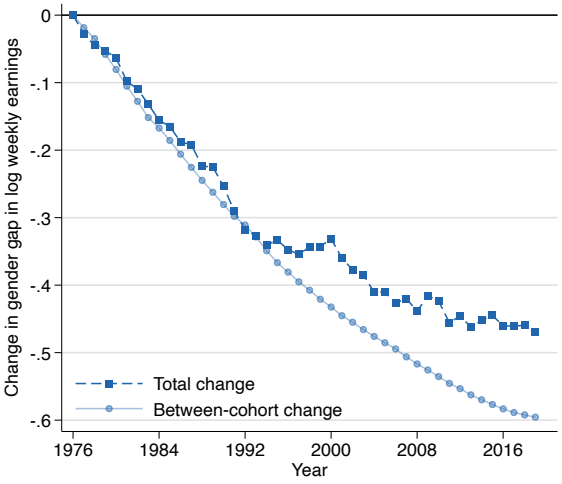
Figure A23: Different Definition of Younger Workers (25-30)



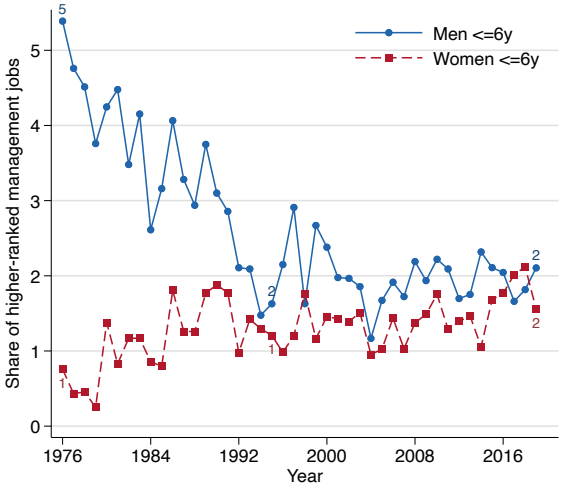
Notes: These figures replicate the main findings (decomposition between and within worker cohorts) and change in mean pay rank over time after changing the definition of younger workers. Instead of analyzing the outcomes of workers at age 25, these figures show outcomes of workers between 25 and 30 years old. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A24: Definition of Younger Workers Based on Potential Experience

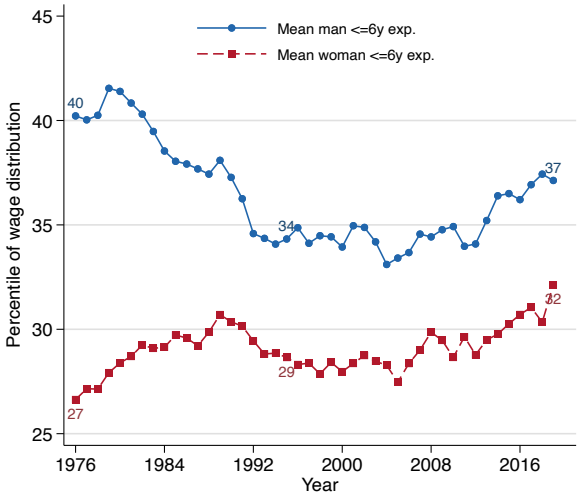
Panel A: Between cohorts



Panel B: Managerial jobs

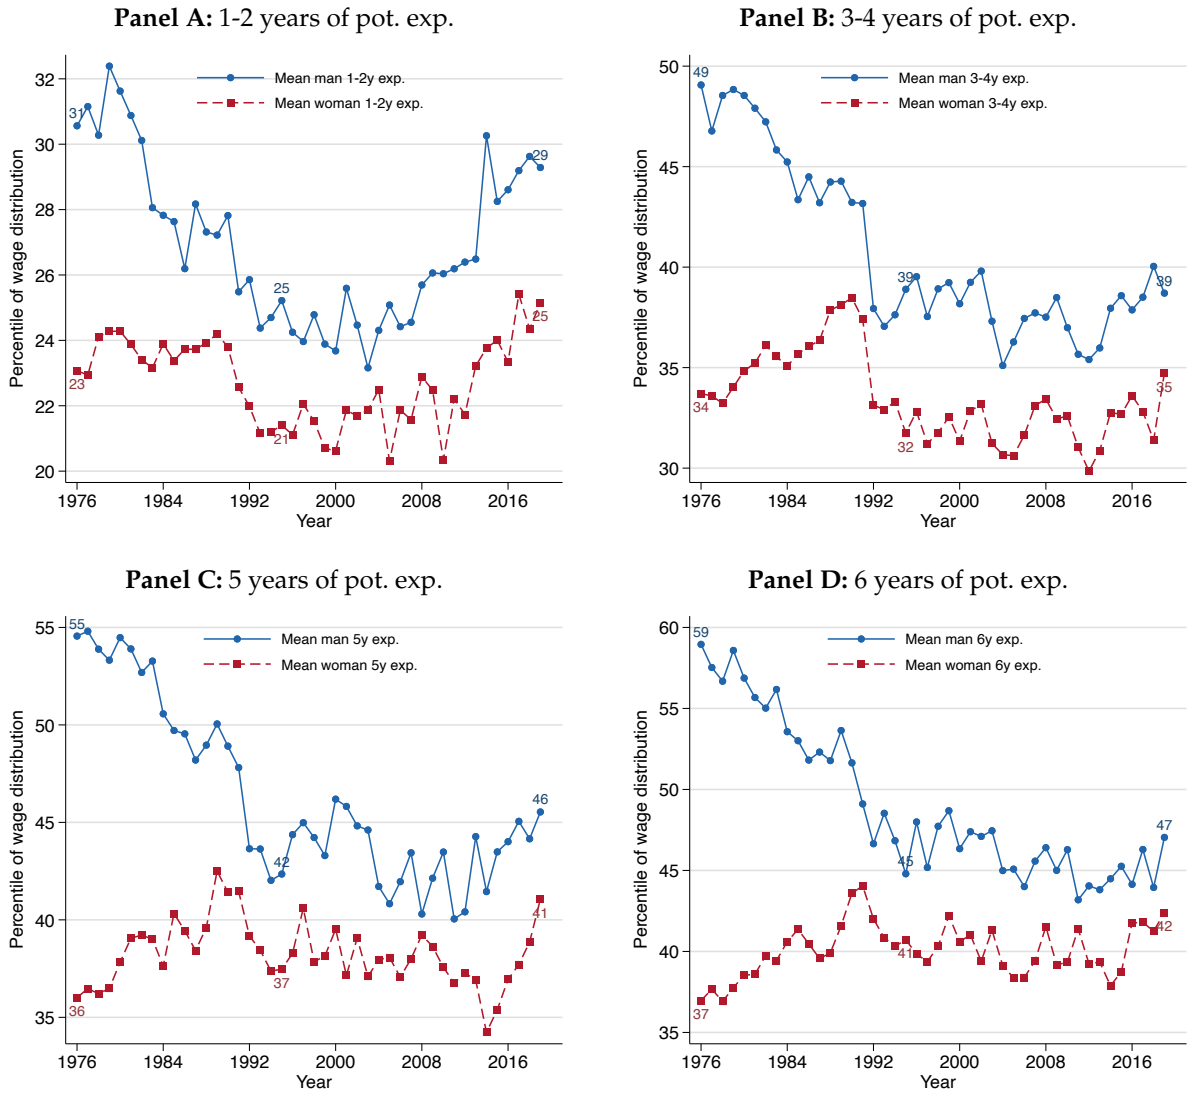


Panel C: Rank in pay distribution



Notes: These figures replicate the main findings after changing the definition of younger workers. In these analysis, younger men and women have at most six years of potential experience in the labor market. Potential experience is calculated as age − years of education − 6. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings. *Source for the United States*: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

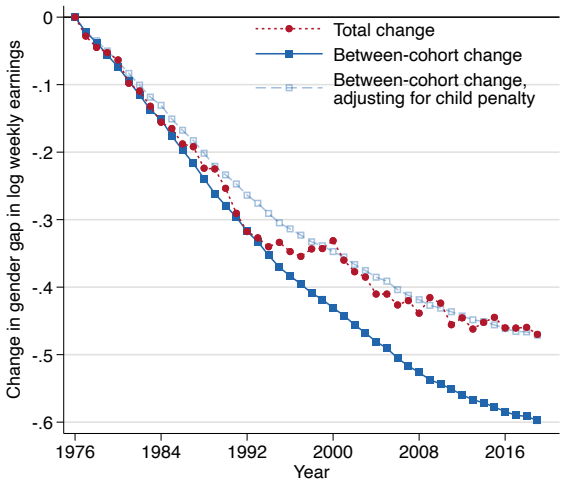
Figure A25: Rank in Pay Distribution Based on Potential Experience



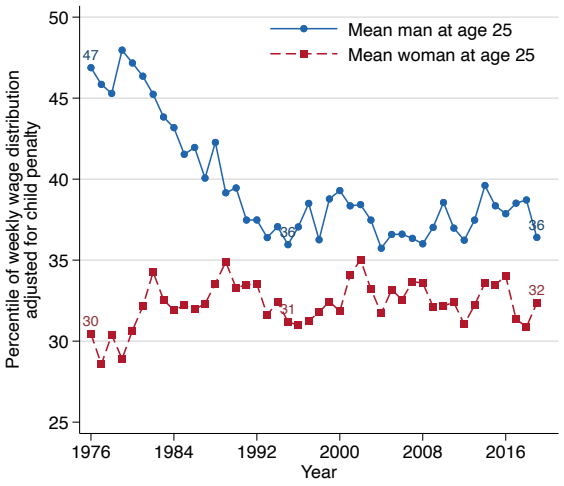
Notes: Panels A to D show the average earning percentile of men and women at different years of potential experience in the United States. Potential experience is calculated as age − years of education − 6. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings.
Source for the United States. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A26: Controlling for Child Penalty

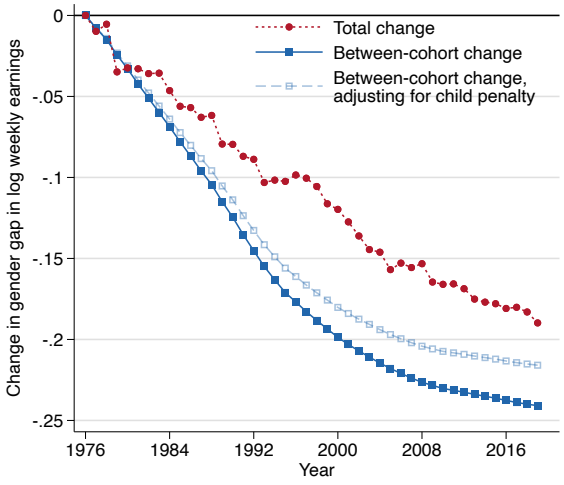
Panel A: USA, between-cohort change



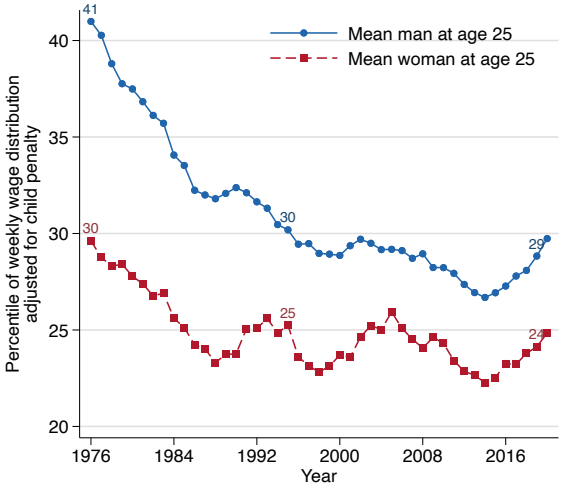
Panel B: USA, rank in pay distribution



Panel C: Italy, between-cohort change



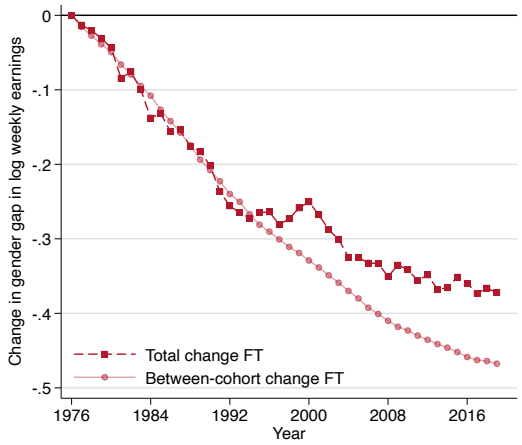
Panel D: Italy, rank in pay distribution



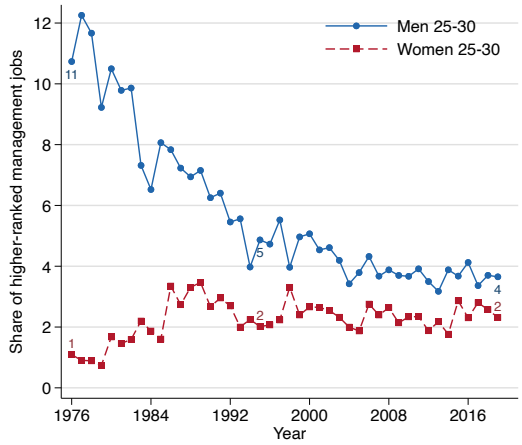
Notes: These figures use weekly earnings that include the child penalty borne by women in the labor market (Kleven, Landais, and Søgaard, 2019). More details are in Appendix D.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A27: Full-Time Workers

