

ONE COHORT AT A TIME: A NEW PERSPECTIVE ON THE DECLINING
GENDER PAY GAP

By

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

October 2026

COWLES FOUNDATION DISCUSSION PAPER NO. 2568



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY

Box 208281

New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

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

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

October 8, 2026

Abstract

This paper studies the interaction between the decrease in the gender pay gap and the stagnation in younger workers' careers, analyzing data from Italy and the United States. The changing mix of cohorts among employed men and women accounts for the entire decline in the aggregate pay gap, holding each cohort's male and female earnings fixed at their initial observed levels. We highlight younger men's positional losses within firms' hierarchies as an important factor driving convergence at labor-market entry. We propose a mechanism by which workforce aging limits career growth for younger workers, disproportionately affecting younger men's trends. Consistent with this explanation, younger men experience the largest positional losses within the firms that are more exposed to workforce aging. We show that these results are robust when accounting for alternative channels, including changes in participation, education, the child penalty, and deindustrialization. Finally, we document that the retirement of older cohorts can account for all of the gender pay gap decline after the mid-1990s.

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, RFBerlin, 87 Trumbull St., New Haven, CT 06511, jaime.arellano-bover@yale.edu; Nicola Bianchi, Kellogg School of Management, Northwestern University, and NBER, 2211 Campus Drive, Evanston, IL 60201, 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 participants and discussants at various conferences and seminars 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), a lower likelihood of being promoted to higher-paying jobs (Bianchi et al., 2023), and lower employment rates in high-skill jobs (Mohnen, 2025). 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 in large part 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. Since then, the retirement of older worker cohorts with higher-than-average gender pay differentials can account for the entire decline in the aggregate gender pay gap, a result that has important consequences for projecting the closure of the pay gap.

We argue that the empirical facts are consistent with a mechanism in which an increased supply of older workers reduces younger workers' promotion opportunities within firms' hierarchies. This aging-driven crowd-out mechanism would operate as follows. Over the past four decades, the relative supply of older workers in high-income economies has increased due to population aging, declining fertility rates, and longer life expectancy (Scott, 2023). Older workers have remained in the labor market longer and have been more likely to advance within firms, benefiting from factors such as backloaded wage schemes (Ke et al., 2018), knowledge spillovers (Sandvik et al., 2020), firm-specific human capital (Lazear, 2009), and employment protection laws (Bentolila and Bertola, 1990). In slot-constrained firms (Lazear et al., 2018) an expansion in the number of senior workers may not be offset by proportional growth in upper-tier positions available for younger workers.

In this environment, workforce aging can create congestion at multiple points along the job ladder and limit younger workers' upward mobility. Crucially, prior to the onset of workforce aging, younger men were substantially more likely than younger women to hold higher-ranked jobs within most career paths. Therefore, aging-driven congestion would tend to displace younger men from these positions at higher rates than younger women, contributing to a narrowing of

gender pay differentials.

Our first empirical result illustrates how the closure of the aggregate gender pay gap over the past four decades has stemmed from reductions happening through labor-market entry and exit of worker cohorts, rather than via changes in the within-cohort, life-cycle profile of the gender gap. This pattern aligns well with our prediction that workforce aging shrinks the gender pay gap among younger workers. Moreover, we show that the 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 the baseline analysis, we assign each worker in our data the mean earnings of workers of the same gender and birth year (a cohort) from the first year in which that cohort appears in our sample, but the results are robust to alternative parametrizations. In both countries, more than the entire decline in the aggregate gender pay gap in weekly earnings can be accounted for by a progressive decrease in this between-cohort component. Our decomposition thus reveals that the trend in the aggregate gender pay gap can be entirely accounted for by the entry of younger cohorts with smaller earnings differentials and the exit of older cohorts with larger differentials.

Next, we show that the cross-cohort component of the gender pay gap has narrowed largely due to worsening labor-market outcomes among younger men. First, we analyze the probability of reaching managerial jobs as a proxy for the likelihood of reaching the upper-tier roles within many career paths inside firms. 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 earnings 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 how the position of 25-year-old men and women evolved 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. The results are robust to alternative choices for the age of labor-market entry and to considering the rank of the gender-specific median entrant rather than the mean of individual entrant ranks.

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 both lower- and higher-paying firms. However, these losses were the largest within the latter group. 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.

After these facts on the evolution of the gender pay gap, 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, for a variety of time windows between 1976 and 2016, we regress the firm-level change in the gender pay gap at ages 25-30 on the change in the firm-level share of workers aged 51-60, controlling for trends correlated with firm age, firm size, sector, and province.

To account for potential endogeneity, we instrument the change in the 51-60 share with the firm-level difference between the share of workers aged 41-50 and the share of workers aged 51-60 at baseline. This instrument leverages cross-firm variation in the exposure to workforce aging stemming from the age distribution of workers over 40 in the initial period. The high tenure of these workers indicates that cross-firm differences in this variable are likely to originate from hiring decisions that firms made decades earlier, rather than firms' recent actions or firm-level shocks. These IVs 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, focusing 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. 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.

Throughout the paper, we also consider a range of alternative mechanisms emphasized in prior work—rising female labor-force participation, sectoral and occupational shifts, trends in education, declining child penalty, and reduced discrimination—to highlight how our analysis complements existing explanations of gender convergence. Our findings are robust to accounting for these trends.

For example, while female labor-force participation rose sharply in the late 1970s and early 1980s, adjusting for this selection ([Blau et al., 2024](#)) leaves our main results unchanged: younger men's positions decline markedly while younger women's remain stable. Similarly, deindustri-

alization plausibly reduced men’s relative earnings, but the convergence we observe takes place almost entirely within sectors, and it holds even when we restrict the sample to those workers predicted to work outside manufacturing. Likewise, differences in educational attainment between men and women explain only a small share of the convergence: the relative decline in younger men’s positions is visible both among college graduates and among those without a college degree. In addition, a decline in the child penalty may have improved women’s outcomes at early career stages, but accounting for this channel only modestly reduces the between-cohort convergence we document.

A vast literature establishes the existence of persistent gender gaps in the labor market (Goldin et al., 2006; Niederle and Vesterlund, 2011; Fortin et al., 2017, Bertrand, 2020; Olivetti et al., 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). Prior work (for example, Black and Spitz-Oener (2010) and Hsieh et al. (2019)) has linked this decline to factors that have improved women’s wages and occupational choices. We complement these findings by showing that part of the narrowing in the gender pay gap at labor-market entry was driven by the relative worsening of younger men’s labor-market outcomes, a trend that we connect to the phenomenon of career spillovers between older and younger workers.

Moreover, prior studies have shown that the gender pay gap has become progressively smaller for younger cohorts while widening within each cohort (Goldin (2014) and Blundell et al. (2025) for the United States; Fortin (2019) for Canada). We build upon this evidence by zooming in on the determinants and consequences of cross-cohort convergence, using data covering a long period and different countries. Our main contribution is to show that labor-market entry and cohort-replacement 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.

A related body of work has demonstrated that economic conditions at the time of labor-market entry have lasting effects on career trajectories (Oyer 2006; Kahn, 2010; Oreopoulos et al., 2012; Arellano-Bover, 2022; Arellano-Bover, 2024). Bovini et al. (2024) and Foliano et al. (2024) have specifically 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 (Guvenen et al., 2022; Dabla-Norris et al., 2023), the macroeconomic consequences of population aging (Jaimovich and Siu, 2009; Jones, 2022; Maestas et al., 2023), and the interconnectedness of the

careers of older and younger workers (Boeri et al., 2022; Bianchi et al., 2023; Bianchi and Paradisi, 2023; Mohnen, 2025), we show both empirically and theoretically that the aging-driven slowdown in the careers of younger workers can disproportionately worsen the labor-market outcomes of younger men, leading to a narrowing of the gender pay gap.

2 A Framework 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 . These two jobs can be thought of as different levels within specific career paths inside firms, rather than strictly denoting top-level executive positions and entry-level jobs. For example, in a manufacturing firm, this framework can capture a blue-collar career path, where an assembly line worker holds a lower-ranked job b and a shift supervisor holds a higher-ranked job t .

Production occurs through the production function $AY(L_y, L_o)$, where A is a productivity shifter, $Y_{L_a} > 0$, and $Y_{L_a, L_o} < 0 \forall a \in \{y, o\}$. 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: $L_a = \theta_{a,t}(m_{a,t} + f_{a,t}) + \theta_{a,b}(m_{a,b} + f_{a,b})$, 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. We introduce three 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, as the firm inherits a stock of older workers from period -1 and cannot change their job assignment and wages.¹ 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

¹ Theoretical and empirical evidence from labor and organizational economics supports this assumption due to factors such as backloaded wage schemes (Lazear, 1979; Ke et al., 2018), firm-specific human capital (Lazear, 2009), knowledge spillovers (Sandvik et al., 2020), and employment protection laws (Bentolila and Bertola, 1990). 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.

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.² Top jobs often entail managerial responsibilities and decision-making authority. Therefore, creating a new top job requires slack in both the firm's organizational capacity (available high-level responsibilities) and its payroll budget. For instance, returning to the earlier example of a blue-collar career path, a firm may not be able to freely create additional shift-supervisor roles—even if it has the payroll budget to do so—unless there are shifts that require an additional supervisory position.

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 ($c_f > c_m$). These parameters, which make younger women less concentrated in top jobs, can be microfounded as either taste-based or statistical discrimination.

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: $w_{a,t}^g = \mu_a w_{a,b}^g$, where $\mu_a > 1$ is the exogenous rent for age group a and $w_{a,j}^g$ is the wage in job j for workers in age group a and gender group $g \in \{m, f\}$.

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 to employ in top and bottom jobs that maximizes its profits, as follows:

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

Appendix [B](#) discusses the full solution of the firm problem.

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

² This idea of constrained slots aligns with models and findings discussed by [Lazear et al. \(2018\)](#), [Bianchi et al. \(2023\)](#), and [Bianchi and Paradisi \(2023\)](#).

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 of younger women to younger men in top jobs below 1: $f_{y,t}/m_{y,t} = c_m/c_f = \delta_f < 1$.

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.³

The comparative static we consider, an increase in the supply of older incumbent workers relative to younger new entrants, matches the evolution of the workforce in both the United States and Italy. Two forces contributed to this change over the period we study: the aging of the baby-boom cohorts and stronger labor-market attachment among incumbent workers. First, as the relatively numerous baby-boom cohorts moved through prime and later-career ages, the age composition of the workforce shifted toward older age groups (Figure A1, Panels A-D). Moreover, these cohorts progressively gained greater representation in higher-ranked positions: older managers, especially those belonging to the baby-boom cohorts, became increasingly numerous relative to managers aged 25-29 (Panels E-F).⁴ Second, employment rates among incumbent workers increased independently of cohort size (Panels G-H).

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

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

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

An increase in the number of older workers induces the firm to add slots at the top 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-

³ 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, or (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 .

⁴ In the United States, these management ratios increased for both older men and older women, but more strongly for older men (Figure A2).

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 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 on younger men and younger women. This framework implies that these negative career spillovers are larger in magnitude for younger men as long as they are more concentrated in top jobs than younger women ($\frac{m_{y,t}}{m_y} > \frac{f_{y,t}}{f_y}$).⁵ 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. 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, as the workforce aged, firm dynamism declined over the period we study (Decker et al., 2014; Hummels and Yue, 2024).

Extensions. In Appendix B, we discuss two extensions to the baseline framework.

The first extension replaces the representative firm with N heterogeneous firms to distinguish within-firm from between-firm variation in the gender pay gap. Higher-paying, more productive firms retain older workers at higher rates, so aging raises the presence of older workers especially in those firms. This generates two results: (i) within every firm, a larger supply of older workers reduces younger men’s probability of holding higher-ranked jobs more than younger women’s, reflecting men’s higher initial concentration in those positions; and (ii) because these spillovers are strongest in higher-paying firms, younger men become disproportionately less likely than younger women to secure employment in these firms and more likely to reallocate toward lower-ranked jobs in lower-paying firms.

The second extension introduces skill heterogeneity to illustrate why the decline in the gender pay gap can slow once younger men and women become similarly represented in top jobs. These skills are acquired before labor-market entry (for example, via education) and map into different tasks in production. In the baseline framework, aging affects the gender pay gap through gender differences in early-career placement across job ranks. When these occupational differences largely vanish, this mechanism cannot generate further gender-differential effects, even if aging continues. In this case, skill composition becomes the main remaining driver of the gap, and further workforce

⁵ 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, given that the relative supply of older workers has increased throughout our sample period, such a model with fixed substitution patterns would not generate the change in trend in younger workers’ outcomes in the mid-1990s/early 2000s that we illustrate in Section 5.

aging can narrow the gender pay gap only if younger women are overrepresented in the higher-skill/higher-return tasks, a condition that is not supported by existing evidence on educational sorting (for example, see [Black et al. \(2008\)](#) and [Sloane et al. \(2021\)](#)).

Summary. In the rest of the paper, we document a set of empirical facts that are consistent with several implications 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).
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 higher-ranked jobs becomes more balanced and predetermined education choices account for a large share of the gender pay gap, workforce aging stops shrinking the gender pay gap at labor-market entry (Section 8).

3 Data

This paper uses administrative and survey data with 375 million observations from the United States and Italy. The combination of Italian and US datasets allows us to exploit the relative strengths of administrative population data (employer-employee linkages) and large-scale survey data (detailed information on all persons regardless of labor force participation).

3.1 US: Labor Force Surveys

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. We impose similar sample restrictions across all the datasets available to us: we limit our sample to individuals between 25 and 64 years old, working in the private sector for at least 24 weeks during the past year, and with strictly positive earnings.⁶

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

⁶ 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.

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 99.9th percentile from above. Our baseline analyses use the wage and salary income variable as provided by IPUMS. In Section 6.6 we show that all our main findings are robust to using historical rank-swapped values for 1976–2010, excluding observations whose earnings are item allocated, and applying both adjustments simultaneously.

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.

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.⁷ This adjustment ensures that the resulting gender gap in weekly earnings is not mechanically driven by women’s greater propensity to work part time or fewer weeks during the year. Both variables are expressed in 2015 euros using the CPI provided by the OECD. Moreover, they are winsorized at the 99.9th percentile from above.⁸ 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.

⁷ INPS reports each worker’s paid weeks and monthly hours paid relative to the full-time hours specified by the applicable collective contract. We construct FTE-adjusted weeks as monthly hours paid/full-time contractual hours $\times$ paid weeks, and FTE weekly earnings as annual earnings divided by FTE-adjusted weeks. We do not separately observe overtime hours, so FTE weekly earnings may still reflect gender differences in overtime intensity.

⁸ Annual earnings are also winsorized at €3,000 from below to address a few instances in which yearly compensation is implausibly low.

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.⁹ 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 (Table A1).

Gap between and within cohorts. We now start to highlight the importance of cross-cohort dynamics in driving the 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 (to avoid time-varying selection into retirement). This graph shows two 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 1971.¹⁰

Second, convergence at labor-market entry across subsequent worker cohorts significantly slowed after 1995. 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 convergence coincided with a slowdown in the closure of the aggregate gender pay gap.

4.2 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: $w_{c,g,t} = w_{c,g}^e + \Delta w_{c,g,t}$, where $w_{c,g}^e$ represents cohort c 's mean log earnings in the

⁹ In the United States, the aggregate gender pay gap was largely stable between 1955 and 1975 (Bailey et al., 2024).

¹⁰ 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.

first year in which it appears in our sample (e for entry into the sample), and $\Delta w_{c,g,t} = w_{c,g,t} - w_{c,g}^e$ is the growth of the average log earnings of cohort c between year t and the first in-sample year. In our baseline results, we compute $w_{c,g}^e$ as the cohort-gender mean earnings either (i) in 1976 for workers who were at least 25 years old in the first sample year or (ii) in the year in which workers who were younger than 25 years old in 1976 turned 25. In Section 6, we show that the results are robust if we choose different entry ages.

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 average log earnings of gender group g is:

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

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 earnings at sample entry 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.

This decomposition between and within cohorts 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—and consequently, between-cohort effects—by the interaction of age (time of labor-market entry) and birth year. This decomposition serves as an initial test of whether the trends in the gender pay gap align with our broader hypothesis, rather than as a formal identification strategy. For this evidence to be a perfect test of our framework, we would need to assume that there are no fundamental, unobserved differences (for example,

men’s ability) across birth cohorts. Our subsequent analyses provide a more direct validation of our hypothesis by establishing a link between workforce aging and worker outcomes at the firm level and by accounting for several cohort-varying factors that have influenced the gender pay gap.

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 as the total decline. These findings hold for different earnings measures (Table A1).¹¹

5 The Wage Ranks of Younger Workers

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 relatively higher-ranked jobs within firms’ hierarchies. Ideally, we would like to have access to detailed personnel data to analyze the probability of younger men and women accessing the higher ranks of specific career paths within firms. Using the administrative and survey data available to us, we focus on managerial jobs as a proxy for the type of higher-ranked positions that represent the culmination of at least some promotion paths. In Section 5.2, we show that our findings do not apply only to these jobs but more broadly to the overall labor market.

We classify higher-ranked positions as all management occupations with associated annual earnings in the top quartile of the year-specific distribution.¹² While this classification is not a necessary condition for our crowd-out mechanism to apply, it is likely a sufficient one.

In both countries, higher-paying managerial jobs have become substantially less likely to be filled by younger workers (Figure A3). 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 these 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 these positions has affected younger

¹¹ The between-cohort component accounts for more than 100 percent of the total decline in the gender pay gap because the widening of the gender gap over the life cycle has increased among more recent cohorts (Bertrand et al., 2010; Goldin and Mitchell, 2017).

¹² 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.

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.

In summary, younger men have progressively held a smaller share of higher-ranked positions, even as the overall prevalence of such roles in the labor market has slightly increased. Younger women were not subject to the same degree of crowd-out because they did not have these promotion opportunities at baseline. If anything, they have become more likely to progress further along firms' hierarchies.

5.2 The Pay Rank Gap at Labor-Market Entry

We next ask whether the pattern documented for higher-ranked managerial jobs extends to younger workers' positions in the overall earnings distribution. Let $F_t(w)$ denote the cumulative distribution of weekly earnings among all workers in our analysis sample in year t , pooling men and women.¹³ For each new entrant i with weekly earnings w_{it}^E , we define the individual's percentile rank as $p_{it} \equiv F_t(w_{it}^E)$. Our baseline statistic is the gender-specific mean of these individual ranks: $\bar{p}_{gt} = \frac{1}{N_{gt}} \sum_{i \in \{g,t\}} p_{it}$, 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 age of labor-market entry in the data at 25.

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 mean rank of men aged 25 was the 50th percentile of the US weekly-earnings distribution and the 47th percentile of the Italian distribution. By 1995, these ranks had fallen to the 39th percentile in the US and to the 36th percentile in Italy. Over the same period, the mean rank of women aged 25 remained close to the 30th percentile in both countries. After the mid-1990s, the ranks of younger men and women followed broadly parallel trajectories. These results are robust to using hourly earnings for the US (Figure A5).

Our approach shares with Bayer and Charles (2018) the idea of evaluating earnings positions relative to a contemporaneous reference distribution, but the baseline estimands differ. Bayer and Charles (2018) converts each prime-age man's annual earnings into its percentile rank in the White male earnings distribution and uses age-adjusted quantile regressions of that rank on race indicators, focusing on the median and the 90th percentile. Our baseline instead considers workers at a common age, ranks their weekly earnings in the pooled distribution of all workers, and reports the

¹³Using workers' positions in the population's pay distribution allows a close mapping between the theoretical framework 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 trends in wage inequality, among other factors.

gender-specific mean of their individual ranks. Because all entrants in our comparison are age 25, no age adjustment is required.

For a quantile-specific measure closer to that in [Bayer and Charles \(2018\)](#), we also consider an alternative metric focusing on the position of the gender-specific *median* entrant. Let $F_t(w)$ be defined as before and $F_{gt}^E(w)$ denote the distribution of weekly earnings among gender- g entrants. The gender-specific quantile- q entry pay is given by w_{qgt}^E , defined so that $F_{gt}^E(w_{qgt}^E) = q$. The rank of the median entrant in the overall distribution is given by P_{qgt} , where $P_{qgt} \equiv F_t(w_{qgt}^E)$ and $q = 0.5$.

Using this measure yields the same qualitative conclusion (Figure [A4](#)): the position of the median male entrant fell substantially relative to that of the median female entrant before the mid-1990s.

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 United States ([Autor et al., 2008](#); [Danieli, 2024](#)) and Italy ([Depalo and Lattanzio, 2025](#)). 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 [A6](#)).

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- and higher-paying firms. For this analysis, we are limited to the Italian context 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 retirement are higher (Figure [A7](#)). 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 those experienced by

younger women 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.¹⁴ Over time, the distribution of younger men across different firm groups moved closer to that of younger women. As expected, there are no remarkable differences across the two genders after 1995 (Figure 4, Panel D).

6 Alternative Explanations and Robustness Checks

6.1 Labor Force Participation

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 A8).¹⁵ Conceptually, changes in female participation could confound the effects we attribute to workforce aging via two channels. First, shifting participation may have altered the average earnings potential of women at labor-market entry: if rising participation draws in more marginal entrants, selection into employment may become more negative, lowering mean female earnings among workers and thus biasing cross-cohort comparisons. However, this channel is not consistent with our finding of a large positional loss for younger men. Second, increases in female participation may have expanded the supply of young workers competing for entry-level positions, potentially compressing equilibrium wages for all

¹⁴Our result that younger women were always overrepresented in lower-paying firms is consistent with the findings in Card et al. (2016) and Casarico and Lattanzio (2024).