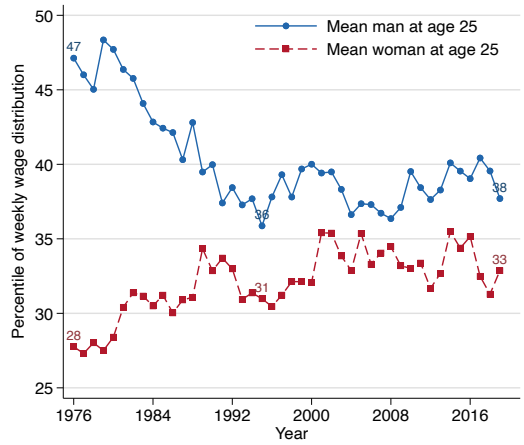
Panel A: Between cohorts, USA



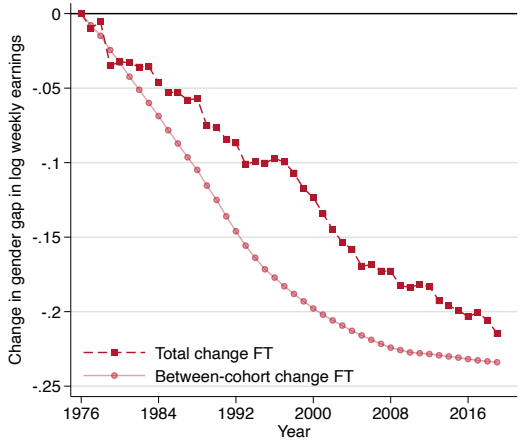
Panel B: Managerial jobs, USA



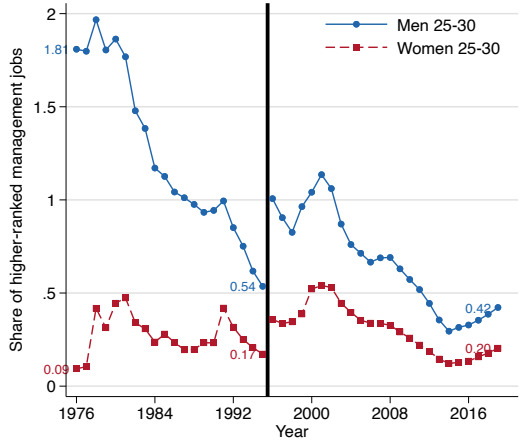
Panel C: Rank in pay distribution, USA



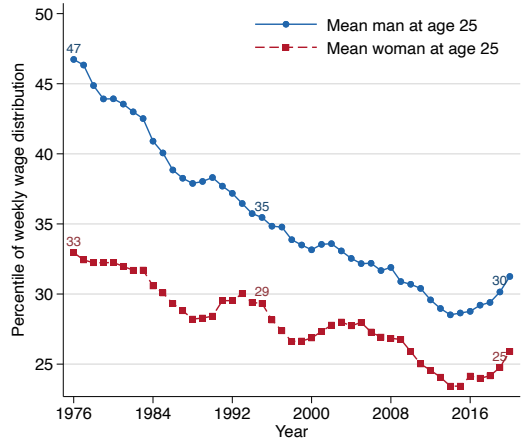
Panel D: Between cohorts, Italy



Panel E: Managerial jobs, Italy

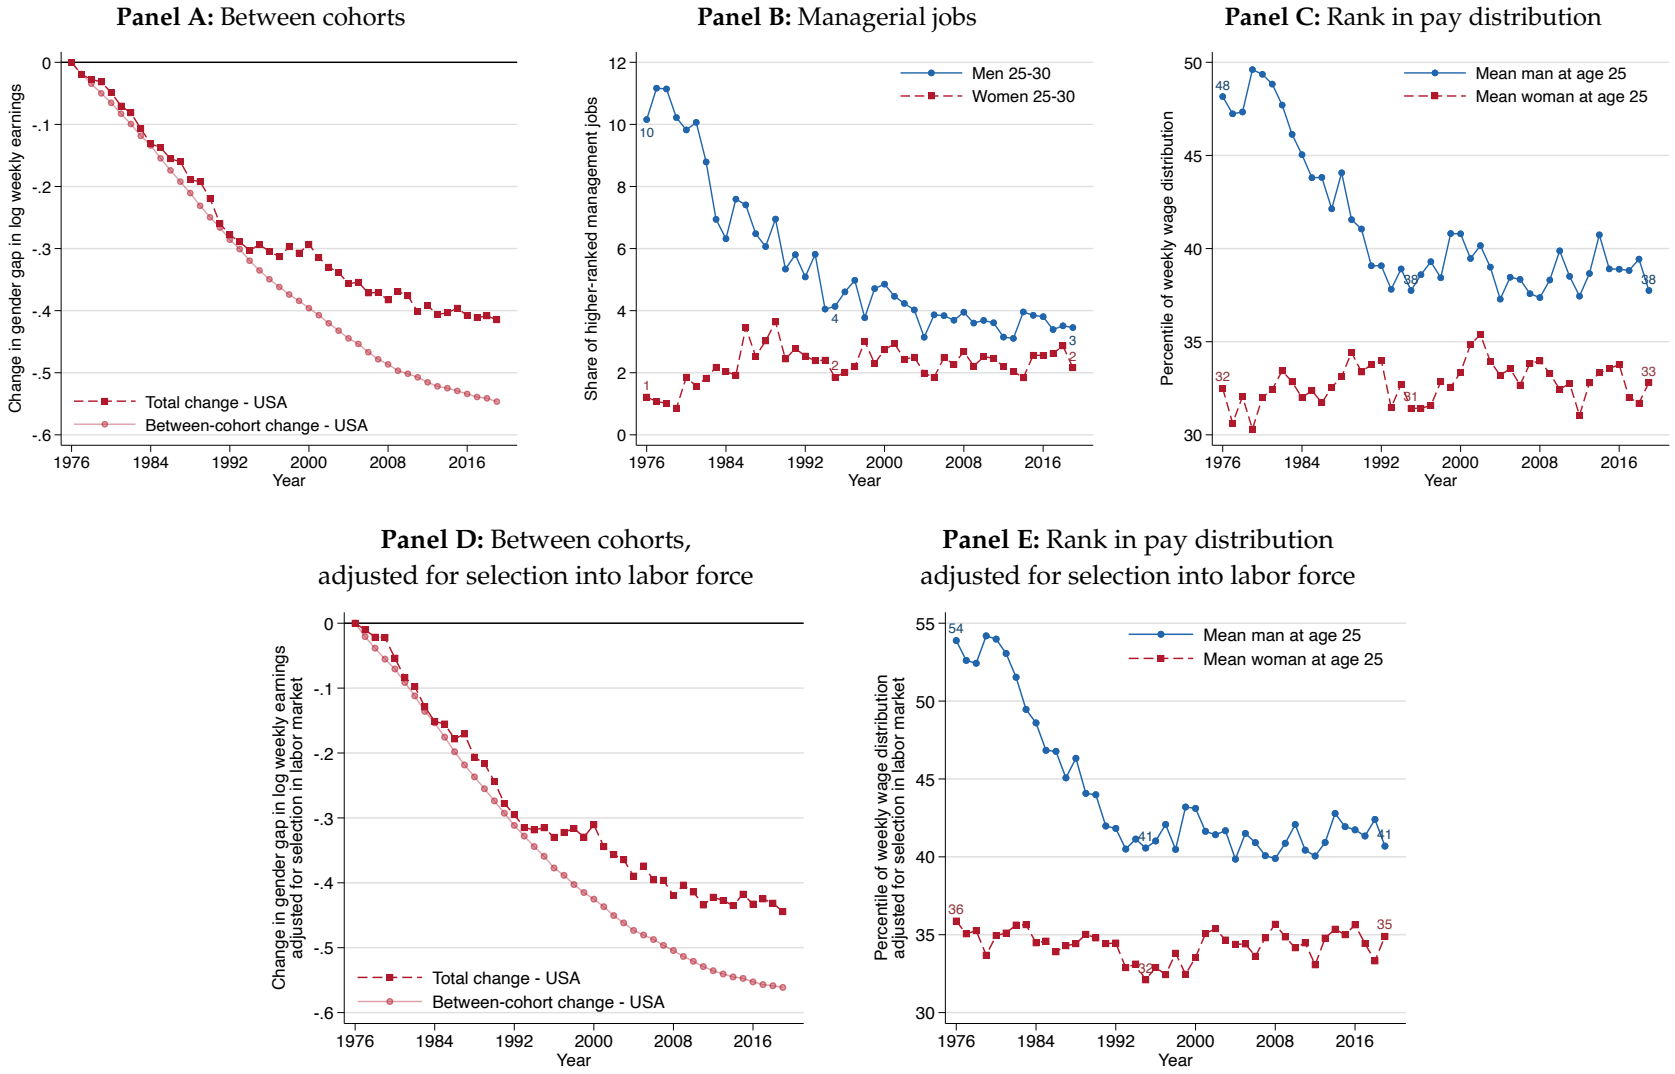


Panel F: Rank in pay distribution, Italy



Notes: These figures replicate the main findings after keeping in the sample only workers with full-time contracts.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

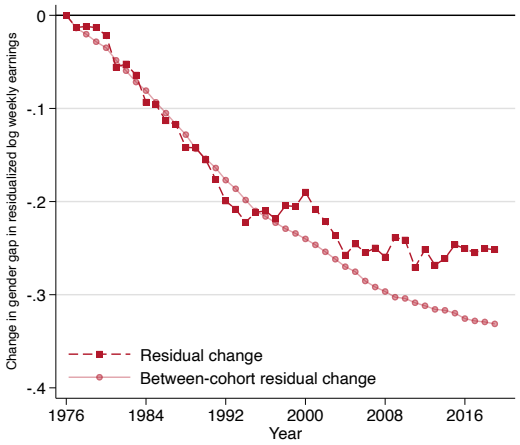
Figure A28: Private-Sector and Public-Sector Employees



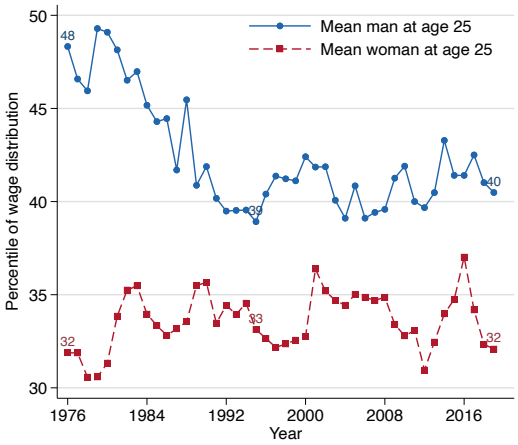
Notes: These figures replicate the main findings including in the sample also public-sector employees in the United States. Panels A-C expand the sample, including public-sector employees. Panels D and E expand the sample with public-sector employees and impute weekly earnings to nonparticipants using the process outlined in Blau et al. (2024). *Source for the United States*: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A29: Residual Earnings

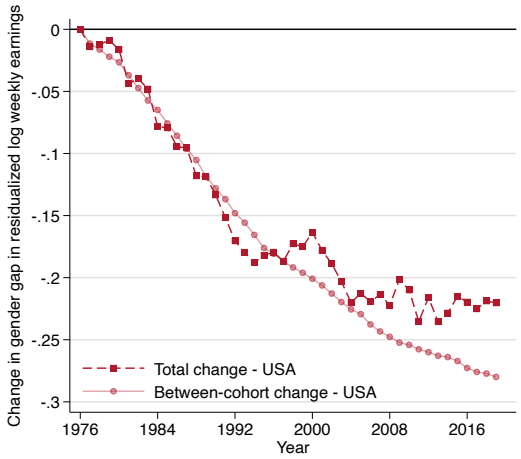
Panel A: Between cohorts, USA



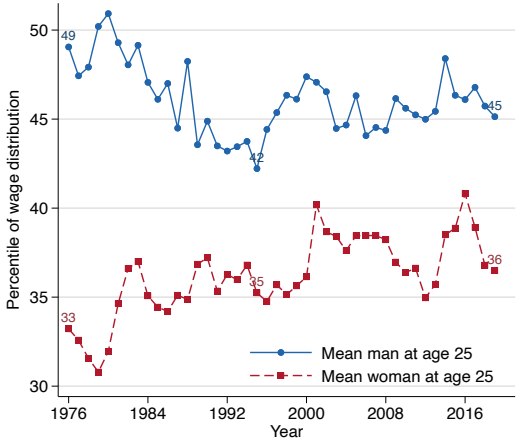
Panel B: Rank in pay distribution, USA



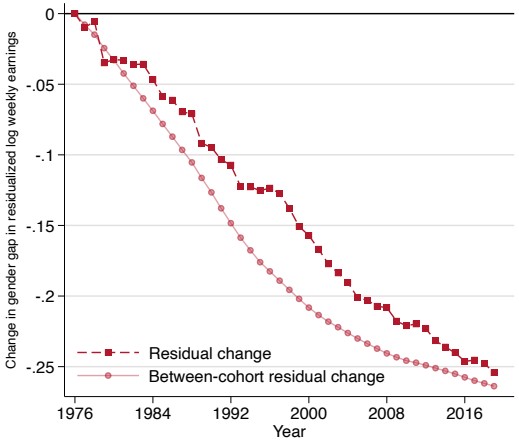
Panel C: Between cohorts, USA



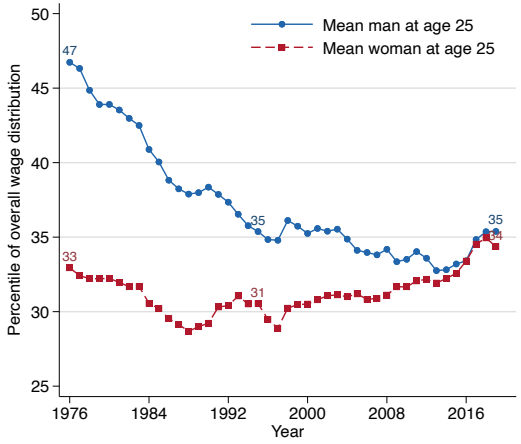
Panel D: Rank in pay distribution, USA



Panel E: Between cohorts, Italy



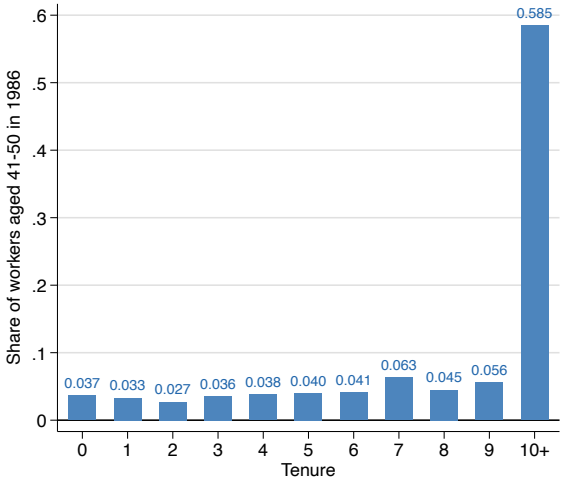
Panel F: Rank in pay distribution, Italy



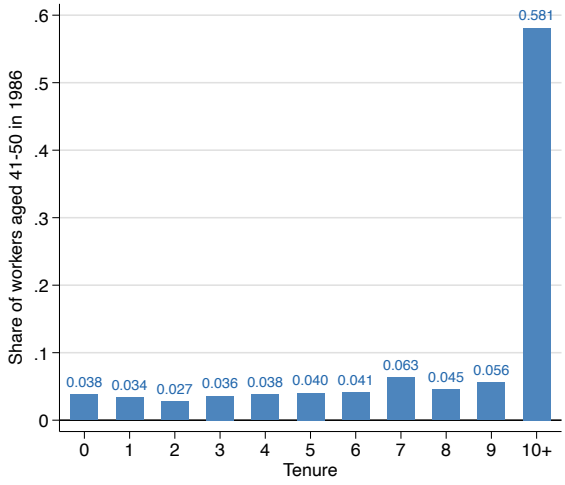
Notes: We regress the log of weekly earnings on a dummy for part-time workers, a dummy for temporary workers (only in Italy), a dummy for domestic-born workers (only in Italy), dummies for race (only in the US), a dummy for college graduation (only in the US), and a dummy for workers with children (only in the US). For the US, we have a second version in which in addition to the previous regressors, we also include a dummy for hispanic ethnicity, dummies for marital status, and fixed effects for Census divisions (Panels C and D). We estimate these regressions separately by year and country. *Source for Italy*: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A30: Tenure of Workers Aged 41-50 and 31-40 in 1986

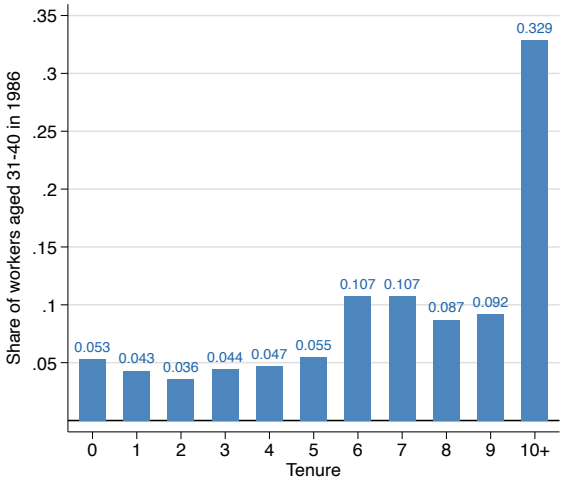
Panel A: Workers 41-50 in larger firms that existed in 1976



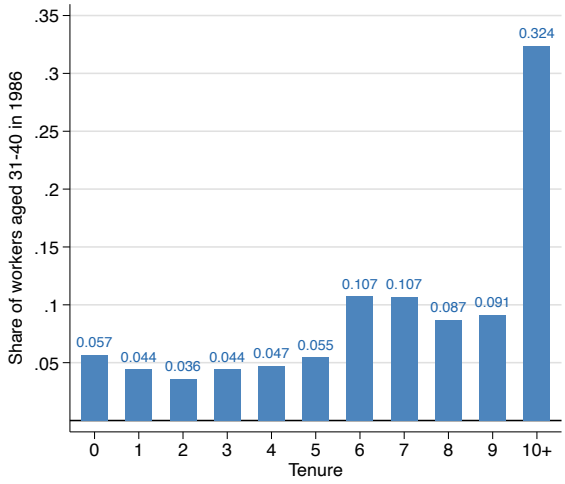
Panel B: Workers 41-50 in firms that existed in 1976



Panel C: Workers 31-40 in larger firms that existed in 1976

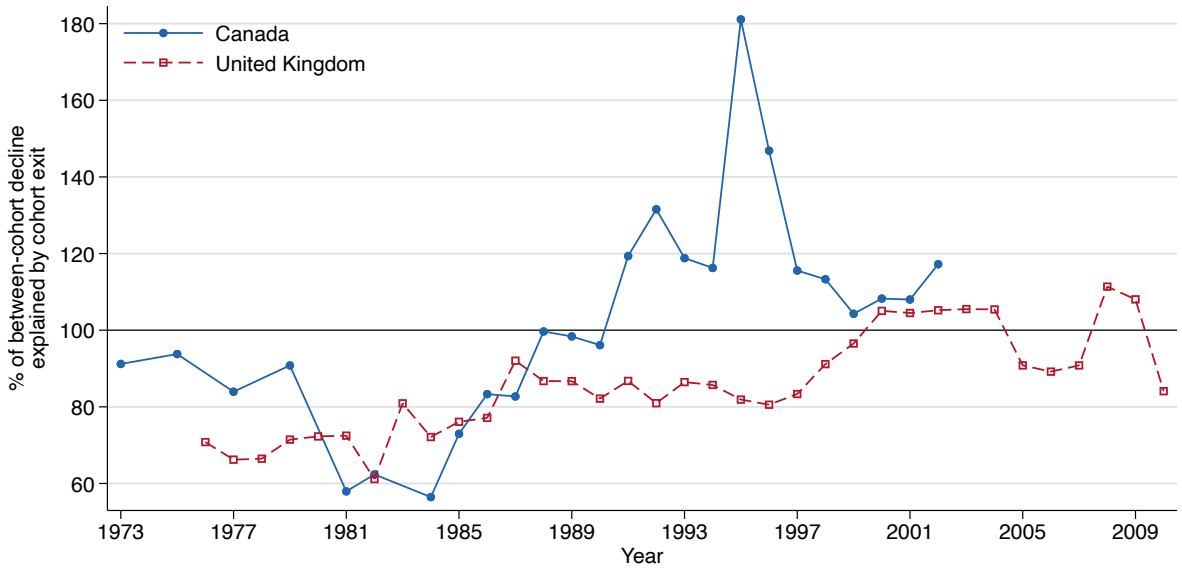


Panel D: Workers 31-40 in firms that existed in 1976



Notes: Panels A and B show the tenure distribution of workers who were between 41 and 50 years old in 1986 in Italy. The distribution of tenure is censored at 10 years because the database starts in 1976. In each year, the data encompass information about all workers who were between 41 and 50 years old, had worked in the private sector for at least 24 weeks, had earned strictly positive earnings, and had not retired by December 31. Moreover, Panel A further restricts the sample to workers who were employed by firms that (i) had more than five total employees in the period 1976-78, (ii) at least one man and one woman under 30 years old in both the periods 1976-78 and 1986-88, and (iii) at least one worker over 40 in the period 1976-78. Panel B drops the requirement for firms to have more than five employees. Panel C and D show the same distributions of tenure for workers who were between 31 and 40 years old in 1986 in Italy. *Source for Italy:* UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

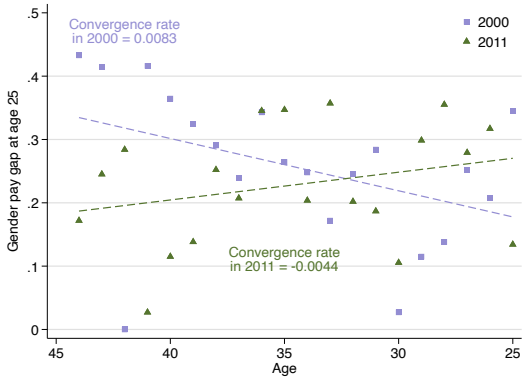
Figure A31: No Convergence in Entry Earnings in Other High-Income Countries



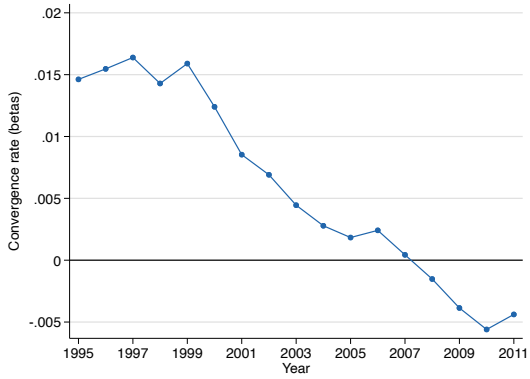
Notes: "Convergence at exit" computes what would have happened to the gender pay gap if the cross-cohort convergence in the earnings of men and women at age 25 had stopped in different benchmark years for Canada and the United Kingdom, respectively. The figure shows the ratio between the change (last year - first year) in the wage gap predicted by this new counterfactual scenario and the change in the gender gap predicted by the total between-cohort component from Equation (2) when the benchmark year t_b moves from the first sample year to 2002 for Canada and to 2010 for the UK. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks (not available in the UK), had earned strictly positive earnings. *Source for LIS data:* Luxembourg Income Study (LIS) Database, last accessed on 06/01/2023 at https://www.lisdatacenter.org/.