¹⁵This 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 Income and Wealth indicate that women's labor force participation was fairly constant in the 1980s and the 1990s and declined only slightly in the 21st century.

young workers and leading to large positional losses for younger men, who started on average from a higher baseline (Goldin et al., 2006). Yet, this explanation conflicts with aggregate supply trends: over the same period, the supply of older workers expanded far more than that of young women (Fry and Braga, 2023), so a pure supply-based mechanism would still predict rising—not falling—wages for younger workers, as captured by the “wage level” component of Equation (1).

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 fewer 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, other) and ethnicity (Hispanic) dummies, and fixed effects for Census divisions. Appendix C provides additional details on the procedure in Blau et al. (2024) and our implementation, while also comparing our approach to other selection-correction methods.

All our prior results hold when we replicate them using this selection-adjusted sample (Figure A9). 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.

6.2 Sectoral and Occupational Shifts

A well-established literature documents how deindustrialization and the shift from manufacturing to services altered employment patterns, especially among younger workers (Autor and Dorn, 2013; Charles et al., 2019). As manufacturing contracted and service industries expanded, women, who historically had lower representation in manufacturing and a comparative advantage in many service and cognitive-skill-intensive jobs, were relatively less exposed to the negative consequences of disappearing industrial jobs. In principle, the cross-cohort decline in the gender pay gap could reflect sectoral reallocation rather than workforce aging.

We first test whether the shrinkage in the gender pay gap is largely accounted for by 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 1-digit NACE Rev. 2 codes in the United States and 2-digit codes in Italy to have enough observations

in each cohort-year-sector cell.

In both countries, most of the between-cohort decline in the gender pay gap has taken place within sectors (Figure A10). 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.

Next, we directly address the progressive decline of the manufacturing sector. Using data from the first five sample years (1976-1980), we estimate a probit model for an indicator for working outside of manufacturing as a function of several worker-level characteristics.¹⁶ We then use the estimated coefficients to (i) predict the probability of working outside of manufacturing for the whole sample, and (ii) retain in the sample only individuals with a predicted probability above the year-specific median (Figure A11). All prior results hold in this more homogenous sample.

The prior evidence on the robustness of our results to sectoral reallocations does not mean that changes in the occupations held by younger men and women did not play a role. Black and Spitz-Oener (2010) shows that technological change reshaped the content of work within occupations in ways that contributed to closing the gender pay gap, with women disproportionately moving away from routine tasks and toward nonroutine analytical and interactive ones. Our analysis examines a different but complementary dimension: how younger workers reallocated across occupations in response to rising congestion from older cohorts (Section 5). Within the framework we study, holding occupations fixed would miss this reallocation.

6.3 Changes in Educational Attainment

Prior work (Autor and Wasserman, 2013; Blau and Kahn, 2017) has emphasized the role of rising female educational attainment—and, more recently, declining male attainment—as a driver of gender pay convergence. We therefore investigate whether diverging trends in education among younger workers can help explain the cross-cohort shrinkage of the gender pay gap.¹⁷ 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 A12). 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

¹⁶We use the following regressors: fixed effects for US Census divisions or Italian regions 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).

¹⁷Here, we focus on the US because the INPS data lack information on completed education for most workers in the sample.

1976 and 2019.

Next, we show that the results in Section 4 and Section 5 hold for individuals with and without a college degree (Figure A13). For example, the finding that younger men have become less likely than younger women to hold higher-ranked positions is robust when analyzing only the jobs held by college graduates or workers without a college education.

While rising college graduation rates among younger women have undoubtedly expanded their career opportunities, part of the gender convergence in pay at labor-market entry stems from a worsening in the careers of younger men *regardless* of educational attainment. We argue that this pattern is consistent with workforce aging being a contributing factor to gender pay equity. Older workers with and without a college education have started occupying a larger share of the upper-tier 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 and supervisory 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 (Table A1, 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 $w_{c,g}^e$ as the earnings observed during the first year in the labor market for each worker who started working after 1976 (Table A1, Columns 7 to 8). Third, the findings hold when we analyze younger workers by years of potential work experience rather than by age (for the US only; Figure A14). Fourth, the managerial-job patterns are similar when workers are observed at ages 30–35 rather than 25–30 (Figure A15).

6.4 Child Penalty

Motherhood carries persistent earnings penalties relative to fatherhood, especially at the time of first childbirth (Kleven et al., 2019; Gallen et al., 2024). Evidence of such penalties already emerged in Figure 2: the gender gap almost always widens over the early career.

Because childbirth typically occurs in early career stages, declining child penalties over time could contribute to cross-cohort convergence in the gender pay gap by improving the relative outcomes of younger women. If this channel were quantitatively important, the between-cohort component of our decomposition might partly capture declining child penalty rather than the effects of

workforce aging. We therefore test whether a decrease in the child penalty borne by mothers in the labor market could fully explain the importance of the between-cohort convergence in earnings.

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 et al. \(2019\)](#) 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. Appendix [D](#) describes in detail how we carry out this exercise in both US and Italian data.

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 small in magnitude (Figure [A16](#)). Even after adjusting for the child penalty in the United States and Italy, the main takeaways of the prior two sections still hold: 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 Changes in Discrimination

[Hsieh et al. \(2019\)](#) documents that reduced barriers to women’s entry into high-skilled jobs account for an important share of aggregate productivity growth in the United States. If gender discrimination in hiring and promotion fell, the gender pay gap may have shrunk even in the absence of workforce aging.

At first glance, our findings may seem inconsistent with this evidence: Figure [3](#) shows that younger women’s relative rank in the earnings distribution at labor-market entry remained largely stable around the 30th percentile across cohorts. However, this stability is best understood as the outcome of offsetting forces. Workforce aging exerted downward pressure on the earnings of all younger workers, and we find strong evidence that this channel disproportionately harmed younger men. Meanwhile, reductions in discrimination likely improved women’s outcomes, preventing their positions from falling in parallel with men’s. The net effect of these countervailing dynamics is a relatively flat trajectory for younger women’s earnings rank.

This interpretation is reinforced by the fact that older women experienced much larger gains over the same period. When we recalculate younger workers’ positions relative to the earnings distribution of men aged over 55, thereby excluding older women from the benchmark group, we observe that younger women in the United States gained some ground between 1976 and 1995

(Figure A17). This pattern is consistent with reduced discrimination improving women’s access to better jobs, even if the mean rank of younger women remains stable once older women’s gains are taken into account.

Taken together, our evidence aligns with the broader literature on gender discrimination: declining discrimination clearly mattered for women’s labor-market advancement, but its effect at the point of labor-market entry was primarily to offset the negative consequences of workforce aging.

6.6 CPS Top-Coding and Allocated Earnings

Two features of the CPS earnings data could affect our US results: the historical treatment of top-coded earnings in the public-use files and the Census Bureau’s allocation of wage-and-salary earnings. Using CPS ASEC records linked to administrative earnings, [Bollinger et al. \(2019\)](#) shows that earnings are not missing at random and that nonresponse is concentrated in both tails of the distribution; it also recommends using historical rank-swapped values for pre-2011 public-use ASEC data. Motivated by these concerns, we replicate our main analyses under three alternative specifications: (i) using Census rank-proximity-swapped earnings for 1976–2010 (the public-use files already incorporate rank swapping from 2011 onward), (ii) excluding records whose wage-and-salary earnings are item allocated, and (iii) applying both adjustments simultaneously. Across all three specifications, the raw gender pay-gap trend, the total and between-cohort changes in the gender pay gap, the gender-specific trends in the share of higher-ranked managerial positions held by workers aged 25–30, and the average earnings ranks of men and women at age 25 closely track their baseline counterparts (Figure A18). Thus, neither the treatment of top-coded earnings nor the inclusion of item-allocated earnings drives our main US findings.

6.7 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 A19). 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 A20, 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 A20, 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 A21).

7 Firm-Level Exposure to Workforce Aging

This section tests the core idea behind the stylized framework in Section 2 by drawing 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:

$$\Delta y_f^{(25-30)} = \alpha + \beta \Delta s_f^{(51-60)} + X_f' \gamma + \varepsilon_f, \quad (3)$$

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 sector, all observed at baseline. Due to its first-difference setup, this specification implicitly controls for any time-invariant firm-level factors, while γ controls for differential trends as a function of covariates included in X_f' . Standard errors are clustered at the province level.¹⁸

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.

¹⁸When 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 $\Delta y_{g,f}^{(25-30)} = \alpha + \beta \Delta s_f^{(51-60)} + \gamma \Delta s_f^{(51-60)} \cdot m_g + \delta m_g + X_f' \theta + \varepsilon_{g,f}$, where the subscript g denotes gender, and m_g is a dummy variable equal to 1 for men's labor-market 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 that is 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 $\tilde{\Delta s}_f^{(51-60)}$, our instrument is defined as:

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

where $s_{f,1976}^{(41-50)}$ is the share of firm f 's workers between 41 and 50 years old in 1976.¹⁹ 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 (2025) to study the effects of differences in retirement rates across US commuting zones on youth employment.

IV Relevance. The variable $\tilde{\Delta s}_f^{(51-60)}$ 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 A2, Panel A). A 1-percentage-point increase in $\tilde{\Delta s}_f^{(51-60)}$ 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 first stage is strong, with F-statistics that range between 2,000 and 11,000 depending on the sample and specification (Table 1).

IV Exclusion. Second, the exogeneity of the instrument relies on the assumption that the projected change in workers over 50 in 1976 ($\tilde{\Delta s}_f^{(51-60)}$) 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 directly testable, we argue that, unlike $\Delta s_f^{(51-60)}$, differences in $\tilde{\Delta s}_f^{(51-60)}$ 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 work-

¹⁹The 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 $x + 2$. In spite of these restrictions, our sample still captures a large share of the total workforce (46 percent of all workers in the Italian data).

ers between 1976 and 1986.²⁰ To support this claim, we show that 59 percent of workers between 41 and 50 years old in 1986 were hired at their employer ten or more years earlier (Figure A22).²¹

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 regressions (Table 1, Panel A). Here, the actual change in the share of older workers is the 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 3.1-percentile additional reduction in the gender gap in pay rank among younger coworkers at the firm. This effect is 77 percent larger than the average decline 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 4.7-percentile faster closure in the gender pay rank gap.²²

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 3.8-percentile positional loss for younger men and with a 1.1-percentile positional gain for younger women. The former effect is statistically significant at the 1-percent level, while the latter is insignificant at the 10 percent level.

²⁰For 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 $\tilde{\Delta s}_f^{(51-60)}$ is predetermined in 1976, it is unlikely to correlate with such shocks once X_f is controlled for.

²¹We 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).