Figure A32: Between-Cohort Convergence in Other High-Income Countries

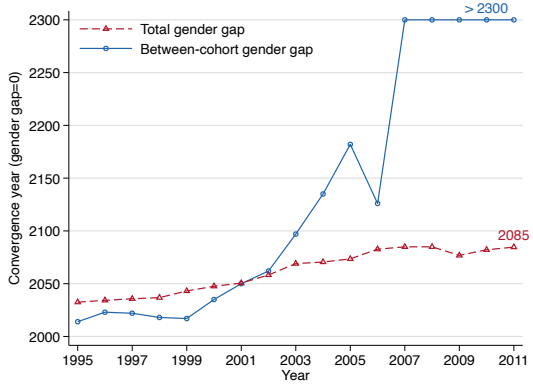
Panel A: Between-cohort trends (Canada, log weekly earnings)



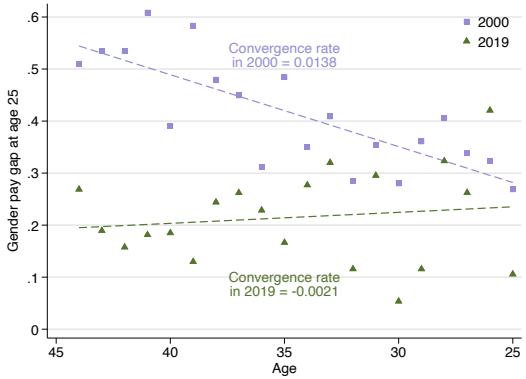
Panel B: Between-cohort convergence rate (Canada, log weekly earnings)



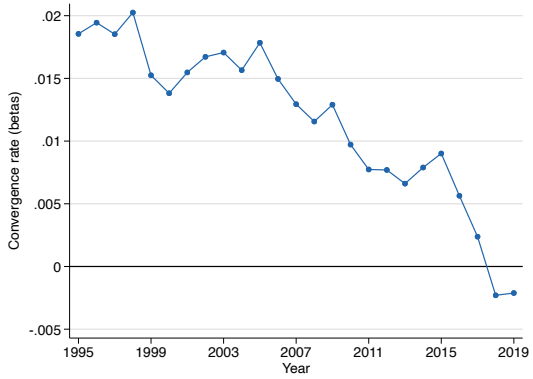
Panel C: Between-cohort convergence year (Canada, log weekly earnings)



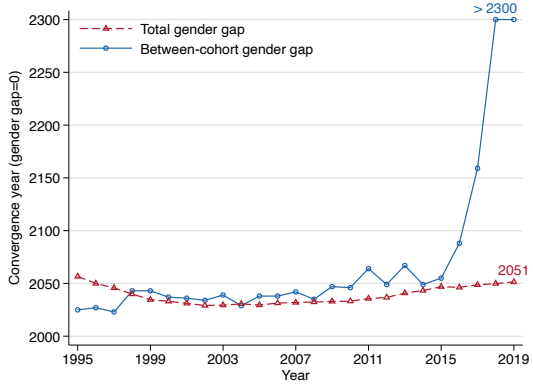
Panel D: Between-cohort trends (UK, log yearly earnings)



Panel E: Between-cohort convergence rate (UK, log yearly earnings)

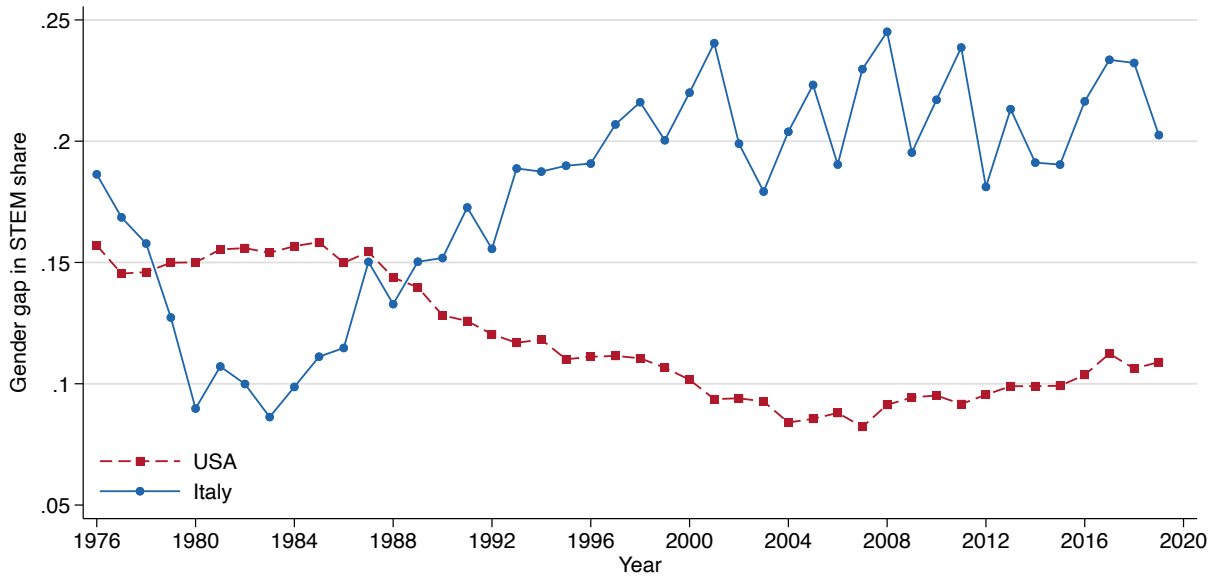


Panel F: Between-cohort convergence year (UK, log yearly earnings)



Notes: In each year t, we compute the gender pay gap at an early career stage for workers in age group a using their (weekly for Canada and yearly for the UK) earnings at age 25. Then, we estimate the linear relationship between the mean gender gap at labor-market entry and age (Equation (3)). Panels A and D show the best fit line in two different years for Canada and the United Kingdom, respectively. Panels B and E show the coefficients of age (β_t) for each year t. Panels C and F show the first cohort with gap at most equal to zero at age 25 predicted by Equation (3) for each year between 1976 and 2019 (1973 and 2011 for Canada). In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks (not available in the UK), had earned strictly positive earnings.
Source for LIS data: Luxembourg Income Study (LIS) Database, last accessed on 06/01/2023 at https://www.lisdatacenter.org/.

Figure A33: Gender Gap in the Share of STEM Graduates



Notes: The figure shows the gender gap (men − women) in the share of graduates in STEM subjects at age 25 in the United States and Italy. STEM subject areas are: Natural Sciences, Physics, Mathematics, Statistics, Computer Science, Information Engineering, Industrial Engineering, Architecture and Civil Engineering. *Source for Italy:* Quarterly Labour Force Survey, Istituto Nazionale di Statistica (ISTAT). *Source for the United States*: Integrated Public Use Microdata Series, American Community Survey. Minneapolis, MN: IPUMS.

Table A1: Characteristics of Data Sources

	# available years	# observations	# workers	# firms	Yearly earnings	Weekly earnings	Hourly earnings	Restrict working weeks
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Current Population Survey								
United States (1976-2019)	44	2,053,131	-	-	Yes	Yes	Yes	Yes
Panel B: Social Security data								
Italy (1976-2019)	44	373,117,856	32,112,786	5,174,323	Yes	Yes	No	Yes
Panel C: Survey data from the Luxembourg Income Study (LIS) Database								
Canada (1973-2019)	42	1,078,555	-	-	Yes	Yes	No	Yes
United Kingdom (1976-2019)	44	565,117	-	-	Yes	No	No	No

Source for Italy: Database UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for United States:* Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS. *Sources for Canada:* Survey of Consumer Finances (1973-1995); Survey of Labour and Income Dynamics (1996-2011); Canadian Income Survey (2012 and later). After 2011, data from Canada include a coarser categorization of age that does not allow us to study the outcomes of younger workers at a specific age (for example, at age 25). Hence, while we can plot the aggregate gender pay gap until 2019, the rest of the analysis can be performed only between 1973 and 2011. *Sources for United Kingdom:* Family Expenditure Survey (1991 and earlier); Family Resources Survey (1994 and later).

Table A2: Alternative Definitions of Labor-Market Entry

	Total change in gender gap (last year - first year)		Between-cohort change in gender gap (last year - first year)				Between-cohort change in gender gap (Shorter time series)			
	First year	Change	Earnings at age 25	Earnings at age 28	Earnings at age 30	Earnings b/w age 25 and age 30	Total change (last y. - 2015)	Earnings at age 25	Total change (last y. - 1980)	True entry earnings
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Log weekly earnings										
United States (1976-2019)	0.800	-0.470	-0.598	-0.606	-0.615	-0.575	-0.025	-0.034	-	-
Italy (1976-2019)	0.333	-0.190	-0.241	-0.218	-0.197	-0.218	-0.012	-0.014	-0.157	-0.280
Canada (1973-2011)	0.672	-0.288	-0.430	-0.507	-0.326	-0.335	-	-	-	-
Panel B: Log yearly earnings										
United States (1976-2019)	0.835	-0.495	-0.621	-0.645	-0.657	-0.598	-0.029	-0.035	-	-
Italy (1976-2019)	0.350	-0.092	-0.204	-0.165	-0.127	-0.166	-0.012	-0.016	-0.061	-0.270
Canada (1973-2011)	0.695	-0.308	-0.460	-0.621	-0.397	-0.359	-	-	-	-
United Kingdom (1976-2019)	1.058	-0.627	-0.750	-0.629	-0.622	-0.669	-0.065	-0.038	-	-
Panel C: Log hourly earnings										
United States (1976-2019)	0.568	-0.361	-0.452	-0.464	-0.458	-0.441	-0.012	-0.024	-	-

Notes: Columns 1 and 2 show the change in gender gap between the first and last available years for each country. Columns 3 to 6 show the between-cohort component of the change in the gender gap between the first and last available years (Equation (2)). In this counterfactual scenario, we assign to each cohort (defined as a combination of birth year and gender) its mean earnings (weekly, yearly, or hourly) at labor-market entry in every year. These columns differ in the definition of labor-market entry: mean earnings at age 25 (Col. 3), at age 28 (Col. 4), at age 30 (Col. 5), and between age 25 and age 30 (Col. 6). Cohorts who are above these age thresholds in the first sample year receive their mean earnings in the first sample year. In Column 8, we assign to each cohort its mean earnings at age 25, starting from the year in which they are observable for all cohorts, including older ones (2015). This analysis can be performed only with data from Italy, the United States, and the United Kingdom. Column 7 shows the total change in gender gap between 2015 and the last available year. In Column 10, we assign to each cohort its true earnings at labor-market entry, rather than its mean earnings at age 25. This analysis is available only for Italy and only from 1980. Column 9 shows the total change in gender gap between 1980 and the last available year. In this analysis. the time series for Canada stops in 2011, the last sample year in which the exact age of each worker is available. In each year, the data encompass information about all workers who were between 25 and 64 years old, had worked in the private sector for at least 24 weeks (not available in the UK), had earned strictly positive earnings, and (only in Italy) had not retired by December 31 (only in Italy). *Source for Italy:* UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS). *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS. *Source for LIS data:* Luxembourg Income Study (LIS) Database, last accessed on 06/01/2023 at https://www.lisdatacenter.org/.

Table A3: First Stage of Firm-Level IV Regressions

	△ share over-50 workers	
	(1)	(2)
Panel A: 1976-1986		
Share workers 41-50 − share workers 51-60 in 1976	0.58***	0.58***
	(0.01)	(0.01)
Mean dep. var.	0.01	0.01
SD dep. var.	0.10	0.10
Obs.	25,536	25,536
Panel B: 1976-1996		
Share workers 31-40 − share workers 51-60 in 1976	0.45***	0.46***
	(0.01)	(0.01)
Mean dep. var.	-0.00	-0.00
SD dep. var.	0.11	0.11
Obs.	14,064	14,064
Panel C: 1996-2016		
Share workers 31-40 − share workers 51-60 in 1996	0.39***	0.40***
	(0.01)	(0.01)
Mean dep. var.	0.16	0.16
SD dep. var.	0.11	0.11
Obs.	14,709	14,709
Controls	No	Yes

Notes: This table shows the results of OLS firm-level regressions in which changes in exposure to workforce aging are regressed on characteristics of the firm-level age distribution at baseline. Specifically, Δ share of workers 51-60 is regressed on the difference between the share of workers aged 41-50 (Panel A) or 31-40 (Panels B and C) and the share of workers 51-60 at baseline, all observed at the firm level. As such, these regressions measure the portion of firm-level workforce aging that stems from the age distribution of a firm's workforce at baseline, rather than from firms' or workers' choices during the period under consideration. Controls include firm age, firm size, mean firm pay, and province fixed effects, all observed at baseline. Firm-level values for year x are computed as three-year averages over x and $x+2$. In each panel, the sample of firms include all firms that had more than five total employees at baseline, at least one man and one woman under 30 years old in both the initial and final period, as well as at least one worker over 40 in the initial period.
Source for Italy: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

Table A4: Firm-Level Workforce Aging and the Gender Pay Gap, Alternative Outcomes

	Gender gap in weekly earnings at 25-30						Gender gap in log weekly earnings at 25-30					
	OLS		Reduced form		2SLS		OLS		Reduced form		2SLS	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: 1976-1986												
△ share older workers	-9.35	-9.73	-11.43**	-11.71**	-19.74**	-20.05**	-0.01	-0.01	-0.03*	-0.02*	-0.05*	-0.04*
	(6.05)	(5.97)	(4.93)	(5.15)	(8.48)	(8.84)	(0.02)	(0.02)	(0.01)	(0.01)	(0.02)	(0.02)
KP F-stat					7,331	11,166					7,331	11,166
Mean dep. var.	-7.41	-7.41	-7.41	-7.41	-7.41	-7.41	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03
SD dep. var.	111.73	111.73	111.73	111.73	111.73	111.73	0.31	0.31	0.31	0.31	0.31	0.31
Obs.	25,536	25,536	25,536	25,536	25,536	25,536	25,536	25,536	25,536	25,536	25,536	25,536
Panel B: 1976-1996												
△ share older workers	-31.21***	-24.12**	-19.81**	-16.02**	-43.85**	-35.18**	-0.04	-0.04	-0.02	-0.03*	-0.05	-0.07*
	(10.34)	(11.20)	(7.67)	(7.19)	(17.02)	(15.82)	(0.03)	(0.03)	(0.02)	(0.02)	(0.05)	(0.04)
KP F-stat					4,597	4,392					4,597	4,392
Mean dep. var.	-16.51	-16.51	-16.51	-16.51	-16.51	-16.51	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07
SD dep. var.	121.12	121.12	121.12	121.12	121.12	121.12	0.30	0.30	0.30	0.30	0.30	0.30
Obs.	14,064	14,064	14,064	14,064	14,064	14,064	14,064	14,064	14,064	14,064	14,064	14,064
Panel C: 1996-2016												
△ share older workers	-8.68	-7.53	-13.56	-12.08	-35.19*	-30.51	-0.02	-0.02	-0.02	-0.02	-0.05	-0.05
	(9.25)	(9.19)	(8.23)	(8.37)	(21.25)	(21.12)	(0.02)	(0.02)	(0.02)	(0.02)	(0.04)	(0.04)
KP F-stat					2,245	3,397					2,245	3,397
Mean dep. var.	-3.50	-3.50	-3.50	-3.50	-3.50	-3.50	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01
SD dep. var.	130.67	130.67	130.67	130.67	130.67	130.67	0.26	0.26	0.26	0.26	0.26	0.26
Obs.	14,709	14,709	14,709	14,709	14,709	14,709	14,709	14,709	14,709	14,709	14,709	14,709
Controls	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes

Notes: This table shows results of OLS and IV firm-level regressions in which changes in the gender pay gap among younger workers (difference in weekly earnings or log weekly earnings between men and women between 25 and 30 years old) are regressed on changes in exposure to workforce aging. The "OLS" columns regress Δ gender gap in pay rank on Δ share of workers 51-60, both observed at the firm level. The "Reduced form" columns regress Δ gender pay rank gap on the difference between the share of workers aged 41-50 (Panel A) or 31-40 (Panels B and C) and the share of workers 51-60 at baseline. The "2SLS" columns regress Δ gender pay rank gap on Δ share of workers 51-60, instrumenting the latter with the difference between the share of workers aged 41-50 (Panel A) or 31-40 (Panels B and C) and the share of workers 51-60 at baseline. Controls include firm age, firm size, mean firm pay, and province fixed effects, all observed at baseline. Firm-level values for year x are computed as three-year averages over x and $x+2$. In each panel, the sample of firms include all firms that had more than five total employees at baseline, at least one man and one woman under 30 years old in both the initial and final period, as well as at least one worker over 40 in the initial period. *Source for Italy:* UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

B A Model of Career Spillovers

Solving the firm problem. The firm problem is

$$\max_{m_{y,t},f_{y,t},m_{y,b},f_{y,b}} AY\left(L_y, L_o\right) - \sum_{g\in\{m,f\}}\sum_{a\in\{y,o\}}\sum_{j\in\{t,b\}}\left(w^g_{a,j}g_{a,j}\right) - \frac{\kappa}{2}K^2 - \sum_{g\in\{m,f\}}\left(\frac{c_g}{2}g^2_{y,t}\right).$$

The FOCs are

$$\begin{aligned}
\left\{f_{y,t}\right\}: &\quad AY_{L_y}\theta_{y,t} - w^f_{y,t} - \kappa K - c_f f_{y,t} &= 0;\\
\left\{m_{y,t}\right\}: &\quad AY_{L_y}\theta_{y,t} - w^m_{y,t} - \kappa K - c_m m_{y,t} &= 0;\\
\left\{f_{y,b}\right\}: &\quad AY_{L_y}\theta_{y,b} - w^f_{y,b} &= 0;\\
\left\{m_{y,b}\right\}: &\quad AY_{L_y}\theta_{y,b} - w^m_{y,b} &= 0.
\end{aligned}$$

The last two first order conditions indicate that the marginal revenue products of labor of younger men and women in the bottom job b are the same, hence their wages in the bottom job are also the same:

$$w^f_{y,b} = w^m_{y,b} = w_{y,b} = AY_{L_y}\theta_{y,b}.$$

Given that the wage in the bottom job pays a constant wedge over the wage in the bottom job, it follows that the wages of younger men and women in the top job t are also the same:

$$w^f_{y,t} = w^m_{y,t} = w_{y,t} = \mu_y w_{y,b}.$$

From the two initial first order conditions, we can pin down the optimal employment of younger men and women in the top job:

$$\begin{aligned}
f^*_{y,t} &= \frac{AY_{L_y}\left(\theta_{y,t} - \mu_y\theta_{y,b}\right) - \kappa K}{c_f};\\
m^*_{y,t} &= \frac{AY_{L_y}\left(\theta_{y,t} - \mu_y\theta_{y,b}\right) - \kappa K}{c_m}.
\end{aligned}$$

Given that $c_f > c_m$, we can conclude that the optimal number of younger women in top jobs is lower than the optimal number of men in top jobs: $f_{t,y} < m_{t,y}$. Furthermore, in equilibrium, the firm keeps the ratio of younger women and men in the top job constant:

$$\frac{f_{y,t}}{m_{y,t}} = \frac{c_m}{c_f} = \delta_f < 1.$$

Next, we consider an increase in the number of older workers in top jobs from period -1. The bottom wage of younger workers (for both gender groups) responds as follows:

$$\frac{\partial w_{y,b}}{\partial l^{-1}_{o,t}} = A\theta_{y,b}\left(Y_{L_yL_y}\frac{\partial L_y}{\partial l^{-1}_{o,t}} + Y_{L_yL_o}\frac{\partial L_o}{\partial l^{-1}_{o,t}}\right).$$