²²The 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 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 $\tilde{\Delta s}_f^{(51-60)} = \left(s_{f,1976}^{(31-40)} - s_{f,1976}^{(51-60)} \right)$, where $s_{f,1976}^{(31-40)}$, the share of firm f 's workers aged 31-40 in 1976, replaces $s_{f,1976}^{(41-50)}$.²³ 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 25,279 in the baseline analysis to 13,993 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 4.9 percentiles in the OLS estimates and 6.2 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 replicate this analysis over the subsequent twenty-year period, 1996-2016, when convergence in the gender pay gap at labor-market entry stalled (Table 1, Panel C).²⁴ Based on our earlier findings, we expect 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 0.2 percentiles in the OLS estimates and 0.9 percentiles in the 2SLS estimates, two effects that are statistically insignificant.

Rather than highlighting a few key periods, we estimate Equation (3) for all ten-year periods ending between 1986 and 2016 (Figure A23). This analysis using rolling ten-year windows confirms our prior findings. The instrument remains a strong predictor of the endogenous variable. Moreover, more severe firm-level workforce aging is correlated with a larger decline in the gender pay gap (Panels A and B) and more negative outcomes for younger men until the mid-1990s (Panels C and D), while the association wanes afterwards.

Our conclusions also do not depend on the specific age band used to define senior workers. Reestimating the specifications in Table 1 using two broader definitions—ages 45-60 and ages 50 and above—yields qualitatively similar results: an increase in the senior-worker share is associated with a decline in the gender pay-rank gap among younger workers during the early part of the sample, with a similarly weaker and less precise relationship later in the sample (Table A3).

²³ Similar to workers aged 41-50, these slightly younger workers had a high mean tenure in 1986 (Figure A22, Panel C). Specifically, 51 percent of them were hired at their current firm eight or more years earlier.

²⁴ In this case, the instrument becomes $\tilde{\Delta s}_f^{(51-60)} = \left(s_{f,1996}^{(31-40)} - s_{f,1996}^{(51-60)} \right)$.

Because the baseline instrument and the endogenous change in the firm’s 51-60 share both depend on the firm’s baseline 51-60 share, we also construct a sector-age shift-share instrument that combines each firm’s initial single-year age composition with age-specific retention rates computed among other firms in the same sector. We also control for the firm’s initial employment shares outside the senior-worker range.²⁵ Table A4 compares the resulting estimates with the baseline 2SLS estimates: the shift-share estimates carry the same sign and a comparable, if somewhat larger, magnitude in both 1976-1986 and 1976-1996, with a first stage that remains strong but less extreme than under the baseline design, and both approaches yield small and statistically insignificant effects after 1996.

Finally, because the firm-level IV analysis uses a selected sample of larger, surviving firms, we also compare trends in the main outcomes in this sample with those in the full population of private-sector employees. Using, in each year, the same firms entering the corresponding rolling ten-year regressions in Figure A23, we replicate the entry-earnings cohort decomposition, the share of managerial jobs held by young workers by gender, and young workers’ mean earnings rank by gender, for both the rolling-regression sample and the full population (Figure A24). Over 1976–2007, the entry-earnings counterfactuals and the deterioration in younger men’s relative position follow similar patterns in both samples, although the observed aggregate gender gap declines less in the rolling-regression firms. The level differences are consistent with the restriction to firms with more than five employees that survive over the ten-year regression window.

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 stagnation 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).

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

²⁵Formally, let ρ_{j,a,t_0}^{-f} denote the retention rate—at the same employer h years later—of workers aged a in sector j at baseline year t_0 , excluding firm f ’s own workers. The instrument is $z_{f,t_0} = \frac{1}{N_{f,t_0}} \sum_{a=41}^{50} n_{f,a,t_0} \rho_{j(f),a,t_0}^{-f} - s_{f,t_0}^{(51-60)}$. As in the underlying firm panel, we average z_{f,t_0} over three consecutive baseline years to smooth firm-level measurement error. For the 1976-1996 window, $h = 20$ and the summation instead runs over ages 31-40 at baseline, corresponding to ages 51-60 after twenty years.

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 labor market entry. Specifically, considering a baseline year t_b , our new counterfactual 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 A25, 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 A25, 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.

8.2 Consequences for Future Convergence

The ongoing importance of convergence through exit carries relevant implications for the future of the gender pay gap. Since retirees with higher-than-average gaps are currently the only source of convergence, the gender pay gap is *not* projected to close.

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 different conclusions. Since the aggregate gender pay gap has been decreasing, any projection would predict full convergence in a few decades in most high-income countries. For example, in its 2023 Global Gender Gap Report, the World Economic Forum predicts that Europe will reach gender parity in 67 years, while North America 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 by 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:

$$g_s^e = \alpha_t - \beta_t (s - \underline{s}_t), \quad (4)$$

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 $s^* = \frac{\alpha_t}{\beta_t} + \underline{s}_t$.²⁶ In addition to estimating Equation (4), 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 (4) in the US and Italy for two years: the year 2000 and the last sample year (Figure 6, Panels A and D). Both countries' convergence rate β_t dramatically decreased from 2000 to 2019. In addition to highlighting two different years, we show the evolution of β_t across all years in the sample (Figure 6, 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 6, 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 to a year after 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.

8.3 Why Did Convergence at Labor-Market Entry Stop?

This section examines why the convergence at entry stopped after the mid-1990s. Our 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

²⁶If 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.

between younger men and women in the share of relatively higher-ranked jobs (Section 5.1) and in mean pay rank (Section 5.2) were small by 1995 and remained 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 earnings 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; Huneus et al., 2021; Sloane et al., 2021; Bovini et al., 2024; Humphries et al., 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 framework'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 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.²⁷

In both countries, the gender pay gap at entry predicted by younger workers' major choices slightly declined until the mid-1980s (Figure A27).²⁸ After this period, it has remained 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.²⁹ These two large shares are likely a lower bound on 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).

²⁷ 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 A26). 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.

²⁸ 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).

²⁹ Our results based on Italian survey data are slightly smaller than estimates based on Italian administrative data (60 percent from Bovini et al. (2024)).

9 Conclusion

In this paper, we document a set of empirical facts about the evolution of the gender pay gap and younger workers' outcomes in Italy and the United States over the last four decades, and propose a mechanism that accounts for them. The entire decline in the aggregate gender pay gap originates from cohort turnover: newer cohorts enter the labor market with smaller-than-average gaps, while older cohorts who exit have larger-than-average gaps. Much of the cross-cohort convergence at labor-market entry up to the mid-1990s reflects younger men falling closer to younger women in both firms' hierarchies and the market-wide pay distribution; after the mid-1990s, convergence at entry largely stops. Since then, the ongoing reduction in the aggregate gap stems from the exit of older cohorts with large gender pay disparities, implying that—in the absence of structural breaks—labor markets in our sample are unlikely to reach full gender pay convergence.

We interpret these descriptive patterns through the lens of career spillovers between older and younger workers. In particular, the evidence is consistent with the idea that shifts in the age composition of the workforce, combined with constraints on expanding higher-ranked positions, can limit early-career advancement, with larger consequences for younger men's trends given their higher initial concentration in upper-tier roles. The evidence in favor of this mechanism is robust to—and in our view, complementary to—other channels emphasized in the literature, including changes in participation, sectoral and occupational shifts, education, the child penalty, and reduced discrimination.

Overall, the cohort-based perspective highlighted in this paper points to three broad implications: within-cohort gender disparities remain persistent; the early years in the labor market appear to be a critical period shaping each cohort's life-cycle pay differentials; and the gender and age dimensions of labor market inequality are closely intertwined.

References

- Acemoglu, Daron and Pascual Restrepo.** 2023. "Automation and Rent Dissipation: Implications for Inequality, Productivity, and Welfare." working paper, MIT.
- 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.
- Arellano, Manuel and Stéphane Bonhomme.** 2017. "Quantile selection models with an application to understanding changes in wage inequality." *Econometrica*, 85(1): 1–28.
- Autor, David and Melanie Wasserman.** 2013. "Wayward Sons. The Emerging Gender Gap in Labor Markets and Education."
- Autor, David H and David Dorn.** 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review*, 103(5): 1553–1597.
- 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.

- 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.
- 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, 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.** 2025. “Labor market inequality and the changing life cycle profile of male and female wages.” *American Economic Journal: Applied Economics*, 17(4): 100–133.
- 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.
- Bollinger, Christopher R., Barry T. Hirsch, Charles M. Hokayem, and James P. Ziliak.** 2019. “Trouble in the Tails? What We Know about Earnings Nonresponse 30 Years after Lillard, Smith, and Welch.” *Journal of Political Economy*, 127(5): 2143–2185.
- 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 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.
- 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.** 2025. “The Increase in Earnings Inequality in Italy: The Role and Persistence of Atypical Contracts.” *Review of Income and Wealth*, 71: e12709.
- Fernández-Val, Iván, Aico van Vuuren, Francis Vella, and Franco Peracchi.** 2023. “Selection and the distribution of female real hourly wages in the United States.” *Quantitative Economics*, 14(2): 571–607.
- 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.
- Fry, Richard and Dana Braga.** 2023. “Older Workers Are Growing in Number and Earning Higher Wages.” *Pew Research Center*.
- Gallen, Yana, Juanna Schrøter Joensen, Eva Rye Johansen, and Gregory F. Veramendi.** 2024. “The Labor Market Returns to Delaying Pregnancy.” *Working Paper, Harris School of Public Policy*.
- 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.
- 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.
- Huneus, 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.
- Jaimovich, Nir and Henry E. Siu.** 2009. “The Young, the Old, and the Restless: Demographics and Business Cycle Volatility.” *American Economic Review*, 99(3): 804–826.
- Jones, Charles I.** 2022. “The End of Economic Growth? Unintended Consequences of a Declining Population.” *American Economic Review*, 112(11): 3489–3527.
- 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 Sogaard.** 2019. “Children and Gender Inequality: Evidence from Denmark.” *American Economic Journal: Applied Economics*, 11(4): 181–209.
- 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.
- Maasoumi, Esfandiar and Le Wang.** 2019. “The gender gap between earnings distributions.” *Journal of Political Economy*, 127(5): 2438–2504.
- Maestas, Nicole, Kathleen J. Mullen, and David Powell.** 2023. “The Effect of Population Aging on Economic Growth, the Labor Force, and Productivity.” *American Economic Journal: Macroeconomics*, 15(2): 306–32.
- Mohnen, Paul.** 2025. “The Impact of the Retirement Slowdown on the U.S. Youth Labor Market.” *Journal of Labor Economics*, 43(1): 203–246.
- 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.
- 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.
- Oyer, Paul.** 2006. “Initial Labor Market Conditions and Long-Term Outcomes for Economists.” *Journal of Economic Perspectives*, 20(3): 143–160.
- Rosolia, Alfonso and Roberto Torrini.** 2007. “The generation gap: relative earnings of young and old workers in Italy.” Bank of Italy, Temi di discussione 639.
- 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.
- Scott, Andrew J.** 2023. “Working Life—Labor Supply, Ageing, and Longevity.” In *The Routledge Handbook of the Economics of Ageing*. 329–348. London: UK:Routledge.
- 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.
- 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 depict 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 trace 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 each cohort (defined as a combination of birth year and gender) its mean log weekly earnings in the first year it appears in the sample (Equation (2)). In the baseline analysis, entry into the sample corresponds to the year in which workers in each cohort turn 25 years old (for workers younger than 25 in 1976, the first sample year). We assign cohorts who were older than 25 at the start of the sample their mean weekly earnings in the first sample year.