The derivatives of the efficiency units of younger and older labor with respect to $l^{-1}_{o,t}$ are:

$$\begin{aligned}
\frac{\partial L_y}{\partial l^{-1}_{o,t}} &= \frac{\partial\left[\theta_{y,t}\left(K - l_{o,t}\right) + \theta_{y,b}\left(l_y - K + l_{o,t}\right)\right]}{\partial l^{-1}_{o,t}} = \left(\theta_{y,t} - \theta_{y,b}\right)\left(\frac{\partial K}{\partial l^{-1}_{o,t}} - \rho_t\right)\\
\frac{\partial L_o}{\partial l^{-1}_{o,t}} &= \frac{\partial\left[\theta_{o,t}\rho_t l^{-1}_{o,t} + \theta_{o,b}\rho_b l^{-1}_{o,b}\right]}{\partial l^{-1}_{o,t}} = \theta_{o,t}\rho_t.
\end{aligned}$$

We can rewrite the change in the bottom wage as follows:

$$\begin{aligned}\frac{\partial w_{y,b}}{\partial l_{o,t}^{-1}} &= A\theta_{y,b}\left(Y_{L_yL_y}\frac{\partial L_y}{\partial l_{o,t}^{-1}} + Y_{L_yL_o}\frac{\partial L_o}{\partial l_{o,t}^{-1}}\right) \\ &= A\theta_{y,b}\left[Y_{L_yL_y}\left(\theta_{y,t}-\theta_{y,b}\right)\left(\frac{\partial K}{\partial l_{o,t}^{-1}}-\rho_t\right) + Y_{L_yL_o}\theta_{o,t}\rho_t\right].\end{aligned}$$

An increase in the supply of older workers causes negative career spillovers (or crowding out of younger workers from top spots) if $\frac{\partial l_{y,t}}{\partial l_{o,t}^{-1}} = \frac{\partial m_{y,t}}{\partial l_{o,t}^{-1}} + \frac{\partial f_{y,t}}{\partial l_{o,t}^{-1}} = \frac{\partial K}{\partial l_{o,t}^{-1}} - \rho_t < 0$, which implies that the total number of top jobs grows less that the number of older workers at the top. The key driver for the sign of this derivative is the change in the number of top jobs. We can write the number of top slots as follows:

$$K = \frac{AY_{L_y}\left(\theta_{y,t}-\mu_y\theta_{y,b}\right) - c_m m_{y,t}}{\kappa}.$$

Next, we write the derivative of K with respect to $l_{o,t}^{-1}$:

$$\frac{\partial K}{\partial l_{o,t}^{-1}} = \frac{A\left(\theta_{y,t}-\mu_y\theta_{y,b}\right)}{\kappa}\left[Y_{L_yL_y}\left(\theta_{y,t}-\theta_{y,b}\right)\left(\frac{\partial K}{\partial l_{o,t}^{-1}}-\rho_t\right) + Y_{L_yL_o}\theta_{o,t}\rho_t\right] - \frac{c_m}{\kappa}\frac{\partial m_{y,t}}{\partial l_{o,t}^{-1}}.$$

We use the fact that the ratio of younger men and women in top jobs is constant in equilibrium to rewrite the derivative of $m_{y,t}$ as a function of the derivative of K:

$$\begin{aligned}\text{[Step 1]}\quad \frac{\partial f_{y,t}}{\partial l_{o,t}^{-1}} &= \frac{c_m}{c_f}\frac{\partial m_{y,t}}{\partial l_{o,t}^{-1}}. \\ \text{[Step 2]}\quad \frac{\partial m_{y,t}}{\partial l_{o,t}^{-1}} &= \frac{\partial K}{\partial l_{o,t}^{-1}} - \frac{\partial f_{y,t}}{\partial l_{o,t}^{-1}} - \rho_t \\ &= \frac{\partial K}{\partial l_{o,t}^{-1}} - \frac{c_m}{c_f}\frac{\partial m_{y,t}}{\partial l_{o,t}^{-1}} - \rho_t \\ &= \frac{c_f}{c_f+c_m}\left(\frac{\partial K}{\partial l_{o,t}^{-1}}-\rho_t\right).\end{aligned}$$

Going back to the derivative of K, we can rewrite it as follows:

$$\frac{\partial K}{\partial l_{o,t}^{-1}} = \frac{1}{\kappa}\left\{A\left(\theta_{y,t}-\mu_y\theta_{y,b}\right)\left[Y_{L_yL_y}\left(\theta_{y,t}-\theta_{y,b}\right)\left(\frac{\partial K}{\partial l_{o,t}^{-1}}-\rho_t\right) + Y_{L_yL_o}\theta_{o,t}\rho_t\right] - \frac{c_m c_f}{c_f+c_m}\left(\frac{\partial K}{\partial l_{o,t}^{-1}}-\rho_t\right)\right\}.$$

We simplify the notation:

$$\begin{aligned}B &= A\left(\theta_{y,t}-\mu_y\theta_{y,b}\right) > 0 \\ D &= \left(\theta_{y,t}-\theta_{y,b}\right) > 0 \\ E &= \frac{c_m c_f}{c_f+c_m} > 0.\end{aligned}$$

Then, we can further simplify the derivative of K as follows:

$$\begin{aligned}\frac{\partial K}{\partial l_{o,t}^{-1}} &= \frac{1}{\kappa}\left\{B\left[Y_{L_yL_y}D\left(\frac{\partial K}{\partial l_{o,t}^{-1}}-\rho_t\right) + Y_{L_yL_o}\theta_{o,t}\rho_t\right] - E\left(\frac{\partial K}{\partial l_{o,t}^{-1}}-\rho_t\right)\right\} \\ &= \underbrace{\frac{1}{\kappa - Y_{L_yL_y}BD - E}\left\{B\left[Y_{L_yL_o}\theta_{o,t} - Y_{L_yL_y}D\right] - E\right\}}_{\text{Negative career spillovers if } <1}\rho_t.\end{aligned}$$

If the term multiplying ρ_t is less than 1, an increase in the number of older workers in top jobs decreases the number of top slots available to younger workers. This scenario happens when the following condition holds:

$$\begin{aligned} 1 &> \frac{1}{\kappa - Y_{L_yL_y}BD - E}\left\{B\left[Y_{L_yL_o}\theta_{o,t} - Y_{L_yL_y}D\right] - E\right\} \\ \kappa - Y_{L_yL_y}BD - E &> B\left[Y_{L_yL_o}\theta_{o,t} - Y_{L_yL_y}D\right] - E \\ \kappa &> \bar{\kappa} = BY_{L_yL_o}\theta_{o,t} = A\left(\theta_{y,t} - \mu_y\theta_{y,b}\right)Y_{L_yL_o}\theta_{o,t} > 0. \end{aligned}$$

This inequality indicates that the cost parameter κ needs to be higher than the productivity gains that younger workers experience from their complementarity with older workers. The term on the right-hand side is greater than zero because $\left(\theta_{y,t} - \mu_y\theta_{y,b}\right) > 0$. The latter follows from the condition that guarantees a positive K (see the formula of the equilibrium K) and from the fact that $f_{y,t} = K - m_{y,t} - l_{t,o}$.

So, when $\kappa > \bar{\kappa}$, we can draw several conclusions. First, as already pointed out, there is crowding out of younger workers in top jobs: $\frac{\partial K}{\partial l_{o,t}^{-1}} - \rho_t < 0$. Second, all younger workers become less likely to hold top jobs, but younger men lose more top spots than younger women because they are more likely to hold top jobs at baseline: $\left|\frac{\partial f_{t,y}}{\partial l_{o,t}^{-1}}\right| = \delta_f \left|\frac{\partial m_{t,y}}{\partial l_{o,t}^{-1}}\right| < \left|\frac{\partial m_{t,y}}{\partial l_{o,t}^{-1}}\right|$ because $\delta_f < 1$. Third, a larger supply of older workers at the top raises the bottom wage of younger workers:

$$\frac{\partial w_{y,b}}{\partial l_{o,t}^{-1}} = A\theta_{y,b}\left[Y_{L_yL_y}\left(\theta_{y,t} - \theta_{y,b}\right)\left(\frac{\partial K}{\partial l_{o,t}^{-1}} - \rho_t\right) + Y_{L_yL_o}\theta_{o,t}\rho_t\right] > 0,$$

because $Y_{L_yL_y} < 0$, $\left(\frac{\partial K}{\partial l_{o,t}^{-1}} - \rho_t\right) < 0$, and $Y_{L_yL_o} > 0$.

Next, we address a change in the mean wage of younger men and women. The mean wages of younger men and women are:

$$\begin{aligned} \{f_y\} \quad \overline{w}_{y,f} &= \frac{f_{y,t}\cdot w_{y,t}^f + f_{y,b}\cdot w_{y,b}^f}{f_y}; \\ \{m_y\} \quad \overline{w}_{y,m} &= \frac{m_{y,t}\cdot w_{y,t}^m + m_{y,b}\cdot w_{y,b}^m}{m_y}, \end{aligned}$$

where $w_{y,b}^f = w_{y,b}^m = w_{y,b}$, $w_{y,t}^f = w_{y,t}^m = w_{y,t} = \mu_y w_{y,b}$, and $g_y = g_{y,t} + g_{y,b}$ for each $g \in \{m, f\}$. Therefore, we can rewrite them as follows:

$$\begin{aligned} \{f_y\} \quad \overline{w}_{y,f} &= \frac{f_{y,t}\cdot\mu_y w_{y,b} + \left(f_y - f_{y,t}\right)\cdot w_{y,b}}{f_y} = \frac{1}{f_y}\left(\mu_y - 1\right)f_{y,t}w_{y,b} + w_{y,b}; \\ \{m_y\} \quad \overline{w}_{y,m} &= \frac{1}{m_y}\left(\mu_y - 1\right)m_{y,t}w_{y,b} + w_{y,b}. \end{aligned}$$

Starting from younger men, we consider the change in the mean wage that stems from a marginal increase in the number of older workers who held top jobs in period -1 ($l_{t,o}^{-1}$). Under the empirically relevant scenario of $\kappa > \bar{\kappa}$, we find that:

$$\frac{\partial \overline{w}_{y,m}}{\partial l_{o,t}^{-1}} = \underbrace{\frac{1}{m_y}\left(\mu_y - 1\right)\frac{\partial m_{y,t}}{\partial l_{o,t}^{-1}}w_{y,b}}_{\text{career spillovers } <0} + \underbrace{\left(\frac{1}{m_y}\left(\mu_y - 1\right)m_{y,t} + 1\right)\frac{\partial w_{y,b}}{\partial l_{o,t}^{-1}}}_{\text{wage level } >0}.$$

The first component is negative because $\frac{\partial m_{y,t}}{\partial l_{o,t}^{-1}} < 0$ and $\mu_y - 1 > 0$. In contrast, the second component is positive because $\frac{\partial w_{b,y}}{\partial l_{t,o}^{-1}} > 0$. Due to the complementarity between the two age groups, a larger number of older workers makes younger men more productive, contributing to raise their mean wage.

The same decomposition applies to the mean wage of younger women:

$$\frac{\partial \bar{w}_{y,f}}{\partial l_{o,t}^{-1}} = \frac{1}{f_y}\left(\mu_y - 1\right)\frac{\partial f_{y,t}}{\partial l_{o,t}^{-1}} w_{y,b} + \left(\frac{1}{f_y}\left(\mu_y - 1\right) f_{y,t} + 1\right)\frac{\partial w_{y,b}}{\partial l_{o,t}^{-1}}.$$

Next, we compare the magnitude of the negative career spillovers between younger men and women. As long as the share of women in the top job is smaller than the equivalent share of men, the career spillovers have a more negative effect on the mean wage of younger men:

$$\begin{aligned}
\frac{1}{m_y}\left(\mu_y - 1\right)\frac{\partial m_{y,t}}{\partial l_{o,t}^{-1}} w_{y,b} &< \frac{1}{f_y}\left(\mu_y - 1\right)\frac{\partial f_{y,t}}{\partial l_{o,t}^{-1}} w_{y,b} \\
\frac{1}{m_y}\left(\mu_y - 1\right)\frac{\partial m_{y,t}}{\partial l_{o,t}^{-1}} w_{y,b} &< \frac{1}{f_y}\left(\mu_y - 1\right)\delta_f \frac{\partial m_{y,t}}{\partial l_{o,t}^{-1}} w_{y,b} \\
\frac{1}{m_y} &> \frac{1}{f_y}\delta_f \\
\frac{m_{y,t}}{m_y}\frac{1}{m_{y,t}} &> \frac{f_{y,t}}{f_y}\frac{1}{f_{y,t}}\delta_f \\
\frac{m_{y,t}}{m_y} \Big/ \frac{f_{y,t}}{f_y} &> \frac{m_{y,t}}{f_{y,t}}\delta_f \\
\frac{m_{y,t}}{m_y} \Big/ \frac{f_{y,t}}{f_y} &> 1.
\end{aligned}$$

The main takeaway from this exercise is that when an increase in the supply of older workers decreases the mean wages of younger men and women, the gender pay gap closes as long as men are more concentrated in top jobs at baseline.

Finally, we can show that an increase in the retention rate of older workers at the top and a decrease in the economy-wide level of economic growth generate similar consequences on the gender pay gap. Starting from an increase in the retention rate of older workers at the top, the bottom wages of younger men and women change as follows:

$$\frac{\partial w_{y,b}}{\partial \rho_t} = A\theta_{y,b}\left[Y_{L_y L_y}\left(\theta_{y,t} - \theta_{y,b}\right)\left(\frac{\partial K}{\partial \rho_t} - l_{o,t}^{-1}\right) + Y_{L_y L_o}\theta_{o,t} l_{o,t}^{-1}\right].$$

We simplify the notation:

$$\begin{aligned}
B &= A\left(\theta_{y,t} - \mu_y \theta_{y,b}\right) > 0 \\
D &= \left(\theta_{y,t} - \theta_{y,b}\right) > 0 \\
E &= \frac{c_m c_f}{c_f + c_m} > 0.
\end{aligned}$$

Then, we can rewrite the derivative of K as follows:

$$\begin{aligned}
\frac{\partial K}{\partial \rho_t} &= \frac{1}{\kappa}\left\{B\left[Y_{L_y L_y} D\left(\frac{\partial K}{\partial l_{o,t}^{-1}} - l_{o,t}^{-1}\right) + Y_{L_y L_o}\theta_{o,t} l_{o,t}^{-1}\right] - E\left(\frac{\partial K}{\partial l_{o,t}^{-1}} - l_{o,t}^{-1}\right)\right\} \\
&= \underbrace{\frac{1}{\kappa - Y_{L_y L_y} BD - E}\left\{B\left[Y_{L_y L_o}\theta_{o,t} - Y_{L_y L_y} D\right] - E\right\}}_{\text{Negative career spillovers if } <1} l_{o,t}^{-1}.
\end{aligned}$$

Therefore, there are negative career spillovers when $\kappa > \bar{\kappa} = A_f\left(\theta_{y,t} - \mu_y\theta_{y,b}\right) Y_{L_y L_o}\theta_{o,t}$, which is the same condition we found for an increase in the number of older workers at the top. All subsequent results follow through.

Furthermore, we can model κ as a decreasing function of the economic economic growth rate: $\kappa(g)$ with $\kappa'(g) < 0$. Bianchi and Paradisi (2023) shows that the condition on the third derivatives of the production function that is needed for a larger κ to generate more negative career spillovers. A decline in g increases κ and, therefore, lowers the response of K to a larger supply of older workers at the top, leading to more crowding out of younger workers in top positions. This result does not change the fact that older men are more affected than younger women as long as they are more concentrated in top jobs at baseline.

Gender gap in productivity. Instead of assuming that the firm faces different costs for employing younger men and women in the top job, we could assume that (i) women are less productive than men and (ii) men and women are imperfect substitutes in the top job. The first assumption could still be micro-founded based on either taste-based or statistical discrimination against women.

In this context, the firm chooses the number of younger men and women in both the top and bottom job in order to maximize its profits,

$$\max_{m_{y,t},f_{y,t},m_{y,b},f_{y,b}} AY\left(L_y,L_o\right)-\sum_{g=m,f}\sum_{a=y,o}\sum_{j=t,b}\left(w_{a,j}^{g}g_{a,j}\right)-\frac{\kappa}{2}K^2.$$

The efficiency units of younger and older labor become $L_a=\theta_{a,t}\left(m_{a,t}^{\frac{\sigma-1}{\sigma}}+\delta_f^{\frac{1}{\sigma}}f_{a,t}^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}+\theta_{a,b}\left(m_{a,b}+f_{a,b}\right)$, where $a\in\{y,o\}$ and $\delta_f<1$.

The first order conditions for the employment level in the bottom jobs are unchanged. Therefore, in equilibrium, we still have that $w_{y,b}^f=w_{y,b}^m=w_{y,b}=AY_{L_y}\theta_{y,b}$. Moreover, $w_{y,t}^f=w_{y,t}^m=w_{y,t}=\mu_y w_{y,b}$. In top jobs, we find that

$$\begin{aligned}\left\{f_{y,t}\right\}: &\quad AY_{L_y}\theta_{y,t}\left(m_{y,t}^{\frac{\sigma-1}{\sigma}}+\delta_f^{\frac{1}{\sigma}}f_{y,t}^{\frac{\sigma-1}{\sigma}}\right)^{\frac{1}{\sigma-1}}\delta_f^{\frac{1}{\sigma}}f_{y,t}^{-\frac{1}{\sigma}}-w_{y,t}^f-\kappa K &&=0;\\ \left\{m_{y,t}\right\}: &\quad AY_{L_y}\theta_{y,t}\left(m_{y,t}^{\frac{\sigma-1}{\sigma}}+\delta_f^{\frac{1}{\sigma}}f_{y,t}^{\frac{\sigma-1}{\sigma}}\right)^{\frac{1}{\sigma-1}}m_{y,t}^{-\frac{1}{\sigma}}-w_{y,t}^m-\kappa K &&=0.\end{aligned}$$

Combining the two first order conditions, we find that:

$$\begin{aligned}\delta_f^{\frac{1}{\sigma}}f_{y,t}^{-\frac{1}{\sigma}} &= m_{y,t}^{-\frac{1}{\sigma}}\\ \frac{f_{y,t}}{m_{y,t}} &= \delta_f<1.\end{aligned}$$

In equilibrium, the firm wants to keep a fixed ratio of younger men and women in top jobs. Moreover, compared with younger women, younger men are more concentrated at the top. All the main results from the baseline specification follow.

Resource constraint. Instead of assuming that there is a given number of costly top jobs, we assume that the firm faces a constraint on the amount of resources that it can spend for top jobs. The firm problem becomes:

$$\begin{aligned}\max_{m_{y,t},f_{y,t},m_{y,b},f_{y,b}} AY\left(L_y,L_o\right)&-\sum_{g\in\{m,f\}}\sum_{a\in\{y,o\}}\sum_{j\in\{t,b\}}\left(w_{a,j}^{g}g_{a,j}\right)\\ &-\sum_{g\in\{m,f\}}\left(\frac{c_g}{2}g_{y,t}^2\right)-\kappa\cdot\left(l_{o,t}+m_{y,t}+f_{y,t}\right),\end{aligned}$$

subject to the organizational resource constraint on top jobs ($\kappa\cdot\left(l_{o,t}+m_{y,t}+f_{y,t}\right)\leq K$). Here, K is the (exogenous) maximum amount of resources that can be spent on maintaining top jobs.

The first order conditions for bottom jobs are unchanged. Therefore, the wages of younger men and women in both bottom and top jobs are the same. The first order conditions for employment in top jobs are:

$$\begin{aligned}\left\{f_{y,t}\right\} &\quad AY_{L_y}\theta_{y,t}-w_{y,t}^f-c_f f_{y,t}-(1+\lambda)\cdot\kappa &&=0;\\ \left\{m_{y,t}\right\} &\quad AY_{L_y}\theta_{y,t}-w_{y,t}^m-c_m m_{y,t}-(1+\lambda)\cdot\kappa &&=0.\end{aligned}$$

In equilibrium, we conclude that:

$$\frac{f_{y,t}}{m_{y,t}} = \frac{c_m}{c_f}=\delta_f<1.$$

All the main results from the baseline specification follow.

Endogenous labor force participation. In this extension, we drop the assumption of fixed labor supply and directly model the choice of participating in the labor market. Specifically, we assume that younger workers work whenever their expected wage is above their reservation wage. Workers draw reservation wages from the following

cumulative distribution function:

$$P_g(w) = \left(\frac{w - w_{r,g}^{LB}}{w_{r,g}^{UB} - w_{r,g}^{LB}} \right)^{\eta_g},$$

for each $g \in \{m, f\}$, where $w_{r,g}^{LB}$ is the lower bound of reservation wages for workers of gender g, and $w_{r,g}^{UB}$ is the corresponding upper bound.

Specifically, we assume that younger workers choose whether to work or not for the representative firm based on the mean wage that the firm offers to younger workers ($\bar{w}_{y,g}$). They base their choice on the mean wage, rather than the actual wage they are going to receive when employed by the firm, because they are randomly assigned to either the top or bottom job after they join the company until the marginal product of labor equates to its cost. In this context, the number of employed younger men is equal to $P_m\left(\bar{w}_{y,m}\right) m_y$, while the number of employed younger women is $P_f\left(\bar{w}_{y,f}\right) f_y$.

The rest of the problem is unchanged. First, the firm receives the legacy older workers from period -1. Then, given a set of wages, the firm decides how many younger men and women to slot in the top and bottom jobs by equating the marginal revenue products of younger labor in the two positions to their marginal costs. In equilibrium, the market clears so that the demand for younger workers equates younger workers' supply: $g_y^d = P_g\left(\bar{w}_{y,g}\right) g_y$. Then, the firm allocates the younger workers randomly between the top and bottom jobs until its labor demands in the two positions are satisfied. Finally, the production is realized, and the firm pays all workers.

The labor-supply response to a change in the mean wage is:

$$\begin{aligned}
\frac{\partial P_g\left(\bar{w}_{y,g}\right) g_y}{\partial \bar{w}_{y,g}} &= \eta_g \frac{1}{w_{r,g}^{UB} - w_{r,g}^{LB}} \left(\frac{\bar{w}_{y,g} - w_{r,g}^{LB}}{w_{r,g}^{UB} - w_{r,g}^{LB}} \right)^{\eta_g - 1} g_y \\
&= \eta_g \frac{1}{w_{r,g}^{UB} - w_{r,g}^{LB}} P_g\left(\bar{w}_{y,g}\right) \left(\frac{\bar{w}_{y,g} - w_{r,g}^{LB}}{w_{r,g}^{UB} - w_{r,g}^{LB}} \right)^{-1} g_y \\
&= \eta_g \frac{P_g\left(\bar{w}_{y,g}\right)}{\bar{w}_{y,g} - w_{r,g}^{LB}} g_y.
\end{aligned}$$

We assume that the lower-bound reservation wage is zero, which is akin to assuming that there is a positive share of workers with no fixed cost from participating in the labor market, to further simplify the derivative:

$$\frac{\partial P_g\left(\bar{w}_{y,g}\right) g_y}{\partial \bar{w}_{y,g}} = \eta_g \frac{P_g\left(\bar{w}_{y,g}\right)}{\bar{w}_{y,g}} g_y.$$

The FOCs are unchanged. In equilibrium, wages of younger men and women are the same in both jobs. Younger men are more represented at the top. Moreover, the firm wants to keep a constant ratio $\delta_f < 1$ between younger women and men in top jobs.

Next, we are going to focus on the change in the mean wage of younger men and women when the number of older workers in the top job increases. The mean wage for younger workers of gender g can be written as follows:

$$\begin{aligned}
\overline{w}_{y,g} &= \frac{g_{y,t} \cdot w_{y,t}^g + g_{y,b} \cdot w_{y,b}^g}{P_g\left(\bar{w}_{y,g}\right) g_y} \\
&= \frac{g_{y,t} \cdot \mu_y w_{y,b} + \left(P_g\left(\bar{w}_{y,g}\right) g_y - g_{y,t}\right) \cdot w_{y,b}}{P_g\left(\bar{w}_{y,g}\right) g_y} \\
&= \frac{\left(\mu_y - 1\right) g_{y,t} \cdot w_{y,b}}{P_g\left(\bar{w}_{y,g}\right) g_y} + w_{y,b},
\end{aligned}$$

because $w_{y,b}^f = w_{y,b}^m = w_{y,b}$, $w_{y,t}^f = w_{y,t}^m = \mu_y w_{y,b}$, and $P_g\left(\bar{w}_{y,g}\right) g_y = g_{y,t} + g_{y,b}$ for each $g \in \{m, f\}$.

An increase in the number of older workers in top jobs changes the mean wage of younger workers of gender g as

follows:

$$\begin{aligned}\frac{\partial \bar{w}_{y,g}}{\partial l_{o,t}^{-1}} &= \frac{\left(P_g\left(\bar{w}_{y,g}\right) g_y\right)\left[\left(\mu_y-1\right)\frac{\partial g_{y,t}}{\partial l_{o,t}^{-1}} w_{y,b}+\left(\mu_y-1\right) g_{y,t}\frac{\partial w_{y,b}}{\partial l_{o,t}^{-1}}\right]}{\left(P_g\left(\bar{w}_{y,g}\right) g_y\right)^2} \\ &- \frac{\frac{\eta_g P_g\left(\bar{w}_{y,g}\right) g_y}{\bar{w}_{y,g}}\frac{\partial \bar{w}_{y,g}}{\partial l_{o,t}^{-1}}\left(\mu_y-1\right) g_{t,y} w_{y,b}}{\left(P_g\left(\bar{w}_{y,g}\right) g_y\right)^2}+\frac{\partial w_{y,b}}{\partial l_{o,t}^{-1}} \\ &= \frac{\left[\left(\mu_y-1\right)\frac{\partial g_{y,t}}{\partial l_{o,t}^{-1}} w_{y,b}+\left(\mu_y-1\right) g_{y,t}\frac{\partial w_{y,b}}{\partial l_{o,t}^{-1}}\right]}{\left(P_g\left(\bar{w}_{y,g}\right) g_y\right)} \\ &- \frac{\frac{\eta_g}{\bar{w}_{y,g}}\frac{\partial \bar{w}_{y,g}}{\partial l_{o,t}^{-1}}\left(\mu_y-1\right) g_{y,t} w_{y,b}}{\left(P_g\left(\bar{w}_{y,g}\right) g_y\right)}+\frac{\partial w_{y,b}}{\partial l_{o,t}^{-1}}.\end{aligned}$$

Further simplifying and rewriting this formula leads to the usual separation of the total effect into a crowding-out component and a relative labor supply effect:

$$\begin{aligned}\frac{\partial \bar{w}_{y,g}}{\partial l_{t,o}} &= \left(\frac{P_g\left(\bar{w}_{y,g}\right) g_y}{P_g\left(\bar{w}_{y,g}\right) g_y+\frac{\eta_g}{\bar{w}_{y,g}}\left(\mu_y-1\right) g_{y,t} w_{y,b}}\right)\left\{\frac{\left(\mu_y-1\right) w_{y,b}}{\left(P_g\left(\bar{w}_{y,g}\right) g_y\right)}\frac{\partial g_{y.t}}{\partial l_{o,t}^{-1}}\right. \\ &\left.+\left(\frac{\left(\mu_y-1\right) g_{y,t}}{\left(P_g\left(\bar{w}_{y,g}\right) g_y\right)}+1\right)\frac{\partial w_{y,b}}{\partial l_{o,t}^{-1}}\right\}.\end{aligned}$$

The crowding-out component for $g \in \{m, f\}$ becomes:

$$\left(\frac{1}{P_g\left(\bar{w}_{y,g}\right) g_y+\frac{\eta_g}{\bar{w}_{y,g}}\left(\mu_y-1\right) g_{y,t} w_{y,b}}\right)\left(\mu_y-1\right) w_{y,b}\frac{\partial g_{y,t}}{\partial l_{o,t}^{-1}}.$$

When there is crowding out of younger workers in top jobs, this component is negative because $\frac{\partial g_{y,t}}{\partial l_{o,t}^{-1}} < 0$. Moreover, the negative crowding-out effect on the mean wages of younger men is larger in magnitude than that for younger women when the following inequality holds:

$$\begin{aligned}\delta_f &< \frac{1+\frac{\eta_f}{\bar{w}_{y,f}}\left(\mu_y-1\right)\frac{f_{y,t}}{P_f\left(\bar{w}_{y,f}\right) f_y} w_{y,b}}{1+\frac{\eta_m}{\bar{w}_{y,m}}\left(\mu_y-1\right)\frac{m_{y,t}}{P_m\left(\bar{w}_{y,m}\right) m_y} w_{y,b}}, \\ 1-\delta_f &> \left(\mu_y-1\right) w_{y,b}\left(\delta_f\frac{\eta_m}{\bar{w}_{y,m}}\frac{m_{y,t}}{P_m\left(\bar{w}_{y,m}\right) m_y}-\frac{\eta_f}{\bar{w}_{y,f}}\frac{f_{y,t}}{P_f\left(\bar{w}_{y,f}\right) f_y}\right) \\ 1-\delta_f &> \left(\mu_y-1\right) w_{b,y}\frac{f_{y,t}}{P_f\left(\bar{w}_{y,f}\right) f_y}\left(\frac{\eta_m}{\bar{w}_{y,m}}\frac{P_f\left(\bar{w}_{y,f}\right) f_y}{P_m\left(\bar{w}_{y,m}\right) m_y}-\frac{\eta_f}{\bar{w}_{y,f}}\right).\end{aligned}$$

This inequality is more likely to hold when (i) the share of younger women in top jobs is lower, (ii) younger men are more likely to be employed than younger women, and (iii) the employment elasticity of younger women is larger than that of younger men. All these conditions are likely to find empirical support. For example, the idea that younger women's employment is more elastic than that of younger men is supported by several prior papers (for example, Bianchi, Gudmundsson, and Zoega (2001), Eissa and Hoynes (2004), and Manoli and Weber (2011)).

No exogenous rents. It is possible to drop the assumption that top-job wages pay an exogenous rent $\mu_y > 1$ over bottom-job wages. However, dropping this assumption requires further modifications to the baseline model.

The key takeaway from a model without exogenous rents is that the difference in the top wages of younger men and women reflects the difference in hiring costs between the two genders. However, since the firm can price discriminate between the two groups, it does not necessarily need to choose difference quantities of younger men and women in top

jobs. In this scenario, there are multiple equilibria that satisfy the following condition:

$$w^m_{y,t} - w^f_{y,t} = c_f f_{y,t} - c_m m_{y,t}.$$

We can restore a positive gender gap in job allocations in favor of younger men by making the labor supply endogenous to the level of wages, as outlined in the section above. Together with a linear cost of hiring, this assumption ensures that a positive difference in the wages of men and women translates into a positive difference in the concentration of younger men and women in top jobs. We further explore this scenario when we discuss an extension with heterogeneous firms.

Introducing skills. When all top jobs are already occupied by older workers, an increase in the supply of older workers has the following effect on the mean wage of younger men:

$$\begin{aligned}\frac{\partial \bar{w}_{y,m}}{\partial l^{-1}_{o,t}} &= \sum_{s\in\{h,l\}} \left(\frac{1}{m_y}\left(\mu_y - 1\right) m_{y,t,s} + 1\right) \frac{\partial w_{y,b,s}}{\partial l^{-1}_{o,t}} \\ &= \sum_{s\in\{h,l\}} \left(\frac{m_{y,s}}{m_y}\right) A\theta_{y,b,s} Y_{L_{y,s}L_o} \theta_{o,t,s}\rho_t,\end{aligned}$$

where $\frac{1}{m_y}\left(\mu_y - 1\right) m_{y,t,s} + 1 = \frac{\mu m_{y,t,s} + m_{y,b,s}}{m_y} = \frac{m_{y,s}}{m_y}$ because all top jobs are occupied by older workers and, therefore, $m_{y,t,s} = 0$, and $m_{y,b,s} = m_{y,s}$. Moreover, $\frac{\partial w_{y,b,s}}{\partial l^{-1}_{o,t}} = A\theta_{y,b,s} \cdot \left(Y_{L_{y,s}L_{y,s}} \frac{\partial L_{y,s}}{\partial l^{-1}_{t,o}} + Y_{L_{y,s}L_o} \frac{\partial L_{o,s}}{\partial l^{-1}_{t,o}}\right) = A\theta_{y,b,s} Y_{L_{y,s}L_o} \theta_{o,t,s}\rho_t$ because $\frac{\partial L_{y,s}}{\partial l^{-1}_{o,t}} = 0$ for all skills/tasks s. Moreover, we assume that the retention rate of older workers in top jobs (ρ_t) is the same across all tasks.

We repeat the same computation for younger women:

$$\frac{\partial \bar{w}_{y,f}}{\partial l^{-1}_{o,t}} = \sum_{s\in\{h,l\}} \left(\frac{f_{y,s}}{f_y}\right) A\theta_{y,b,s} Y_{L_{y,s}L_o} \theta_{o,t,s}\rho_t.$$

For both younger men and women, the change in mean wages is positive. If we compare the two derivatives, we find that an increase in the number of older workers in top jobs shrinks the gender pay gap if the wage increase in larger for younger women:

$$\sum_{s\in\{h,l\}} \left(\frac{f_{y,s}}{f_y}\right) A\theta_{y,b,s} Y_{L_{y,s}L_o} \theta_{o,t,s}\rho_t > \sum_{s\in\{h,l\}} \left(\frac{m_{y,s}}{m_y}\right) A\theta_{y,b,s} Y_{L_{y,s}L_o} \theta_{o,t,s}\rho_t.$$

This inequality crucially depends on the distribution of younger men and women across different skills/tasks. We can further show this point by assuming that the cross-complementarity between younger and older workers is proportional to a task's marginal product: $Y_{L_{y,s}L_o} = Y_{L_{y,s}} \cdot C\left(\mathbf{L_y}, L_o\right)$, where $C\left(\cdot,\cdot\right)$ is the same across tasks.[42] In this case, the

[42] This is true if the production function is a Cobb-Douglas or a nested CES. Indeed, for the Cobb-Douglas we have $Y = L_o^{\alpha_o} \cdot \prod_{s\in\{h,l\}} L_{a,s}^{\alpha_{a,s}}$ and

$$Y_{L_{y,s}L_o} = \alpha_{y,s}\alpha_o \frac{Y}{L_{y,s}L_o} = Y_{L_{y,s}} \cdot \frac{\alpha_o}{L_o},$$

where $\frac{\alpha_o}{L_o}$ is the same across all skills. In the nested CES case $Y = \left[\beta_y \left(L_{y,h}^{\rho_y} + L_{y,l}^{\rho_y}\right)^{\frac{\rho}{\rho_y}} + \beta_o L_o^{\rho}\right]^{\frac{1}{\rho}}$ and

$$\begin{aligned}Y_{L_{y,s}L_o} &= \beta_y L_{y,s}^{\rho_y - 1}\left(L_{y,h}^{\rho_y} + L_{y,l}^{\rho_y}\right)^{\frac{\rho}{\rho_y}-1}\left[\beta_y\left(L_{y,h}^{\rho_y} + L_{y,l}^{\rho_y}\right)^{\frac{\rho}{\rho_y}} + \beta_o L_o^{\rho}\right]^{\frac{1}{\rho}-2}\left(\frac{1}{\rho} - 1\right)\beta_o\rho L_o^{\rho-1} \\ &= \underbrace{\beta_y L_{y,s}^{\rho_y - 1}\left(L_{y,h}^{\rho_y} + L_{y,l}^{\rho_y}\right)^{\frac{\rho}{\rho_y}-1}\left[\beta_y\left(L_{y,h}^{\rho_y} + L_{y,l}^{\rho_y}\right)^{\frac{\rho}{\rho_y}} + \beta_o L_o^{\rho}\right]^{\frac{1}{\rho}-1}}_{=Y_{L_{y,s}}}\underbrace{\left(\frac{1}{\rho} - 1\right)\beta_o\rho L_o^{\rho-1}\left[\beta_y\left(L_{y,h}^{\rho_y} + L_{y,l}^{\rho_y}\right)^{\frac{\rho}{\rho_y}} + \beta_o L_o^{\rho}\right]^{-1}}_{=C\left(\mathbf{L_y}, L_o\right)},\end{aligned}$$

which proves the result.

change in mean wages for younger women is larger than that for younger men if:

$$\frac{f_{y,h}}{f_y}Y_{L_{y,h}}\theta_{y,b,h}\theta_{o,t,h}+\frac{f_{y,l}}{f_y}Y_{L_{y,l}}\theta_{y,b,l}\theta_{o,t,l}>\frac{m_{y,h}}{m_y}Y_{L_{y,h}}\theta_{y,b,h}\theta_{o,t,h}+\frac{m_{y,l}}{f_y}Y_{L_{y,l}}\theta_{y,b,l}\theta_{o,t,l}$$

$$\left(\frac{f_{y,h}}{f_y}-\frac{m_{y,h}}{m_y}\right)\left(Y_{L_{y,h}}\theta_{y,b,h}\theta_{o,t,h}-Y_{L_{y,l}}\theta_{y,b,l}\theta_{o,t,l}\right)>0.$$

Assuming that high-skill tasks have higher marginal product than low-skill tasks ($Y_{L_{y,h}}\theta_{y,b,h}\theta_{o,t,h}-Y_{L_{y,l}}\theta_{y,b,l}\theta_{o,t,l}>0$), the inequality holds if women are overrepresented in high-skill tasks: $\frac{f_{y,h}}{f_y}>\frac{m_{y,h}}{m_y}$.