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 earnings 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




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


Panel A: Gender pay gap, 1986-1976




Panel B: Gender pay gap, 1996-1976




Panel C: Gender pay gap, 2016-1996




Panel D: Rank in pay distribution, 1986-1976



Panel E: Rank in pay distribution, 1996-1976




Panel F: Rank in pay distribution, 2016-1996



Notes: These figures show binned scatterplots that correlate firm-level changes in the gender pay gap of younger workers with 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 for 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 for 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 from x to $x+2$. In each panel, the sample of firms includes 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: Projected Convergence in the Gender Pay Gap



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 (4)). 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 a gap at most equal to zero at age 25) predicted by Equation (4) 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.

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

	Gap in pay rank between men and women at ages 25-30					
	OLS		Reduced form		2SLS	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 1976-1986						
Δ share older workers	-2.78** (1.06)	-3.08*** (1.08)	-2.83*** (0.97)	-2.73*** (0.98)	-4.88*** (1.67)	-4.65*** (1.68)
KP F-stat					7,168	11,259
Mean dep. var.	-1.75	-1.75	-1.75	-1.75	-1.75	-1.75
SD dep. var.	19.94	19.94	19.94	19.94	19.94	19.94
Obs.	25,279	25,279	25,279	25,279	25,279	25,279
Panel B: 1976-1996						
Δ share older workers	-5.02*** (1.77)	-4.88*** (1.81)	-3.50*** (1.27)	-2.86** (1.19)	-7.74*** (2.78)	-6.24** (2.57)
KP F-stat					4,389	4,734
Mean dep. var.	-3.89	-3.89	-3.89	-3.89	-3.89	-3.89
SD dep. var.	20.41	20.41	20.41	20.41	20.41	20.41
Obs.	13,993	13,993	13,993	13,993	13,993	13,993
Panel C: 1996-2016						
Δ share older workers	-0.42 (1.43)	-0.24 (1.42)	-0.43 (1.23)	-0.37 (1.28)	-1.11 (3.18)	-0.93 (3.19)
KP F-stat					2,286	4,364
Mean dep. var.	-0.76	-0.76	-0.76	-0.76	-0.76	-0.76
SD dep. var.	20.20	20.20	20.20	20.20	20.20	20.20
Obs.	14,644	14,644	14,644	14,644	14,644	14,644
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 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, sector, and province fixed effects, all observed at baseline. Firm-level values for year x are computed as three-year averages from x to $x + 2$. In each panel, the sample of firms includes 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. Standard errors are clustered at the province level. Source: UNIMENS, 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)
<u>Panel A: 1976-1986</u>								
Δ share older workers	-2.67*** (0.80)	-2.20*** (0.72)	-4.60*** (1.37)	-3.77*** (1.24)	0.16 (0.85)	0.63 (0.81)	0.28 (1.46)	1.11 (1.39)
KP F-stat			3,584	3,584			3,584	3,584
Mean dep. var.	-2.19	-2.19	-2.19	-2.19	-2.19	-2.19	-2.19	-2.19
SD dep. var.	16.50	16.50	16.50	16.50	16.50	16.50	16.50	16.50
Obs.	50,558	50,558	50,558	50,558	50,558	50,558	50,558	50,558
<u>Panel B: 1976-1996</u>								
Δ share older workers	-6.07*** (1.10)	-5.74*** (0.85)	-13.42*** (2.53)	-12.58*** (1.86)	-2.57* (1.49)	-2.24* (1.23)	-5.68* (3.33)	-4.84* (2.72)
KP F-stat			2,194	2,194			2,194	2,194
Mean dep. var.	-3.65	-3.65	-3.65	-3.65	-3.65	-3.65	-3.65	-3.65
SD dep. var.	18.25	18.25	18.25	18.25	18.25	18.25	18.25	18.25
Obs.	27,986	27,986	27,986	27,986	27,986	27,986	27,986	27,986
<u>Panel C: 1996-2016</u>								
Δ share older workers	-1.77 (1.07)	-1.09 (0.89)	-4.57 (2.83)	-2.75 (2.27)	-1.34 (1.18)	-0.67 (1.00)	-3.47 (3.09)	-1.64 (2.54)
KP F-stat			1,143	1,143			1,143	1,143
Mean dep. var.	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03
SD dep. var.	17.92	17.92	17.92	17.92	17.92	17.92	17.92	17.92
Obs.	29,288	29,288	29,288	29,288	29,288	29,288	29,288	29,288
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 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, sector, and province fixed effects, all observed at baseline. Firm-level values for year x are computed as three-year averages from x to $x + 2$. In each panel, the sample of firms includes 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. Standard errors are clustered at the province level. *Source:* UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

Online Appendix


A Additional Figures and Tables

Figure A1: Workforce Aging




Notes: Panels A-B report survey-weighted age shares among private-sector employees aged 16-80; Panels C-D report each five-year group's share relative to ages 25-29. Panels E-F report analogous ratios for higher-ranked managerial jobs, using the country-specific definitions in Figure 3, Panels A-B. Panel F has a series break in 1996 when *quadri* positions enter the classification. In Panels C-F, approximate cohort is year minus the midpoint of the numerator group. Panels G-H report private-sector employment rates; 50+ covers ages 50-80 and the denominator is private-sector employees plus the unemployed. Panels A-E require nonmissing wages and at least 24 weeks worked in the prior year (CPS) or six months (SHIW). U.S. ages above 80 are coded as 80. *Sources for Italy*: Survey of Household Income and Wealth, Bank of Italy. UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS) (Panel F). *Source for the United States*: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A2: Management Ratios by Gender




Notes: Panels A and B report, respectively for men and women, each five-year age group's share of U.S. higher-ranked management jobs relative to the corresponding share for ages 25-29. Approximate cohort is year minus the midpoint of the numerator group. Managers are defined by two-digit SOC code 11. 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. The sample consists of private-sector employees aged 16-80 with nonmissing wages and at least 24 weeks worked in the prior year. Ages above 80 are coded as 80. Source: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A3: Share of Higher-Ranked Managerial Positions in Economy




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 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 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: these are apprenticeships, blue-collar jobs, white-collar jobs, and managerial jobs. The 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 A4: Percentile Rank of the Median Entrant.



Notes: Panels A and B show the percentile rank of the median male and female entrant at age 25 in the United States and Italy, 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, 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 A5: 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, and had earned strictly positive earnings.

Source: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.


Figure A6: Younger Workers' Distance from Mean Earnings



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 A7: Share of Workers Over 50 in Lower-Paying and Higher-Paying Firms



Notes: The 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 A8: 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 A9: Results Adjusted for Selection into Labor Market, Weekly Earnings

Panel A: Between-cohort change




Panel B: 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 fewer 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 earnings distributions at the start and end of the imputation process become sufficiently similar.

Source: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.


Figure A10: Between-Cohort Decline in Gender Pay Gap, Between and Within Sectors



Notes: Panels A and B decompose both the total change in the gender pay gap and just its between-cohort component between and within sectors (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 A11: Workers With High Probability of Working Outside Manufacturing



Notes: These figures replicate the main findings after keeping in the sample only workers with a high predicted probability of working outside manufacturing. We predict the probability of working outside of manufacturing using the first five sample years (1976-1980) and the following observable characteristics: Census divisions (US) or region (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 a predicted probability above the year-by-year median. *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 A12: Between-Cohort Decline in Gender Pay Gap, Between and Within College Education



Notes: The figure decomposes both the total change in the gender pay gap and just its between-cohort component between and within groups defined by college graduation status in the United States.

Source: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.


Figure A13: 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. In Panels B and E, higher-ranked managerial jobs are managerial positions with annual earnings above the year- and education-specific 75th percentile, calculated among private-sector employees aged 16 and older who worked at least 24 weeks. Each series reports the share of these jobs held by men or women aged 25–30; the denominator comprises all such jobs held by workers aged 25–64 in the same education group, pooling both genders. Panels C and F report the mean percentile ranks of men and women at age 25 in the contemporaneous weekly-earnings distribution of workers aged 25–64 within the same education group.

Source: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.


Figure A14: Rank in Pay Distribution Based on Potential Experience



Notes: Panels A to D show the average earnings percentile of men and women with 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, and had earned strictly positive earnings.


Source: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A15: Managerial Positions at Ages 30–35



Notes: These figures reproduce the analysis on the share of higher-ranked management jobs held by younger people, focusing on ages 30–35 instead of the baseline age range 25–30. *Source for the US:* Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS. *Source for Italy:* UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).


Figure A16: Controlling for Child Penalty



Notes: These figures use weekly earnings adjusted for the child penalty borne by women in the labor market (Kleven et al., 2019).

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 A17: Rank in Pay Distribution of Over-55 Men




Notes: Panels A and B show the average earnings 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 A18: Adjusting for changes in CPS top-coding and nonresponse


Panel A: Gender gap in weekly earnings




Panel B: Gender pay gap decomposition



Panel C: Share of mgmt. jobs




Panel D: Mean rank at age 25



Notes: These figures replicate the main findings after using Census rank-proximity-swapped earnings for 1976-2010 (the public-use files already incorporate rank swapping from 2011 onward), excluding records whose wage-and-salary earnings are item allocated, and applying both adjustments simultaneously.


Source: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure A19: Full-Time Workers

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 A20: 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:* Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.


Figure A21: Residual Earnings




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 A22: 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. Panels C and D show the same distributions of tenure for workers who were between 31 and 40 years old in 1986 in Italy.


Source: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

Figure A23: Firm-Level Workforce Aging and the Gender Pay Gap, Rolling 10-Year Changes




Notes: Each point on the x-axis represents the end year of a ten-year rolling window used to estimate the correlation between changes in the gender pay gap and workforce aging at the firm level. For example, the point labeled 1986 corresponds to regressions on the change from 1976 to 1986, the point labeled 1987 corresponds to regressions from 1977 to 1987, and so forth. Panel A regresses Δ gender pay rank gap on Δ share of workers 51-60. Panel B regresses 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 and the share of workers 51-60 at baseline. Panel C regresses Δ in the mean pay rank of younger men and women on Δ share of workers 51-60. Panel D regresses Δ in the mean pay rank of younger men and women on Δ share of workers 51-60, instrumenting the latter with differences between the share of workers aged 41-50 and the share of workers 51-60 at baseline. Controls include firm age, firm size, sector, and province fixed effects, all observed at baseline. Firm-level values for year x are computed as three-year averages from x to $x + 2$. In each panel, the sample of firms includes 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. Standard errors are clustered at the province level. *Source:* UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

Figure A24: Rolling-Regression Firm Samples versus the Full Population



Notes: The figure compares workers employed by the year-specific rolling-regression firms with the full population of private-sector employees. In each year the firm sample consists of the firms entering the corresponding rolling ten-year regression reported in Figure A23. Those regressions are estimated on firms with more than five employees, so the rolling-regression sample consists of larger firms by construction. Panel A reports the entry-earnings cohort decomposition. For every cohort, gender, and year, we compute mean log weekly earnings and employment shares among workers aged 25-64; earnings are held fixed at the cohort-gender mean observed at age 25, or at the 1976 mean for cohorts already older than 25 in the base year, while cohort employment weights are allowed to evolve. The figure reports both the observed and counterfactual changes in the gender gap for the rolling-regression firms and for the full population. Panel B reports, separately for men and women aged 25-30, the share of higher-paying managerial jobs in each sample held by each group. Higher-paying managerial jobs are those with weekly earnings in the top quartile of the national distribution. As in the baseline analysis, the managerial category includes *quadri* from 1996 onward. Panel C reports the mean earnings rank of 25-year-old men and women. All workers are ranked in the same contemporaneous national weekly-earnings distribution, so the two samples share a common reference point. In Panels B and C, blue denotes men and red women; solid lines with circular markers denote the full population and dashed lines with triangular markers the rolling-regression firms. *Source*: UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).


Figure A25: Between-Cohort Change Without Convergence in Entry Earnings



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 A26: 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.

Figure A27: Major-Predicted Gender Pay Gap



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 A1: 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. - 1980)	True entry earnings
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<u>Panel A: Log weekly earnings</u>								
United States (1976-2019)	0.800	-0.470	-0.598	-0.606	-0.615	-0.575	-	-
Italy (1976-2019)	0.333	-0.190	-0.241	-0.218	-0.197	-0.218	-0.157	-0.280
<u>Panel B: Log yearly earnings</u>								
United States (1976-2019)	0.835	-0.495	-0.621	-0.645	-0.657	-0.598	-	-
Italy (1976-2019)	0.350	-0.092	-0.204	-0.165	-0.127	-0.166	-0.061	-0.270
<u>Panel C: Log hourly earnings</u>								
United States (1976-2019)	0.568	-0.361	-0.452	-0.464	-0.458	-0.441	-	-

Notes: Column 1 shows the gender gap in the first available year, and Column 2 shows its change 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)). These columns differ in the definition of entry into the sample: 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 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. *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.

Table A2: First Stage of Firm-Level IV Regressions

	Δ share over-50 workers	
	(1)	(2)
<u>Panel A: 1976-1986</u>		
Share workers 41-50 – share workers 51-60 in 1976	0.58*** (0.01)	0.59*** (0.01)
Mean dep. var.	0.01	0.01
SD dep. var.	0.11	0.11
Obs.	25,279	25,279
<u>Panel B: 1976-1996</u>		
Share workers 31-40 – share workers 51-60 in 1976	0.45*** (0.01)	0.46*** (0.01)
Mean dep. var.	-0.00	-0.00
SD dep. var.	0.11	0.11
Obs.	13,993	13,993
<u>Panel C: 1996-2016</u>		
Share workers 31-40 – share workers 51-60 in 1996	0.39*** (0.01)	0.40*** (0.01)
Mean dep. var.	0.16	0.16
SD dep. var.	0.12	0.12
Obs.	14,644	14,644
Controls	No	Yes

Notes: This table shows the results of OLS firm-level regressions in which changes in workforce aging are regressed on 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, sector, and province fixed effects, all observed at baseline. Firm-level values for year x are computed as three-year averages from x to $x + 2$. In each panel, the sample of firms includes 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. Standard errors are clustered at the province level. *Source:* UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

Table A3: Firm-Level IV Estimates Using Alternative Definitions of Senior Workers

	Change in gender gap in pay rank		
	Ages 51–60 (1)	Ages 45–60 (2)	Ages 50–64 (3)
Panel A: 1976–1986			
△ Senior-worker share	-4.65*** (1.68)	-3.92** (1.66)	-5.18*** (1.60)
KP F-stat	11,259	6,888	14,574
Mean dep. var.	-1.75	-1.75	-1.75
SD dep. var.	19.94	19.94	19.94
Obs.	25,279	25,279	25,279
Panel B: 1976–1996			
△ Senior-worker share	-6.24** (2.57)	-4.87*** (1.62)	-5.03** (2.23)
KP F-stat	4,734	9,880	8,064
Mean dep. var.	-3.89	-3.89	-3.89
SD dep. var.	20.41	20.41	20.41
Obs.	13,993	13,993	13,993
Panel C: 1996–2016			
△ Senior-worker share	-0.93 (3.19)	-1.83 (1.77)	-0.29 (2.57)
KP F-stat	4,364	6,390	5,429
Mean dep. var.	-0.76	-0.76	-0.76
SD dep. var.	20.20	20.20	20.20
Obs.	14,644	14,644	14,644
Controls	Yes	Yes	Yes

Notes: The dependent variable is the change in the gender gap in mean pay rank among workers aged 25–30. Each column reports a 2SLS estimate for the indicated definition of senior workers using firms with more than five employees at the beginning of the period. All specifications control for baseline firm size and firm age and include baseline province and sector fixed effects. Standard errors clustered by province are reported in parentheses. KP F-stat denotes the Kleibergen–Paap first-stage F-statistic. * p<0.10, ** p<0.05, *** p<0.01. *Source:* UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

Table A4: Baseline and Sector-Age Shift-Share IV Estimates

	Change in gender gap in pay rank	
	Baseline IV (1)	Shift-share IV (2)
<u>Panel A: 1976–1986</u>		
△ Share aged 51–60	-4.68*** (1.68)	-7.52*** (2.83)
KP F-stat	11,317	4,157
Obs.	25,279	25,258
<u>Panel B: 1976–1996</u>		
△ Share aged 51–60	-6.32** (2.57)	-8.38*** (2.54)
KP F-stat	4,621	3,630
Obs.	13,993	13,976
<u>Panel C: 1996–2016</u>		
△ Share aged 51–60	-1.22 (3.18)	-0.87 (3.99)
KP F-stat	3,363	2,364
Obs.	14,644	14,630
Baseline firm controls	Yes	Yes
Province and sector FE	Yes	Yes
Baseline age-composition controls	No	Yes

Notes: The dependent variable is the change in the gender gap in mean pay rank among workers aged 25–30. The sample contains firms with more than five employees at the beginning of the period. Column 1 reports the baseline IV estimates. Column 2 uses the sector-age shift-share instrument and additionally controls for the baseline employment shares of workers aged 30–50 and 61–64. Both columns control for baseline firm size and firm age and include baseline province and sector fixed effects. Standard errors clustered by province are reported in parentheses. KP F-stat denotes the Kleibergen–Paap first-stage F-statistic. * p<0.10, ** p<0.05, *** p<0.01. *Source:* UNIEMENS, Istituto Nazionale della Previdenza Sociale (INPS).

B A Framework of the Gender Pay Gap with Career Spillovers

Solving the firm problem. The firm problem is

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

The FOCs are

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

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:

$$w_{y,b}^f = w_{y,b}^m = w_{y,b} = AY_{L_y} \theta_{y,b}.$$

Given that the wage in the top 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:

$$w_{y,t}^f = w_{y,t}^m = w_{y,t} = \mu_y w_{y,b}.$$

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

$$f_{y,t}^* = \frac{AY_{L_y} (\theta_{y,t} - \mu_y \theta_{y,b}) - \kappa K}{c_f}; m_{y,t}^* = \frac{AY_{L_y} (\theta_{y,t} - \mu_y \theta_{y,b}) - \kappa K}{c_m}.$$

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 to 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:

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

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

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

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

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

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 than 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. The number of top slots is:

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

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

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

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 :

$$[\text{Step 1}] \quad \frac{\partial f_{y,t}}{\partial l_{o,t}^{-1}} = \frac{c_m}{c_f} \frac{\partial m_{y,t}}{\partial l_{o,t}^{-1}}.$$

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

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

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

We simplify the notation:

$$B = A(\theta_{y,t} - \mu_y\theta_{y,b}) > 0; \quad D = (\theta_{y,t} - \theta_{y,b}) > 0; \quad E = \frac{c_m c_f}{c_f + c_m} > 0.$$

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

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

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:

$$1 > \frac{1}{\kappa - Y_{LyLy} BD - E} \{ B [Y_{LyLo} \theta_{o,t} - Y_{LyLy} D] - E \}$$

$$\kappa > \bar{\kappa} = BY_{LyLo} \theta_{o,t} = A (\theta_{y,t} - \mu_y \theta_{y,b}) Y_{LyLo} \theta_{o,t} > 0.$$

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 $(\theta_{y,t} - \mu_y \theta_{y,b}) > 0$. 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:

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

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

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

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

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:

$$\{f_y\} \quad \bar{w}_{y,f} = \frac{f_{y,t} \cdot \mu_y w_{y,b} + (f_y - f_{y,t}) \cdot w_{y,b}}{f_y} = \frac{1}{f_y} (\mu_y - 1) f_{y,t} w_{y,b} + w_{y,b};$$

$$\{m_y\} \quad \bar{w}_{y,m} = \frac{1}{m_y} (\mu_y - 1) m_{y,t} w_{y,b} + w_{y,b}.$$

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:

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

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 $\frac{\partial w_{y,b}}{\partial l_{o,t}^{-1}} > 0$. Due to the complementarity between the two age groups, a larger number of older workers makes younger men more productive, contributing to raising their mean wage. The same decomposition applies to the mean wage of younger women. 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:

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

Extension 1: 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. Second, consistent with the idea that the opportunity cost of retirement is increasing with wages, we assume that $\rho_{j,n}$ increases with firm-level productivity A_n of firm n . Third, firms set wages for the bottom and top jobs, instead of taking them as given. Fourth, the ratio of top and bottom wages is not equal to a fixed rent and is not necessarily constant across firms. Fifth, we assume that $\kappa(K_n)$ is 0 up to a threshold level $\bar{K}_n$ and then is ∞ beyond $\bar{K}_n$. In practice, this is equivalent to a binding constraint on the number of top jobs: $l_{o,t,n} + m_{y,t,n} + f_{y,t,n} = \bar{K}_n$. Sixth, to make the computations more tractable, we assume that the cost of hiring younger workers in the top job is linear, instead of quadratic.

In this extension, the timing 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, production takes place, 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,

$$\max_{w_{y,t,n}^m, w_{y,t,n}^f, w_{y,b,n}^m, w_{y,b,n}^f} A_n Y(L_{y,n}, L_{o,n}) - \sum_{g \in \{m,f\}} \sum_{a \in \{y,o\}} \sum_{j \in \{t,b\}} \left(w_{a,j,n}^g g_{a,j,n} \right) - \sum_{g \in \{m,f\}} c_g g_{y,t,n},$$

subject to the constraint: $l_{o,t,n} + m_{y,t,n} + f_{y,t,n} \leq \bar{K}_n$.