Heterogeneous firms. First, we replace the representative firm with N firms, but each firm is small and does not internalize the consequences of its actions on other firms. We further assume that $\rho_{j,n}$ increases with firm-level productivity A_n of firm n (Antwi and Phillips, 2013; Ruffini, 2022). Second, firms set wages for the bottom and top jobs, instead of taking them as given. Third, the ratio of top and bottom wages is not equal to a fixed rent and is not necessarily constant across firms. Fourth, we assume that $\kappa(K_n)$ is 0 up to a threshold level $\overline{K}_n$ and then is ∞ beyond $\overline{K}_n$. In practice, this is equivalent to a binding constraint on the number of top jobs: $l_{o,t,n}+m_{y,t,n}+f_{y,t,n}=\bar{K}_n$. Fifth, to make the computations more tractable, we assume that the cost of hiring younger workers in the top job is linear, instead of quadratic.

The timing of the model is as follows. First, each firm receives legacy older workers from period -1. Then, each firm posts wage offers for its bottom and top jobs, and each younger worker joins the firm and job that maximizes her utility. Finally, the production is realized, and the firm makes payments to all workers. In this scenario, the firm problem is to choose the wages of younger men and women in the top and bottom job that maximize its profits,

$$\max_{w^m_{y,t,n},w^f_{y,t,n},w^m_{y,b,n},w^f_{y,b,n}} A_nY\left(L_{y,n},L_{o,n}\right)-\sum_{g\in\{m,f\}}\sum_{a\in\{y,o\}}\sum_{j\in\{t,b\}}\left(w^g_{a,j,n}g_{a,j,n}\right)-\sum_{g\in\{m,f\}}c_gg_{y,t,n},$$

subject to

$$l_{o,t,n}+m_{y,t,n}+f_{y,t,n}\leq\bar{K}_n.$$

On the worker side, we assume that a worker i of age group $a\,(i)$ and gender $g\,(i)$ derives the following utility when working in job j and firm n:

$$U_{i,a,j,n}=\log\left(w^g_{a,j,n}\right)+\xi_{i,a,j,n},$$

where $\xi_{i,a,j,n}$ represents the idiosyncratic preference of worker i over job j of firm n. We assume that $\xi_{i,a,j,n}$, which is unobserved by firms, follows a type-1 extreme distribution with a parameter σ that captures the degree of substitutability across jobs and firms in workers' preferences. In this context, firm n faces the following labor supply function for its job j from younger workers of gender g:

$$g_{y,j,n}=\frac{\left(w^g_{y,j,n}\right)^{\frac{1}{\sigma}}}{\sum_{n=1}^{N}\sum_{j\in\{t,b\}}\left(w^g_{y,j,n}\right)^{\frac{1}{\sigma}}}g_y.$$

The marginal change in the wage of job j and firm n has the following effect on the labor supply of younger workers of gender g anticipated by the firm:

$$\begin{aligned}\frac{\partial g_{y,j,n}}{\partial w^g_{y,j,n}}&=\frac{1}{\sigma}\frac{\left(w^g_{y,j,n}\right)^{\frac{1}{\sigma}-1}}{\sum_{n=1}^{N}\sum_{j\in\{t,b\}}\left(w^g_{y,j,n}\right)^{\frac{1}{\sigma}}}g_y\\&=\frac{1}{\sigma}\frac{g_{y,j,n}}{w^g_{y,j,n}}.\end{aligned}$$

In the first row, the denominator is left unchanged due to the assumption that firms do not anticipate the effect of a change in their wage on the other wages in the economy (Card et al., 2018; Lamadon, Mogstad, and Setzler, 2022).

Firms choose optimal wages as follows:

$$\left\{w^g_{y,b,n}\right\}:A_nY_{L_{y,n}}\theta_{y,b,n}\frac{\partial g_{y,b,n}}{\partial w^g_{y,b,n}}-w^g_{y,b,n}\frac{\partial g_{y,b,n}}{\partial w^g_{y,b,n}}-g_{y,b,n}=0$$

$$\left\{w^g_{y,t,n}\right\} : A_n Y_{L_{y,n}} \theta_{y,t,n} \frac{\partial g_{y,t,n}}{\partial w^g_{y,t,n}} - w^g_{y,t,n} \frac{\partial g_{y,t,n}}{\partial w^g_{y,t,n}} - g_{y,t,n} - c_g \frac{\partial g_{y,t,n}}{\partial w^g_{y,t,n}} - \lambda_n \frac{\partial g_{y,t,n}}{\partial w^g_{y,t,n}} = 0,$$

where λ_n is the multiplier on the constraint on the quantity of top jobs. Therefore, wages in the bottom job are equal to:

$$w^g_{y,b,n} = \underbrace{\frac{1}{1+\sigma}}_{\text{Markdown}} A_n Y_{L_{y,n}} \theta_{y,b,n},$$

so that bottom wages pay a markdown below the marginal product of labor, and $w^f_{y,b,n} = w^m_{y,b,n}$ for all firms.

Instead, wages in the top job are equal to:

$$w^g_{y,t,n} = \underbrace{\frac{1}{1+\sigma}}_{\text{Markdown}} \left(A_n Y_{L_{y,n}} \theta_{y,t,n} - c_g - \lambda_n\right).$$

Again, wages in the top job pay a markdown below the marginal product of labor in the top job. Moreover, as we discussed in the version with a representative firm, younger men are paid more than younger women in top jobs ($c_m < c_f$) and, therefore, are more likely to hold these positions. Finally, all firms have $\overline{K_n}$ top jobs.

Next, we consider the effect of an increase in the economy-wide number of older workers on wages and the number of top slots. Specifically, we study a marginal increase in $l^{-1}_{o,t}$, the total number of older workers in top jobs in period -1. We assume that this increase affects all firms proportionately to the share of the total number of older workers they employ in top jobs. So, in firm n, a marginal increase in $l^{-1}_{o,t}$ increases period-0 older workers in top jobs by $\rho_{t,n} l^{-1}_{o,t,n}/l^{-1}_{o,t}$.

Wages at the bottom change as follows (by the same amount for both genders):

$$\begin{aligned}
\frac{\partial w_{y,b,n}}{\partial l^{-1}_{o,t}} &= \frac{1}{1+\sigma} A_n \theta_{y,b,n} \left(Y_{L_{y,n}L_{o,n}} \theta_{o,t,n} \rho_{t,n} \frac{l^{-1}_{o,t,n}}{l^{-1}_{o,t}} \right. \\
&\quad \left. + Y_{L_{y,n}L_{y,n}} \left(\frac{1}{\sigma} \frac{\left(m_{y,b,n} + f_{y,b,n}\right)}{w_{y,b,n}} \frac{\partial w_{y,b,n}}{\partial l^{-1}_{o,t}} \theta_{y,b,n} + \left(\frac{\partial m_{y,t,n}}{\partial l^{-1}_{o,t}} + \frac{\partial f_{y,t,n}}{\partial l^{-1}_{o,t}} \right) \theta_{y,t,n} \right) \right) \\
&= \frac{\frac{1}{1+\sigma} A_n \theta_{y,b,n} \left(Y_{L_{y,n}L_{o,n}} \theta_{o,t,n} - Y_{L_{y,n}L_{y,n}} \theta_{y,t,n} \right)}{1 - \frac{1}{1+\sigma} \frac{1}{\sigma} A_n \theta^2_{y,b,n} Y_{L_{y,n}L_{y,n}} \frac{\left(m_{y,b,n} + f_{y,b,n}\right)}{w_{y,b,n}}} \rho_{t,n} \frac{l^{-1}_{o,t,n}}{l^{-1}_{o,t}} > 0,
\end{aligned}$$

Wages at the top change as follows (again, by the same amount for both genders):

$$\frac{\partial w_{y,t,n}}{\partial l^{-1}_{o,t}} = \frac{\partial w_{y,b,n}}{\partial l^{-1}_{o,t}} \frac{\theta_{y,t,n}}{\theta_{y,b,n}} - \frac{1}{1+\sigma} \frac{\partial \lambda_n}{\partial l^{-1}_{o,t}}.$$

If the marginal productivities of younger workers in top and bottom jobs are not too dissimilar, an increase in the number of older workers increases wages in the top job less than wages in the bottom job. This is due to the fact that $\frac{\partial \lambda_n}{\partial l^{-1}_{o,t}} > 0$ because the shadow value of relaxing the quantity constraint increases with the number of older workers in the economy.

Next, we show that the restricted access to top jobs leads to larger limitations in the number of younger men who hold these positions. We start by considering the ratio of younger workers of gender g between the top and bottom job in firm n :

$$\frac{g_{y,t,n}}{g_{y,b,n}} = \left(\frac{w^g_{y,t,n}}{w^g_{y,b,n}} \right)^{\frac{1}{\sigma}}.$$

The derivative of this ratio with respect to the number of older workers is as follows:

$$\begin{aligned}
\frac{\partial \frac{g_{y,t,n}}{g_{y,b,n}}}{\partial l^{-1}_{o,t}} &= \frac{1}{\sigma} \left(\frac{w^g_{y,t,n}}{w^g_{y,b,n}} \right)^{\frac{1}{\sigma}-1} \frac{\frac{\partial w^g_{y,t,n}}{\partial l^{-1}_{o,t}} w^g_{y,b,n} - \frac{\partial w^g_{y,b,n}}{\partial l^{-1}_{o,t}} w^g_{y,t,n}}{\left(w^g_{y,b,j}\right)^2} \\
&= \frac{1}{\sigma} \frac{g_{y,t,n}}{g_{y,b,n}} \left(\frac{\partial w^g_{y,t,n}}{\partial l^{-1}_{o,t}} \frac{1}{w^g_{y,t,n}} - \frac{\partial w^g_{y,b,n}}{\partial l^{-1}_{o,t}} \frac{1}{w^g_{y,b,n}} \right).
\end{aligned}$$

When $\frac{\partial \lambda_n}{\partial l_{o,t}^{-1}}$ is between $\frac{\theta_{y,t,f}-\theta_{y,b,f}}{\theta_{y,b,f}}(1+\sigma)\frac{\partial w_{y,b,f}}{\partial l_{o,t}^{-1}}$ and $\frac{\theta_{y,t,f}}{\theta_{y,b,f}}(1+\sigma)\frac{\partial w_{y,b,f}}{\partial l_{o,t}^{-1}}$, the change in the top wage is positive but lower than the change in the bottom wage. Hence, the difference within the parentheses is negative and the employment of younger workers in the top job decreases relative to their employment in the bottom job.

Moreover, it is possible to compare $\partial \frac{m_{y,t,n}}{m_{y,b,n}}/\partial l_{o,t}^{-1}$ to $\partial \frac{f_{y,t,n}}{f_{y,b,n}}/\partial l_{o,t}^{-1}$. The wage change at the top $\frac{\partial w^g_{y,t,n}}{\partial l_{o,t}^{-1}}$ is the same for both genders, but younger men are paid more in top jobs. Therefore, $\frac{\partial w^m_{y,t,n}}{\partial l_{o,t}^{-1}}\frac{1}{w^m_{y,t,n}}$ is smaller than $\frac{\partial w^f_{y,t,n}}{\partial l_{o,t}^{-1}}\frac{1}{w^f_{y,t,n}}$. In contrast, the percentage change at the bottom is the same for both younger men and women. Therefore, the difference in parentheses is more negative for younger men. In addition, it is also true that $\frac{m_{y,t,n}}{m_{y,b,n}} > \frac{f_{y,t,n}}{f_{y,b,n}}$ at baseline because the wages in top jobs are higher for younger men (due to the higher cost of hiring women in top jobs). Therefore, it follows that, if top wages increase in response to a larger supply of older workers (due to positive relative supply effects), then the share of younger men in top jobs decline more than the share of women in the same positions.

In short, workforce aging can shrink the gender pay gap by blocking more younger men from reaching higher-paying positions. And the main reason why younger men are more affected is that they are more represented in top jobs at baseline, and, thus, have more to lose from congested access at the top of firms' hierarchies.

Finally, we consider whether more older workers can change the distribution of younger workers across different types of firms. Under the assumption that the retention rates are higher in higher-productivity firms, an increase in the total number of older workers in the economy leads to larger increases in the number of older workers in the top jobs of higher-productivity firms. The fact that younger workers (both men and women) experience more negative career spillovers in higher-productivity firms does not necessarily mean that they find employment in the bottom jobs of the same set of firms. Bianchi and Paradisi (2023) shows that an increase in the number of older workers induces younger workers to move toward firms with higher percentage increases in the bottom wages (the increase is again a result of complementarity with older workers):

$$\frac{\partial \frac{g_{y,b,n}}{g_{y,b,n'}}}{\partial l_{o,t}^{-1}} = \frac{1}{\sigma}\frac{g_{y,b,n}}{g_{y,b,n'}}\left(\frac{\partial w^g_{y,b,n}}{\partial l_{o,t}^{-1}}\frac{1}{w^g_{y,b,n}} - \frac{\partial w^g_{y,b,n'}}{\partial l_{o,t}^{-1}}\frac{1}{w^g_{y,b,n'}}\right).$$

The derivative is positive if $\frac{\partial w^g_{y,b,n}}{\partial l_{o,t}^{-1}}\frac{1}{w^g_{y,b,n}} > \frac{\partial w^g_{y,b,n'}}{\partial l_{o,t}^{-1}}\frac{1}{w^g_{y,b,n'}}$. Bianchi and Paradisi (2023) also derives under what conditions a larger supply of older workers in top jobs coincides with younger workers moving toward lower-productivity firms, a fact that finds empirical support.

However, to better understand the consequences of this migration toward lower-productivity firms for the gender pay gap, we need to compare the formula above for younger men and women. The two derivatives $\partial \frac{m_{y,b,n}}{m_{y,b,n'}}/\partial l_{o,t}^{-1}$ and $\partial \frac{f_{y,b,n}}{f_{y,b,n'}}/\partial l_{o,t}^{-1}$ are the same because all their components are the same between younger men and women. This result is due to the fact that younger men and women always earn the same wages in the bottom job when they are employed by the same firm. Thus, the number of younger workers who migrate across firms is entirely driven by how many younger workers are blocked from accessing top jobs. Given that (i) younger men are always more negatively impacted than younger women, and (ii) there is more congestion within higher-productivity firms, there are more younger men who move from higher-productivity to lower-productivity firms. In short, workforce aging can contribute to shrinking the gender pay gap by making the distributions of younger men and women across higher- and lower-paying firms more even.

C Correcting for Selection into Employment

As a result of the differential selection of men and women into labor market participation, the literature often considers to what extent are gender gap trends and levels in observed wages reflective of similar patterns in offered wages. Depending on the selection mechanism, features of the distributions of observed and offered wages could be quite distinct from each other. Several methods and applications have been proposed to address this distinction (see Heckman, 1979; Blundell et al., 2007; Mulligan and Rubinstein, 2008; and Blau et al., 2024). We use the most recent imputation method proposed in Blau et al. (2024) to asses whether our empirical findings on the gender gap in observed wages plausibly extend to the gender gap in offered wages.

C.1 Summary of the Imputation Method

The empirical challenge is that offered wages of non-employed people are unobserved. Blau et al. (2024), building on Olivetti and Petrongolo (2008), proposes a method of imputing offered wages for non-employed individuals applying a three-step approach to PSID panel data. Each step augments the baseline sample of full-time employed workers who worked for at least 26 weeks during the previous year:

1. Add employed workers who work part-time and/or worked for less than 26 weeks but more than 100 hours during the previous year. Assign them their actual, observed wage.
2. Add currently non-employed persons who were employed in past or future recent years. Assign them the wage from those recent employment spells.
3. For the remainder persons in ages 25--54 who are still not assigned an offered wage, use a probability-weighted imputation method where observables are used to predict the probability that each person belongs to each one of ten deciles of the gender-year-specific wage distribution.

The third step is iterated until the decile cutoffs used in the estimation of probabilities coincide with the cutoffs in the resulting, post-imputation wage offer distribution. After the ten probabilities are assigned to each person of the last group, observations are tentuplicated and weighted by the predicted probability of belonging to each decile. Then, the midpoint of the wage decile is assigned to each of the ten observations. This method results in unbiased estimates of the median wage gap as long as the probability mass is correctly distributed between above and below the median. In practice, Blau et al. (2024) show that the selection correction results for the median and mean pay gaps are very similar to each other.

Blau et al. (2024) argue that the benefits of this probability-weighted imputation method relative to other selection correction approaches is that (i) it does not assume positive selection into employment, (ii) corrects for the employment selection of men as well as women, and (iii) leads to less scope for bias thanks to its probabilistic nature.

C.2 Implementation

We follow Blau et al. (2024) as closely as possible, with the caveat that our CPS data is cross-sectional and the variables at our disposal to impute probabilities are slightly different than the ones in the PSID. Given that this method is well suited for settings where the data include comprehensive and accurate information on the non-employed, we do not apply this method to the Italian administrative data.

C.2.1 Samples

Throughout the selection correction analysis, our target population are persons between 25 years old and 54 years old, who are either non-employed or employed as workers in the private sector. That is, we do not aim to include the self-employed nor public-sector employees in the distribution of offered wages.[43] Whereas our analyses in the paper consider employed workers between 25 years old and 64 years old, here we drop those between 55 years old and 64 years old to limit high shares of imputed wages in ages close to retirement. As in our main analyses, our baseline compensation measure is weekly earnings, which in the remainder we refer to as wages for short.

Baseline sample As in our main analysis, our core sample consists of employed private-sector workers who worked for at least 24 weeks during the past year.

Step 1 Similarly to step 1 in Blau et al. (2024), we first augment the baseline sample with people who worked for less than 24 weeks but more than 100 hours during the last year. These are people who had lower attachment to the labor market but for whom we directly observe wage information. We denote as Sample 2 the sample that results from the union of these newly added persons and the baseline sample (or Sample 1).

[43] In Section 6.5, we replicate all our main findings using a sample that (i) corrects for time-varying selection into the labor market and (ii) includes public-sector employees.

Step 2 In the absence of panel data, we skip the second step in Blau et al. (2024) and proceed with the probabilistic imputation procedure for those in our target population who are not included in Sample 2. The following imputation algorithm is carried out separately for each gender-year cell.

(a) Using Sample 2, we compute the deciles of the wage distribution. Let the nine wage cutoffs that partition deciles be denoted by w_q, for $q = 1, \ldots, 9$.