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 :

$$U_{i,a,j,n} = \log(w_{a,j,n}^g) + \xi_{i,a,j,n},$$

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-value distribution with a parameter $\sigma > 0$ 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 :

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

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:

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

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 et al., 2022).

Firms choose optimal wages as follows:

$$\begin{aligned} \left\{w_{y,b,n}^g\right\} : A_n Y_{L_{y,n}} \theta_{y,b,n} \frac{\partial g_{y,b,n}}{\partial w_{y,b,n}^g} - w_{y,b,n}^g \frac{\partial g_{y,b,n}}{\partial w_{y,b,n}^g} - g_{y,b,n} &= 0 \\ \left\{w_{y,t,n}^g\right\} : A_n Y_{L_{y,n}} \theta_{y,t,n} \frac{\partial g_{y,t,n}}{\partial w_{y,t,n}^g} - w_{y,t,n}^g \frac{\partial g_{y,t,n}}{\partial w_{y,t,n}^g} - g_{y,t,n} - c_g \frac{\partial g_{y,t,n}}{\partial w_{y,t,n}^g} - \lambda_n \frac{\partial g_{y,t,n}}{\partial w_{y,t,n}^g} &= 0, \end{aligned}$$

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

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

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

Instead, wages in the top job are equal to:

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

As in the representative-firm framework, the higher cost of employing younger women in top jobs ($c_m < c_f$) favors younger men's access to these positions. In this extension, this cost difference leads firms to offer younger men higher wages in top jobs, while bottom-job wages remain equal across genders. Consequently, younger men have a higher top-to-bottom employment ratio within each firm. Finally, all firms have $\bar{K}_n$ top jobs.

Next, we study a marginal increase in $l_{o,t}^{-1}$, 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.

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

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

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

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

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 $\frac{\partial \lambda_n}{\partial l_{o,t}^{-1}} > 0$ 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 : $\frac{g_{y,t,n}}{g_{y,b,n}} = \left(\frac{w_{y,t,n}^g}{w_{y,b,n}^g} \right)^{\frac{1}{\sigma}}$. The derivative of this ratio with respect to the number of older workers is as follows:

$$\frac{\partial \frac{g_{y,t,n}}{g_{y,b,n}}}{\partial l_{o,t}^{-1}} = \frac{1}{\sigma} \left(\frac{w_{y,t,n}^g}{w_{y,b,n}^g} \right)^{\frac{1}{\sigma}-1} \frac{\frac{\partial w_{y,t,n}^g}{\partial l_{o,t}^{-1}} w_{y,b,n}^g - \frac{\partial w_{y,b,n}^g}{\partial l_{o,t}^{-1}} w_{y,t,n}^g}{\left(w_{y,b,n}^g \right)^2} = \frac{1}{\sigma} \frac{g_{y,t,n}}{g_{y,b,n}} \left(\frac{\partial w_{y,t,n}^g}{\partial l_{o,t}^{-1}} \frac{1}{w_{y,t,n}^g} - \frac{\partial w_{y,b,n}^g}{\partial l_{o,t}^{-1}} \frac{1}{w_{y,b,n}^g} \right).$$

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 $\partial \frac{m_{y,t,n}}{\partial l_{o,t}^{-1}}$ to $\partial \frac{f_{y,t,n}}{\partial l_{o,t}^{-1}}$. The wage change at the top $\frac{\partial w_{y,t,n}^g}{\partial l_{o,t}^{-1}}$

is the same for both genders, but younger men are paid more in top jobs. Therefore, $\frac{\partial w_{y,t,n}^m}{\partial l_{o,t}^{-1}} \frac{1}{w_{y,t,n}^m}$ is smaller than $\frac{\partial w_{y,t,n}^f}{\partial l_{o,t}^{-1}} \frac{1}{w_{y,t,n}^f}$. 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 declines more than the share of women in the same positions.

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):

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

The derivative is positive if $\frac{\partial w_{y,b,n}^g}{\partial l_{o,t}^{-1}} \frac{1}{w_{y,b,n}^g} > \frac{\partial w_{y,b,n'}^g}{\partial l_{o,t}^{-1}} \frac{1}{w_{y,b,n'}^g}$. [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 $\frac{\partial \frac{m_{y,b,n}}{m_{y,b,n'}}}{\partial l_{o,t}^{-1}}$ and $\frac{\partial \frac{f_{y,b,n}}{f_{y,b,n'}}}{\partial l_{o,t}^{-1}}$ are the same because all their components are the same between younger men and women. 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.

Extension 2: Introducing skills. For this extension, we go back to the simpler framework with a representative firm. 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 $\mathbf{L}_y = (L_{y,h}, L_{y,l})$, and workers with different skills are complements in the production function. The rest of the firm problem is unchanged.

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:

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

where $\frac{1}{m_y} (\mu_y - 1) m_{y,t,s} + 1 = \frac{\mu m_{y,t,s} + m_{y,b,s}}{m_y} = \frac{m_{y,s}}{m_y}$ 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, $\frac{\partial w_{y,b,s}}{\partial l_{o,t}^{-1}} = A \theta_{y,b,s} \cdot \left(Y_{L_{y,s} L_{y,s}} \frac{\partial L_{y,s}}{\partial l_{o,t}^{-1}} + Y_{L_{y,s} L_o} \frac{\partial L_{o,s}}{\partial l_{o,t}^{-1}} \right) = A \theta_{y,b,s} Y_{L_{y,s} L_o} \theta_{o,t,s} \rho_t$ because $\frac{\partial L_{y,s}}{\partial l_{o,t}^{-1}} = 0$ 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:

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

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 is 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: $Y_{L_{y,s} L_o} = Y_{L_{y,s}} \cdot C(\mathbf{L}_y, L_o)$, where $C(\cdot, \cdot)$ is the same across tasks.³⁰ In this case, the change in mean wages for younger women

³⁰This 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

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

is larger than that for younger men if:

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

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

Assuming that high-skill tasks have a 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: $\frac{f_{y,h}}{f_y} > \frac{m_{y,h}}{m_y}$.

C Sensitivity to Selection into Employment

Differential selection of men and women into employment raises the question of whether trends in observed weekly earnings also characterize trends in potential (or offered) weekly earnings. If the relationship between potential earnings and employment differs by gender or changes over time, the gaps in observed and potential earnings need not evolve similarly. The literature has addressed this problem using parametric selection models (Heckman, 1979; Mulligan and Rubinstein, 2008), bounds (Blundell et al., 2007), quantile-selection models (Maasoumi and Wang, 2019; Blundell et al., 2025), control-function methods (Fernández-Val et al., 2023), and imputation (Blau et al., 2024). We implement the probability-weighted imputation procedure of Blau et al. (2024) as a sensitivity exercise and examine whether our main conclusions survive under its identifying assumptions.

C.1 Summary of the Imputation Method

The empirical challenge is that potential weekly earnings of non-employed people are unobserved. Blau et al. (2024), building on Olivetti and Petrongolo (2008), proposes a method of imputing potential earnings for non-employed individuals by applying a three-step approach to PSID panel data. Each step augments its 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 at least 100 hours during the previous year. Assign them their actual, observed wage.
2. Add currently non-employed persons who were employed in nearby past or future years. Assign them the wage from those recent employment spells.
3. For the remaining individuals aged 25–54 who have not yet been assigned potential earnings, use a probability-weighted imputation method in which observable characteristics 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 in the last group, observations are tentuplicated and weighted by the predicted probability of belonging to each decile. Then, the median 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 on either side of the median. While the mean does not share this robustness property, in practice, Blau et al. (2024) shows that the selection adjustment results for the median and mean pay gaps are very similar to each other.

C.2 Implementation

We follow Blau et al. (2024) as closely as possible, with the caveat that we use ASEC repeated cross-sections and that the variables available to impute probabilities differ slightly from those 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 it to the Italian administrative data.

C.2.1 Samples

Throughout the selection adjustment analysis, our target population consists of individuals 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 or public-sector employees in the distribution of potential earnings.³¹ 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 at ages close to retirement. As in our main analyses, our baseline compensation measure is weekly earnings, which we refer to as wages for short in the remainder of this appendix.

Baseline sample 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 at least 100 hours during the last year. These are people who had weaker attachment to the labor market but for whom we directly observe wage information. We refer to the baseline sample as Sample 1 and to the sample combining Sample 1 with these additional individuals as Sample 2.

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 $q = 1, \dots, 9$.
- (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, a Hispanic dummy, and region of residence dummies (the nine Census Divisions).
- (c) We refer to our final sample, which includes potential earnings for everyone in our target population, as Sample 3. For each person in Sample 3 who is not in Sample 2 (Sample 3 “entrants”), we assign 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

³¹In Section 6.7, we replicate all our main findings using a sample that (i) adjusts for time-varying selection into the labor market and (ii) includes public-sector employees.

replicated observations a weight π_{id} that is equal to $\pi_{id} = \hat{p}_{i,d} \times c_i$, where c_i is the CPS survey weight for individual i , for observations/deciles $d = 1, \dots, 10$. 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 median of the corresponding wage decile.
- (e) Using Sample 3 and weights c_i and π_{id} (for those with observed and imputed wages, respectively), 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 “adjusted” for selection (labeled BKBC, short for [Blau et al. \(2024\)](#)), and to statistics computed in Sample 1 as “baseline.”

C.2.2 Identifying Assumptions and Choice of Method

To state the identifying restriction precisely, fix a gender-year cell ($G_i = g, T_i = t$) and let $R_i = 1$ indicate that individual i belongs to Sample 2 and therefore has observed weekly earnings, while $R_i = 0$ indicates that earnings must be imputed. Let w_i denote potential weekly earnings, and let $D_i \in \{1, \dots, 10\}$ denote the decile containing w_i in the gender-year-specific potential-earnings distribution. The central selection-on-observables assumption is:

$$\Pr(D_i = d \mid X_i = x, R_i = 0, G_i = g, T_i = t) = \Pr(D_i = d \mid X_i = x, R_i = 1, G_i = g, T_i = t), \quad d = 1, \dots, 10.$$

Thus, conditional on X_i , Sample 2 is assumed to provide the appropriate distribution of potential-earnings decile assignments for otherwise comparable individuals without observed earnings. The distribution of X_i may differ between the two groups, and the procedure does not impose either positive or negative selection. It does, however, rule out residual selection on unobserved determinants of potential earnings after conditioning on X_i .

Alternative approaches allow employment and earnings unobservables to be related, but replace conditional selection on observables with identifying restrictions often based on excluded variables. [Maasoumi and Wang \(2019\)](#) and [Blundell et al. \(2025\)](#) use the quantile-selection framework of [Arellano and Bonhomme \(2017\)](#), which links the unobserved ranks in the earnings and employment equations through a parametric copula. Their applications also use excluded variables: [Maasoumi and Wang \(2019\)](#) uses the presence of young children, while [Blundell et al. \(2025\)](#) uses unemployment rates and simulated disposable-income measures constructed from tax and transfer rules. [Fernández-Val et al. \(2023\)](#) instead uses annual hours as a continuous selection variable in a control-function framework with household-composition excluded variables in its baseline specification.

The alternative approaches address selection on unobservables through substantive modeling restrictions and may also require substantial additional data construction. Specifications that use young children as excluded instruments require children not to affect potential earnings directly,


conditional on the included controls. Approaches that use hours or weeks worked to model selection instead rely on assumptions about the working-time process and its relationship with unobserved earnings determinants. [Blundell et al. \(2025\)](#) uses variation in tax and transfer policies to help identify selection into employment. Implementing its approach in our setting would require assembling historical tax and transfer rules and constructing simulated disposable-income measures. This would be a valuable extension, but it goes beyond the scope of our sensitivity analysis. For our purposes, the [Blau et al. \(2024\)](#) procedure provides a transparent and relatively parsimonious sensitivity exercise. It adjusts the distributions of men and women separately in every year, does not impose a particular sign of selection, and does not require an excluded variable that affects employment but not potential earnings. The procedure also generates a complete gender-year earnings distribution from which we can recompute the mean gender gap, the cohort decomposition, and our positional statistics. It is therefore well matched to our question of whether the paper's findings survive a reasonable adjustment for selection into employment.

C.2.3 Results

Figure [A9](#) shows the robustness of our main findings in the selection-adjusted sample. Below, we show additional results on the adjusted and baseline wage distributions. Panel A of Figure [C1](#) shows the evolution of the adjusted and baseline 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 adjusted gap representing a slight parallel upward shift (between 0.01 and 0.04 log points) relative to the baseline gap. We find similar results when zooming into the gender pay gap at age 25. Panel B of Figure [C1](#) shows that the adjusted gap at age 25 follows a very similar trend to the baseline gap: steep convergence from 1976 until the mid-1990s/early-2000s, and little systematic change thereafter. The similarity in the adjusted and baseline trends in Figure [C1](#) underpins the fact that our main results are robust to implementing this selection adjustment method and, as a result, unlikely to be driven by changes over time in the nature of selection into employment.

However, 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 adjustment 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 adjustment 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-adjusted 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 adjusted 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: Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]. Minneapolis, MN: IPUMS.

Figure C2: Baseline and Selection-Adjusted Earnings at Selected Percentiles



Notes: Trends in the 10th, 50th, and 90th percentiles of the raw and selection-adjusted 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: 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 for adjusting weekly earnings for the child penalty in both the United States and Italy. The methodology closely follows [Kleven et al. \(2019\)](#), [Kleven \(2023\)](#), and [Casarico and Lattanzio \(2023\)](#). It relies on comparing the effects of childbirth on women and men in the United States and on comparing mothers with 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:³²

$$\tilde{w} = \frac{w}{(1 - p\pi)}, \quad (5)$$

where w denotes 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. In the following paragraphs, we describe how we calculate each of these components.

A decline in the estimated contribution of the child penalty across cohorts could in principle reflect either a smaller per-mother penalty p or a smaller share of women who have already become mothers by a given age, π — for instance, if childbearing is delayed to older ages in more recent cohorts. To rule out the latter as a purely mechanical explanation, we allow both p and π to vary flexibly by age, cohort, and calendar year, rather than holding either fixed. In the United States, as described below, this is achieved by assigning each identified mother her own cohort- and event-time-specific penalty, so the correction reflects each woman’s actual fertility timing rather than an assumed, constant rate of motherhood. In Italy, as described below, we directly estimate the age profile of motherhood in each cohort and year, which reveals a postponement of childbearing toward older ages — precisely the pattern that could otherwise confound a declining child penalty with delayed fertility.

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 she were not a mother:

$$\tilde{w}_{it} = \frac{w_{it}}{1 - p_{jc(i,t)}}, \quad (6)$$

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 from the March CPS for 1968–2020 and the American Community Survey (ACS) for 2000–2019. We follow its 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

³²To simplify notation, we omit time and cohort subscripts here.

(see Kleven (2023) for more details). We then estimate the following earnings equation, separately for men and women:

$$w_{it}^g = \alpha^g D_{it}^{Event} + \beta^g D_{it}^{Age} + \gamma^g D_{it}^{Year} + v_{it}^g, \quad (7)$$

where w_{it}^g denotes 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 childbirth. The event-time coefficients α^g measure the impact of childbirth 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:

$$P_t^g = \frac{\alpha_t^g}{E[\tilde{w}_{it}^g | t]}. \quad (8)$$

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:

$$p = E[P_t^m - P_t^w | t \geq 0] - E[P_t^m - P_t^w | t < 0] \quad (9)$$

Results. Figure D1 presents the pooled event study for men and women surrounding the birth of their first child. The event-study horizon displayed in this and subsequent figures ranges from $t = -5$ years to $t = 10$ years. The average child penalty on weekly earnings conditional on working, over this event-time horizon, is 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 (6) 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, 1960–64, 1965–69, 1970–74, 1975–1979, and 1980–95.³³ Figure D2 shows the estimated child penalties for each of these nine birth cohort groups. The child penalty on weekly earnings conditional on working ranges from 11% to 23%. Even though the pattern is not monotonic, the size of the child penalty tends to decrease across birth cohorts.

The estimates for event times 0–10 in each of the nine birth cohort groups represent our adjustment factors p_{jc} in Equation (6). 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 child.³⁴ This cell assignment determines the adjustment factor p_{jc} assigned to each mother and her resulting counterfactual weekly earnings net of the child penalty. Women who are not observed as mothers at a given age receive no adjustment. The correction is therefore implemented at the individual level rather than through an aggregate π : it automatically reflects each woman’s own, cohort-specific fertility timing, so a shift toward later childbearing in more recent cohorts does not by itself generate a spurious decline in the estimated child penalty.

³³The first and last birth cohort groups span a larger number of birth years because these are the groups with fewer observations.

³⁴We apply the 10-year child penalty to mothers whose eldest child is older than 10.

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 leaves. 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 begin one to two months before the expected childbirth and continue for three to four months afterward. The contribution archive is only available for a sample of workers born after 1950. We therefore restrict the data to workers included in this 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 ages 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, we recover a control group of non-mothers directly from women born between 1950 and 1974, who are observed through age 45 and are therefore not subject to right-censoring. Because women born after 1974 had not yet turned 45 by the end of our sample (2019), we predict their probability of remaining childless from a linear probability model estimated on the non-truncated cohorts, using earnings-quartile and province dummies as predictors,³⁵ and assign the women with the highest predicted probability to the control group until the predicted share of never-mothers among truncated cohorts matches the observed share among non-truncated cohorts—an assumption that only marginally affects the size of the child penalty (Casarico and Lattanzio, 2023). The control group therefore combines actual never-mothers from non-truncated cohorts with predicted never-mothers from truncated cohorts. We then assign each control-group woman a placebo age at first birth, drawn from the age-at-childbirth distribution of actual mothers within her birth-cohort-by-earnings-quartile cell.³⁶ For truncated cohorts, we shift this distribution’s mean upward using a time trend estimated on actual mothers’ age at first birth, so that predicted never-mothers are assigned a placebo birth age consistent with the postponement of childbearing documented below, rather than with the (right-censored) raw distribution.

Estimating Child Penalties by Calendar and Childbirth Year. Our goal is to estimate the child penalty by calendar and first childbirth year. To do so, we run the following event study, separately for mothers and non-mothers:

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

where we interact event-time dummies (year relative to first childbirth) with year dummies in order to estimate the year-specific coefficients α_{yj}^g (which are equivalent to childbirth-year-specific coefficients, as year of childbirth $c = s - t$). The equation also controls for age and year dummies. Figure D3 reports the average child penalty estimates (the average difference in coefficients for

³⁵We estimate $\text{Never-Mother}_{it} = \alpha + X'_{it}\beta + \varepsilon_{it}$, where Never-Mother_{it} is a dummy for never-mothers in birth cohorts 1950–1974 and X_{it} includes quartiles of the cohort-specific log weekly earnings distribution and province of residence.

³⁶For actual never-mothers, we draw from a log-normal distribution fitted to the observed age at childbirth of mothers in the same cohort-by-earnings-quartile cell; for predicted never-mothers, we use the same variance but a higher, trend-adjusted mean, since we do not observe their (right-censored) completed fertility.

mothers and non-mothers) in the first 10 years after first childbirth for different years of childbirth. The estimates hover around 6-8%. We assume that the relative stability in the child penalty also applies to the earlier childbirths. To this end, we fit a linear trend and assign to the childbirths before 1990 the predicted child penalty from later ones.³⁷


Time-Varying Prevalence and Timing of Motherhood. This is also where a decline in the estimated contribution of the child penalty could be confounded with a purely mechanical shift in when women become mothers, so—rather than assuming a constant prevalence of motherhood—we estimate the prevalence of motherhood separately by age, 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 D4 shows the share of new mothers by year and for different age groups in Panel A. As expected, the share of new mothers displays an inverted U-shaped relationship with age: it is small at earlier ages, peaks around 30 years old, and then declines. The shift toward later childbearing is reflected in 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 use the fitted values to predict the share of new mothers, reported as lines in Figure D4. 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 adjust weekly earnings in Equation 5. Figure D5 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 growth in the wage gap: the gender gap would be 0.13 log points at 40 years old, close to 20% lower than that observed in the data.

This correction speaks directly to whether the decline in the child penalty’s contribution to the gender pay gap is itself an artifact of delayed motherhood. Because the correction already nets out the changing prevalence of motherhood at each age—individually in the United States and through the directly estimated, time-varying π in Italy—before evaluating whether the per-mother penalty has declined, a small remaining role for the child penalty cannot simply reflect fewer women having had children by a given age. If postponed motherhood were driving the between-cohort convergence in the gender pay gap, a correction that incorporates the shifting age profile of motherhood so directly would have removed most of it; instead, it leaves the main pattern largely intact.


³⁷ We choose 1990 as a threshold because this is the earliest first-childbirth cohort for which we observe enough observations per mother.

Figure D1: Child Penalty Event Study: Weekly Earnings Conditional on Working




Notes: Event study of weekly earnings for men and women around the birth of their first child at $t = 0$. The series show the percentage impact of childbirth on men and women at each event time t , i.e., $\hat{P}_t^m$ and $\hat{P}_t^w$ estimated from Equations (7)-(8). The figure also displays the average child penalty over event times 0 – 10 defined as in Equation (9). Age at first birth is restricted to 25–45 years. 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 D2: Child Penalty by Birth Cohort: Weekly Earnings Conditional on Working




Notes: Event studies of first childbirth 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 (7) is estimated separately for each cohort group. Each panel displays the average child penalty over event times 0-10 (defined in Equation 9) for the birth cohort group 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 D3: Child Penalty Estimates by Childbirth Year




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

Figure D4: Share of Mothers by Year and Age



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

Figure D5: 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 adjusting women's weekly earnings for 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).