(b) Using an ordered probit, we estimate the conditional probability of being in each wage decile as a function of observables X_i. That is, we estimate:

- $Pr(w_i \leq w_1 | X_i)$
- $Pr(w_1 < w_i \leq w_2 | X_i)$
- ...
- $Pr(w_8 < w_i \leq w_9 | X_i)$, and
- $Pr(w_9 < w_i | X_i)$, where X_i includes years of education, dummies for years of education corresponding to the attainment of BAs and advanced degrees, quadratic potential experience, race dummies, hispanic dummy, and region of residence dummies (the nine Census Divisions).

(c) Denote as Sample 3 our final sample that includes offered wages for everyone in our target population. For each person in Sample 3 who is not in Sample 2 (Sample 3 "entrants"), we assign them predicted probabilities of belonging to each wage distribution decile: $\hat{p}_{i,1} \equiv \hat{Pr}(w_i \leq w_1 | X_i)$, $\hat{p}_{i,2} \equiv \hat{Pr}(w_1 < w_i \leq w_2 | X_i)$, ..., $\hat{p}_{i,10} \equiv \hat{Pr}(w_9 < w_i | X_i)$; we then tentuplicate each Sample 3 entrant, and assign each of the replicated observations a weight π_{id} that is equal to $\pi_{id} = \hat{p}_{i,d} \times c_i$, where c_i are the CPS survey weights, for observations/deciles $d = 1, \ldots, 10$. This implies that $\sum_{d=1}^{10} \pi_{id} = c_i$.

(d) To each Sample 3-entrant tentuplicated observation, we assign a wage that is equal to the midpoint of the corresponding wage decile.

(e) Using Sample 3 and weights c_i (for those with observed wages) and π_{id} (for those with imputed wages), we calculate the nine wage cutoffs $\tilde{w}_q$ that split the sample into deciles. If these nine cutoffs are different from the cutoffs w_q in step (a), we iterate the algorithm. Whenever $\tilde{w}_q$ and w_q converge and are equal to each other (up to a small tolerance), the algorithm stops and we are left with our Sample 3, with weights c_i and π_{id}.

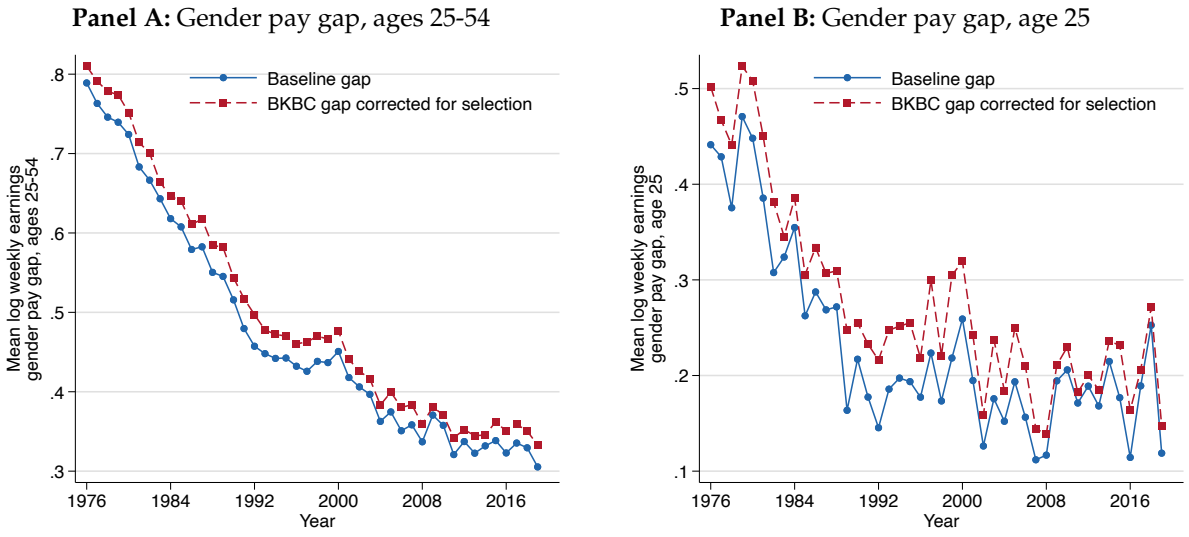
We refer to statistics computed in this augmented sample as "corrected" for selection (labeled BKBC, short for Blau et al. (2024)), and statistics computed among our baseline Sample 1 as "uncorrected" for selection. We now show some results on the corrected and uncorrected wage distributions.

C.2.2 Results

Panel A of Figure C1 shows the evolution of the uncorrected and corrected mean gender gap in weekly earnings for those between the ages of 25 and 54. The main takeaway is that, similarly to Blau et al. (2024), we find that the two series follow a very similar trend, with the corrected gap representing a slight parallel upward shift (between 0.01 and 0.04 log points) relative to the uncorrected gap. We find similar results when zooming into the gender pay gap at age 25. Panel B of Figure C1 shows that the corrected gap at age 25 follows a very similar trend to the uncorrected gap: steep convergence from 1976 until the mid-1990s/early-2000s, and a flat profile thereon. The similarity in the corrected and uncorrected trends in Figure C1 underpins the fact that our main results are very robust to implementing this selection correction method and, as a result, unlikely to be driven by changes over time in the nature of selection into employment.

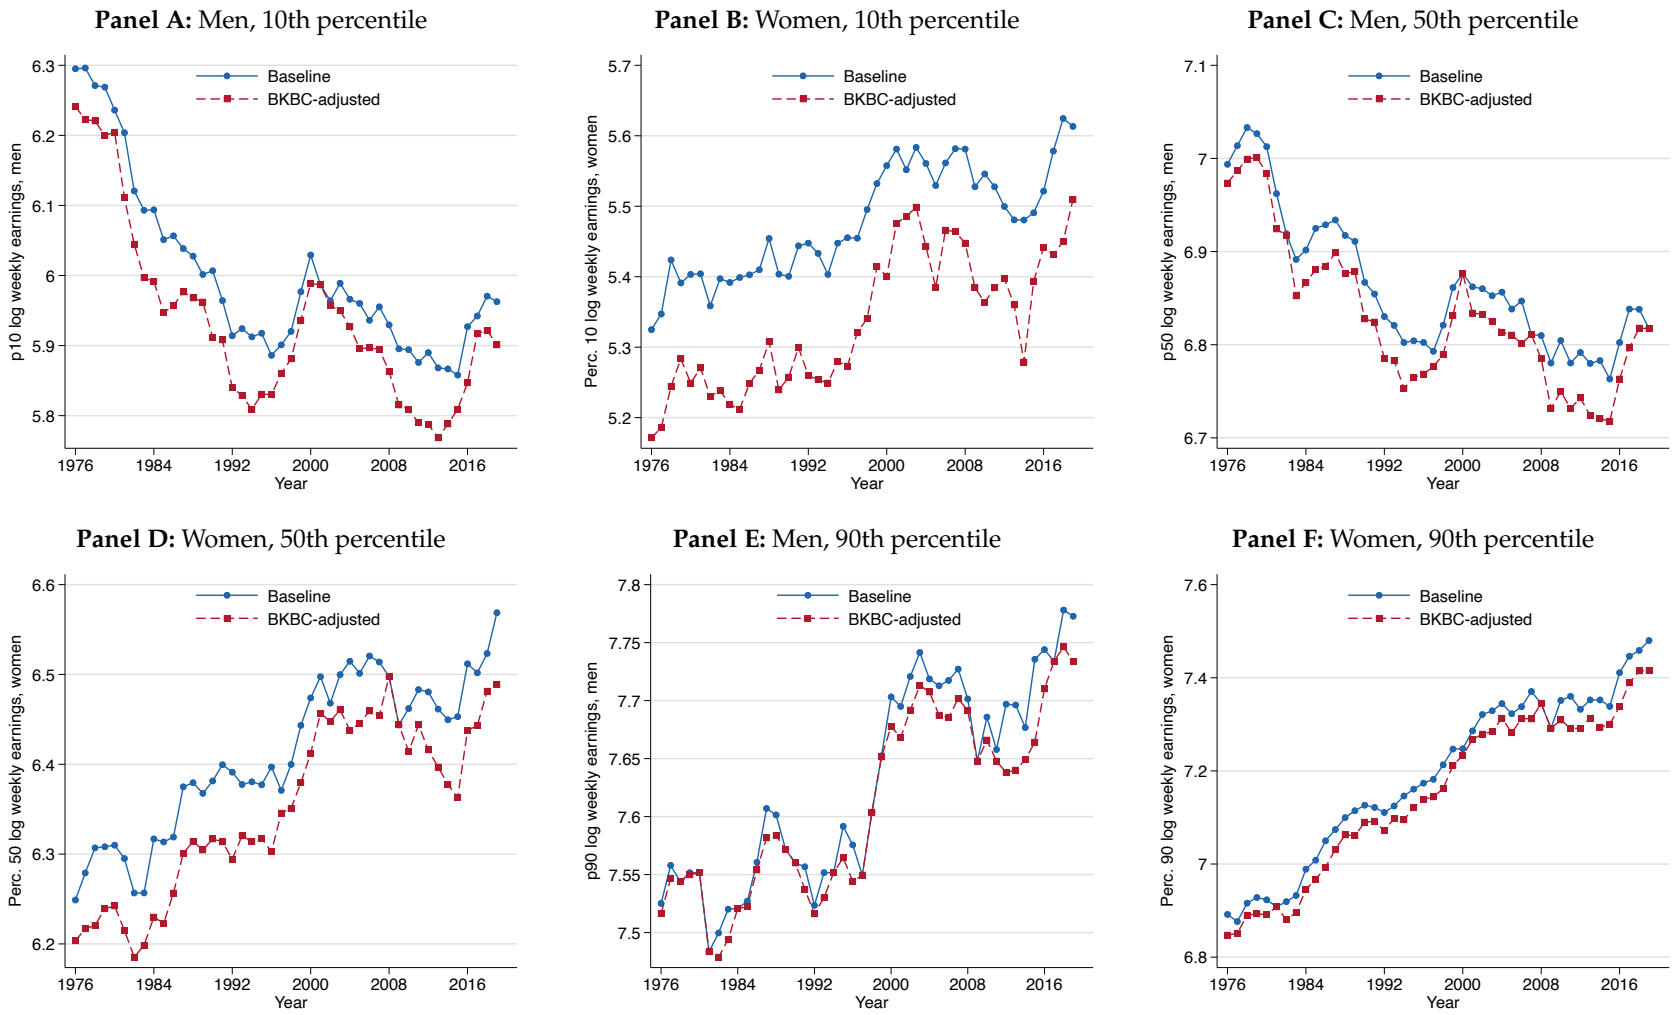
Note, however, that the fact that the relevant gender pay gap trends and our main results are quite similar after accounting for selection does not imply that the correction procedure has little impact on the wage distribution. In fact, we find that it meaningfully impacts the parts of the wage distribution one would expect, both for men and women. Figure C2 shows that the selection correction lowers the 10th percentiles of the male and female distributions by a substantial amount, the 50th percentiles by a moderate amount (especially among women), and only barely impacts the 90th percentiles.

Figure C1: Raw and Selection-Adjusted Gaps



Notes: Trends in the raw and selection-corrected mean gender pay gaps (log weekly earnings of men - log weekly earnings of women) in the United States. Panel A shows the gaps among those between ages 25 and 54. Panel B shows the gaps among people who are 25 years old. The BKBC gap corrected for selection is computed after imputing weekly earnings to nonparticipants using the process outlined in Blau et al. (2024) and described in the text.
Source for the United States. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure C2: Raw and Selection-Adjusted Gaps at Different Percentiles



Notes: Trends in the 10th, 50th, and 90th percentiles of the raw and selection-corrected log weekly earnings distributions, separately for men and women. The BKBC-adjusted percentiles are computed after imputing weekly earnings to nonparticipants using the process outlined in Blau et al. (2024) and described in the text. *Source for the United States*. Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

D Adjusting Weekly Earnings by the Contribution of the Child Penalty

In this section, we describe the procedure to compute the adjustment of weekly earnings for the child penalty in both the United States and Italy. The methodology follows closely Kleven, Landais, and Søgaard (2019), Kleven (2023), and Casarico and Lattanzio (2023). It relies on estimating the effects of childbirth on women and men in the United States and women without children in Italy. Specifically, we are interested in the counterfactual weekly earnings of women, absent the child penalty, in each birth cohort and calendar year:[44]

$$\widetilde{w} = \frac{w}{(1 - p\pi)}, \tag{4}$$

where w are the average weekly earnings of women, p is the cohort- and year-specific child penalty, and π is the average fraction of mothers in each cohort-year pair. We describe in the following paragraphs how we calculate each of these components.

D.1 United States

For each mother i in our CPS sample with year-t weekly earnings w_{it}, our goal is to construct $\tilde{w}_{it}$, the counterfactual weekly wage we would expect if that person were not a mother:

$$\tilde{w}_{it} = \frac{w_{it}}{1 - p_{jc(i,t)}}, \tag{5}$$

where p_{jc} is the child penalty, which is specific to birth cohort c and time-since-childbirth j. We estimate the set of child penalties p_{jc} following Kleven (2023) very closely, with the key difference that we estimate the weekly wage child penalty conditional on working.

Data. Following Kleven (2023), we estimate the child penalties using repeated cross sections of the March CPS 1968–2020 and American Community Survey (ACS) 2000–2019. We follow his sample restrictions and further limit the sample to include workers employed for at least half of the weeks of the preceding year.

Empirical Approach. We follow the pseudo-event-study approach of Kleven (2023) in order to estimate child penalties using repeated cross sections. The initial step in the approach involves matching to generate a pseudo-panel of men and women before and after the birth of their first child (see Kleven (2023) for more details). We then estimate the following earnings equation, separately for men and women:

$$w_{it}^g = \boldsymbol{\alpha}^g \boldsymbol{D}_{\boldsymbol{it}}^{\boldsymbol{Event}} + \boldsymbol{\beta}^g \boldsymbol{D}_{\boldsymbol{it}}^{\boldsymbol{Age}} + \boldsymbol{\gamma}^g \boldsymbol{D}_{\boldsymbol{it}}^{\boldsymbol{Year}} + \nu_{it}^g, \tag{6}$$

where w_{it}^g is the weekly earnings for (pseudo-)individual i of gender $g = w, m$ at event time t. On the right-hand side, boldface denotes vectors. The first term includes dummies for each event time t, omitting a base year before child birth. The event time coefficients α^g measures the impact of child birth on gender g in event year t, relative to the base year. The second and third terms include a full set of age and year dummies.

The estimated level effects are transformed into percentage effects by calculating:

$$P_t^g = \frac{\alpha_t^g}{E\left[\tilde{w}_{it}^g | t\right]}. \tag{7}$$

Here, $\tilde{w}_{it}^g$ is the predicted counterfactual weekly wage in the absence of children.

Finally, the child penalty p is defined as the average effect of having children on women relative to men over a specified event time horizon:

$$p = E\left[P_t^m - P_t^w | t \geq 0\right] - E\left[P_t^m - P_t^w | t < 0\right] \tag{8}$$

Results. Figure D3 presents the pooled event study for men and women surrounding the birth of their first child. The event study horizon displayed in this and ensuing graphs ranges from $t = -5$ years to $t = 10$ years. The average child penalty on weekly earnings conditional on working, across this event time spectrum, is calculated to be 17%. Without conditioning on employment, Kleven (2023) finds an annual earnings child penalty of 33%.

Crucially for our purposes—studying the gender gap dynamics over a long period—we allow the child penalties that feed into counterfactual wages (5) to vary over time. To this end, we estimate child penalties separately across nine birth cohort groups that span the women in our sample: birth cohorts 1930–44, 1945–49, 1950–54, 1955–59, 1960–

[44] To simplify notation, we ignore here time and cohort subscripts.

64, 1965–69, 1970–74, 1975–1979, and 1980–95.[45] Figure D4 shows the estimated child penalties for each of these nine birth cohort groups. The child penalty on weekly earnings conditional on working fluctuates between 11%–23%. Even though the pattern is not monotonic, there is a decreasing trend of the child penalty size across birth cohorts.

The event-time estimates 0–10, one for each of the nine birth cohort groups, represent our adjustment factors p_{jc} in Equation (5). That is, we assign each of the mothers in our CPS analysis sample to one of 99 cells based on her birth year and the age of her eldest son.[46] This cell assignment determines the adjustment factor p_{jc} assigned to each mother and her resulting counterfactual weekly earnings net of the child penalty.

D.2 Italy

Defining the Sample to Estimate Child Penalties. In Italy, we first recover childbirth episodes from the "contribution archive," which reports the full history of workers' Social Security contributions from their first employment spell. This archive not only records actual contributions paid by employers but also imputed contributions related to leaves of absence, sick leaves, unemployment benefit receipt and, crucially, maternity leave. The latter allow us to identify childbirth episodes based on the first month of maternity leave, which has a mandatory duration of five months and can be taken one to two months before the expected childbirth and lasts until three to four months after. The contribution archive is only available for a sample of workers born after 1950. We, therefore, restrict the data to workers included in such sample and extrapolate our estimates of the child penalty to the full set of workers. We further restrict the sample to women who had their first child (their first maternity leave episode) between 25 and 45. As we follow them for at most 10 years after childbirth, our sample comprises women between 25 and 55 years old.

As there is no information on fathers, the first step is to recover a suitable control group of non-mothers. In our sample of 25 to 55-year-old women, we first focus on those born between 1950 (the first birth year available in the contribution archive) and 1974, who were not yet 45 years old by 1976 (the first year in our sample) and turned 45 by 2019. Among these women, we identify mothers from maternity leave take-up both during our observation period (1976-2019) and before, as we have workers' full Social Security contribution histories. Women born between 1950 and 1974 who do not have a child enter the group of never-mothers. Women born after 1974 are subject to right-censoring, as they were not yet 45 by the end of the observation period (2019) and might have had a child after. We solve this issue by assigning a birth probability to the truncated cohort. Specifically, we estimate a linear probability model in the non-truncated cohorts 1950-1974 by regressing a dummy taking value one for never-mothers on the following set of dummy controls: quartiles of the cohort-specific log weekly earnings distribution and province of residence.[47] We then assign all women in the truncated birth cohorts the predicted probability of giving birth based on the coefficients estimated in the linear probability model. We then sort women born after 1974 based on such predicted probability and, starting from the largest value, assign them to the control group up to the point in which the fraction of "predicted" never-mothers in the truncated cohort post-1974 equals the fraction of actual never-mothers in the non-truncated cohorts 1950-1974.[48] The final sample consists of three groups of women: actual mothers, actual never-mothers from non-truncated birth cohorts and predicted never-mothers from truncated birth cohorts. The latter two groups constitute the control group.

The second step is to assign a placebo year of birth to the control group of never-mothers. We do so by assigning a placebo age at childbirth to non-mothers, drawing from the actual distribution of age at childbirth for mothers. We distinguish again between actual and predicted never-mothers. For actual never-mothers, we assume that the distribution of age at childbirth $A_{c,q}$ follows a log-normal distribution within cells of birth cohort c and quartiles of log weekly earnings q, $A_{c,q} \sim \mathcal{LN}\left(\widehat{\mu}_{c,q}, \widehat{\sigma}_{c,q}\right)$, where mean $\widehat{\mu}_{c,q}$ and variance $\widehat{\sigma}_{c,q}$ are obtained from the actual within-cell distribution for mothers. We assign a random draw from this distribution to actual never-mothers. For predicted never-mothers, we use random draws from a distribution with same variance $\widehat{\sigma}_{c,q}$ but different mean $\widetilde{\mu}_{c,q}$, which is obtained by predicting age at childbirth from the estimation of a regression on a quadratic time trend for actual mothers to allow women born after 1974 to have their first child at an older age.[49]

---

[45] The first and last birth cohort groups span a larger number of birth years because these are the groups with fewer observations.

[46] We apply the 10-year child penalty to mothers whose eldest son is older than 10.

[47] In other words, we estimate the following regression $\text{Never-Mother}_{iT} = \alpha + X'_{it}\beta + \epsilon_{it}$, where Never-Mother_{iT} is a dummy equal to 1 for never-mothers in birth cohorts 1950-1974, and X_{it} includes the dummy controls indicated in the text.

[48] The assumption that the fraction of never-mothers is constant contrasts with the secular reduction in fertility rates. However, this assumption is rather innocuous, as it only marginally affects the size of the child penalty, as highlighted in Casarico and Lattanzio (2023).

[49] This adjustment is necessary because of the truncation issue. As we do not observe completed fertility for truncated cohorts, age at childbirth would be skewed to the right if we did not make any adjustment.

Estimating Child Penalties by Year and Birth Cohort. Our goal is to estimate the child penalty in each year and birth cohort. To do so, we run the following event study, separately for mothers and non-mothers:

$$w_{ist}^{g} = \sum_{y} \sum_{j \neq -1} \alpha_{yj}^{g} \mathbf{I}[\mathbf{j} = \mathbf{t}] \mathbf{I}[y = s] + \sum_{k} \beta_{k}^{g} \mathbf{I}[k = \text{age}_{is}] + \sum_{y} \gamma_{y}^{g} \mathbf{I}[y = s] + \varepsilon_{ist}^{g},$$

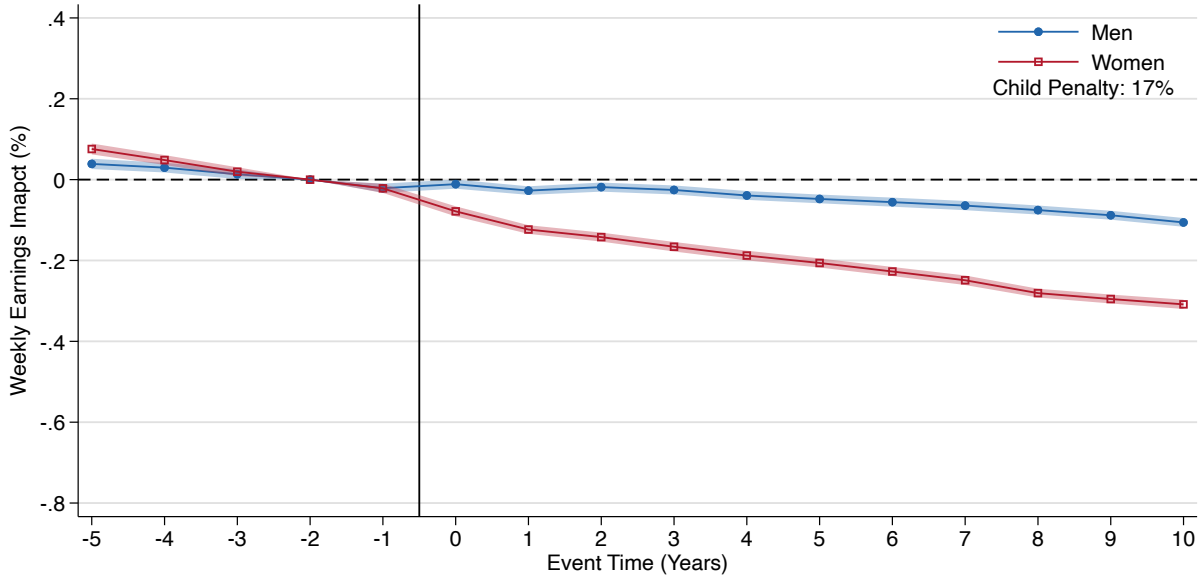
where we interact event time dummies (year to childbirth) with year dummies in order to estimate the year-specific coefficients α_{yj}^{g} (which are equivalent to birth-cohort-specific coefficients, as birth cohort $c = s - t$). The equation also controls for age and year dummies. Figure D5 reports the average child penalty estimates (the average difference in coefficients for mothers and non-mothers) in the first 10 years after childbirth for different birth cohorts. The estimates hover around 6-8%. We assume that the relative stability in the child penalty also applies to the earlier birth cohorts. To this end, we fit a linear trend and assign to the birth cohorts before 1990 the predicted child penalty from later cohorts.[50]

Fraction of Mothers in Each Cohort and Year. The second element needed to correct the weekly earnings of women is an estimate of the fraction of mothers at each event time and for each cohort. For the years in which we have observations, we compute the share of new mothers and the total share of mothers in the data at each age. Figure D6 shows the share of new mothers by year and different age groups in Panel A. As expected, the share of new mothers displays an inverse U-shape relationship with age: it is small at earlier ages, peaks around 30 years old, and then declines. The peak in age at childbirth has increased over time, as can be seen by the upward-sloping trend in the share of new mothers at age 28, especially after the 1990s, mirrored by a decline in the share of new mothers at age 25. For the years in which we do not have enough observations to compute the share of new mothers, we fit a quadratic time trend by age and assume the share of new mothers equals the predicted values of the fit, reported as lines in Figure D6. We also estimate the share of total mothers in a given year at each age. Again, we fit a quadratic trend to retrieve the fraction of mothers by age in the years in which we have no observations.

Correcting Weekly Earnings. We now have all the elements to perform the correction of weekly earnings in Equation 4. Figure D7 reports the life cycle profile, averaged over calendar years, of the gender gap in weekly earnings with and without the adjustment for the child penalty. The unadjusted gender wage gap starts at around 0.12 log points at age 25 and increases to 0.16 log points by age 40. Correcting for the child penalty removes most of the life cycle wage growth in the gap: the gender gap would be 0.13 log points at 40 years old, close to 20% lower than that observed in the data.

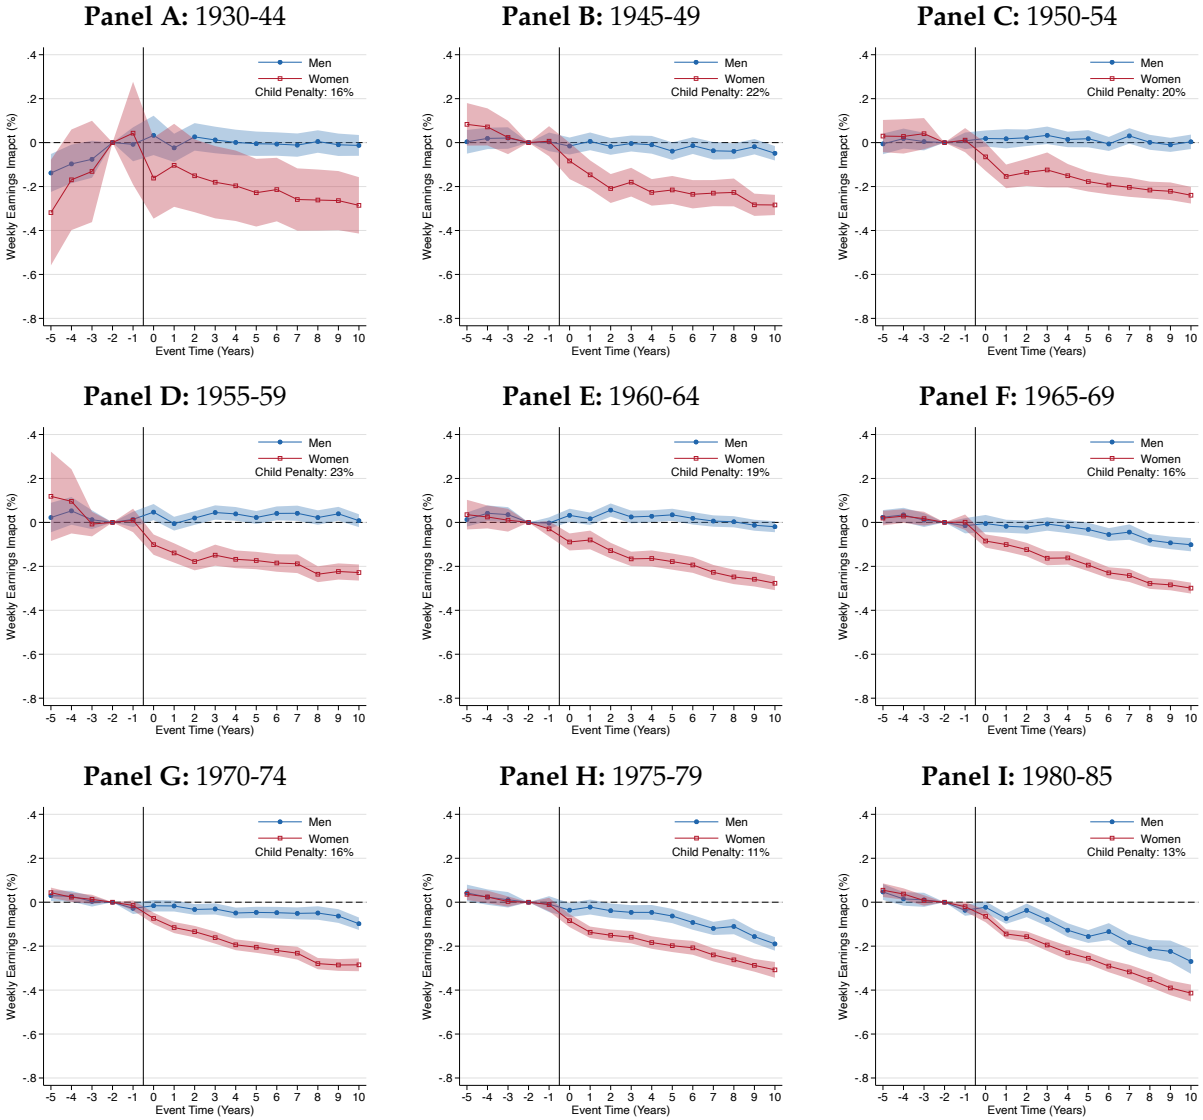
[50] We choose 1990 as a threshold as this is the first birth cohort for which we observe enough observations per mother.

Figure D3: Child Penalty Event Study: Weekly Earnings Conditional on Working



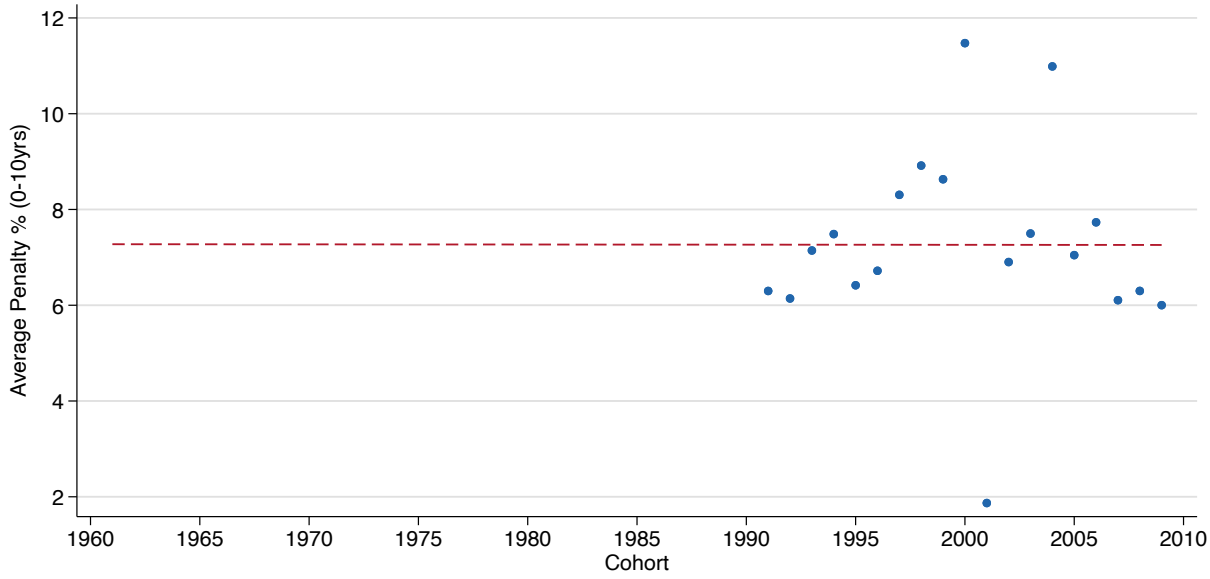
Notes: Event study for weekly earnings for men and women around the birth of their first child at $t = 0$. The series show the percentage impact of child birth on men and women at each event time t, i.e., $\hat{P}_t^m$ and $\hat{P}_t^w$ estimated from Equations (6)-(7). The figure also displays the average child penalty over event times $0 - 10$ defined as in Equation (8). Age at first birth is restricted to be between ages 25–45. The 95% confidence intervals are based on robust standard errors. *Source:* Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure D4: Child Penalty by Birth Cohort: Weekly Earnings Conditional on Working



Notes: Event studies of first child birth for weekly earnings conditional on working across birth cohort groups. The sample of parents is split by birth cohort group and the event study specification (6) is estimated separately for cohort group. Each panel displays the average child penalty over event times 0-10 (defined in Equation 8) for the time period in question. The 95% confidence intervals are based on robust standard errors. *Source:* Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

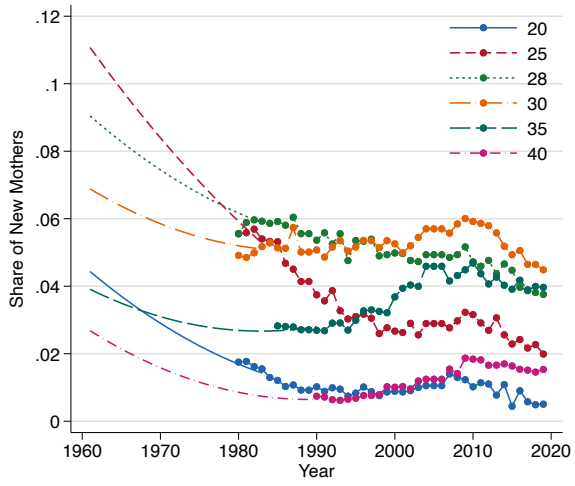
Figure D5: Child Penalty Estimates by Birth Cohort



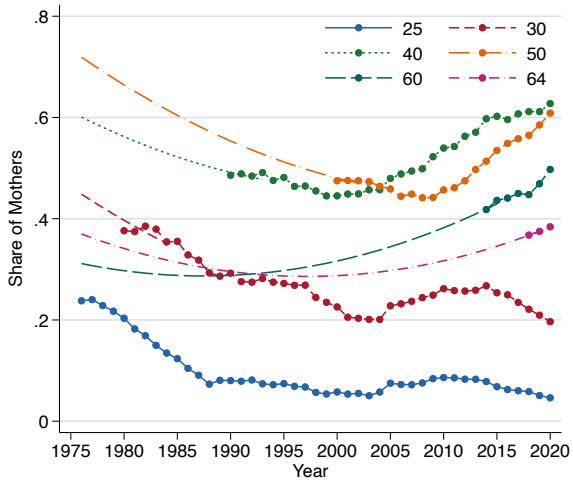
Notes: Average child penalty estimates in the first 10 years following childbirth (dots) and linear trend (dashed lines) by birth cohort. *Source:* Database UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

Figure D6: Share of Mothers by Year and Cohort

Panel A: Share of New Mothers

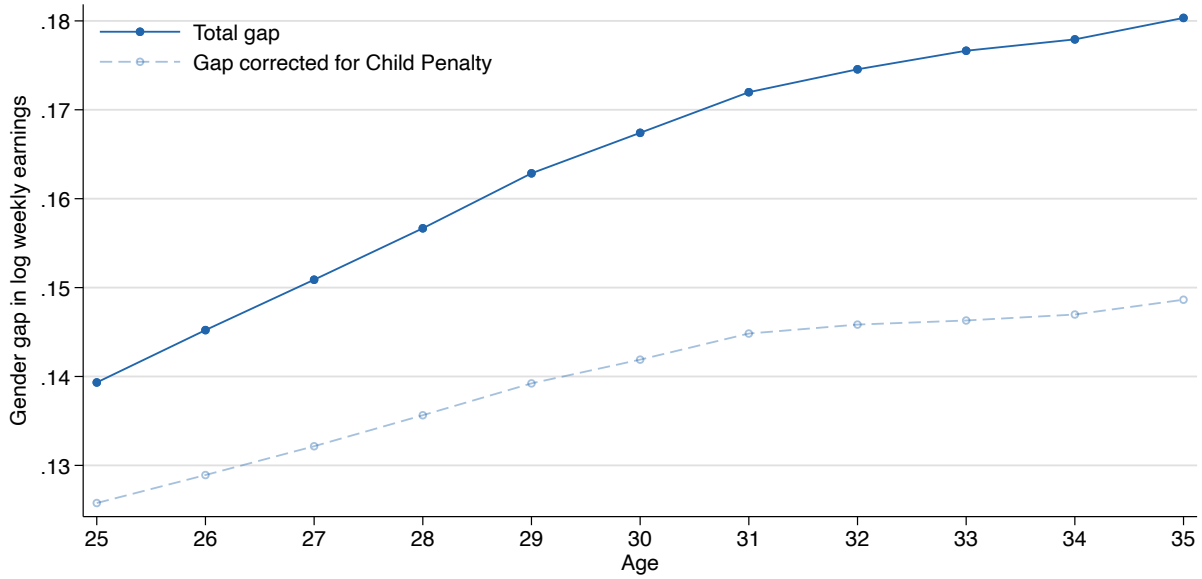


Panel B: Share of Total Mothers



Notes: Panel A shows the observed share of new mothers in any given year (dots), computed as the fraction of women having children over total women. Panel B shows the observed share of total mothers in any given year, computed as the share of women with children over total women. In both panels dashed lines are quadratic trends by age. *Source:* Database UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

Figure D7: Raw and Child-Penalty-Adjusted Weekly Earnings Over the Life Cycle



Notes: The figure reports the gender gap in log weekly earnings over the life cycle averaged over time and cohorts. The solid line is the observed gap. The dashed line is the gap obtained after correcting women's weekly earnings by the estimated child penalty. The sample includes the first ten years after labor-market entry for all cohorts who entered starting from 1995. *Source:* Database UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).