Prime Years, Hidden Bias: Middle Ageism in the Labor Market

Lizi Yu Shuo Zhang
Simin Yuan Taoxiong Liu Kebin Dai

April 14, 2025

Abstract

This paper examines age-based hiring bias using 14 million applications from a major Chinese job platform. We address three empirical challenges: (1) providing evidence of age bias *at the margin*, complementing *average* discrimination identified by correspondence studies (Heckman, 1998); (2) leveraging the platform's process to present limited applicant information to reduce omitted variable bias; (3) estimating age bias among closely comparable cohorts, where applicants are just one year older than their peers. We find hiring chance declines after age 30, with a pronounced effect between 35 and 50. This "middle-ageism" effect emerges earlier for women and is prevalent across industries. (*JEL* J14, J16, J23, J71, M51)

Yu: University of Queensland (email: lizi.yu@uq.edu.au); Zhang: Northeastern University (email: s.zhang@northeastern.edu); Yuan: Tsinghua University (email: ysm22@mails.tsinghua.edu.cn); Liu: Tsinghua University (email: liutx@mail.tsinghua.edu.cn); Dai: Tongdao Liepin Group. The authors wish to thank Joanna Lahey, Patrick Button, Peter Kuhn, Xincheng Qiu, Heather Royer, Kailing Shen, Anna Aizer, Andrey Fradkin, Mindy Marks, and participants at the ANU General Seminar, UQ Applied Seminar, WEAI Conference, Australian Conference of Economists, UNSW CEPAR Seminar for helpful comments.

"Young people are just smarter."

– Mark Zuckerberg, 2007

1 Introduction

Older workers often have fewer opportunities to find or maintain employment. In the past, employers openly stated age requirements in job postings.[1] More recently, age requirements in job ads are banned in many countries, but age bias still persists in the recruitment process. Correspondence resume studies provide the most credible evidence of identifying age discrimination in recruitment in the existing literature (Rosen and Jerdee, 1976; Bendick Jr et al., 1997; Lahey, 2008; Büsch et al., 2009; Riach and Rich, 2010; Ahmed et al., 2012; Carlsson and Eriksson, 2019; Neumark et al., 2019; Van Borm et al., 2021). These studies create fictitious applicants that differ only in age, make older and younger applicants apply for the same jobs, and compare their callback rates.[2] Recognizing their advances in identifying *average* discrimination, researchers often note that correspondence studies are not informative about discrimination *at the margin* (Heckman, 1998). Discrimination *at the margin* refers to the idea that the impact of discrimination is determined not by the average market-level disparities (e.g. wage gaps or employment rates), but by the "level of discrimination at the firms where ethnic minorities or women actually end up". In the real world, job seekers apply for jobs where they believe they have the best chance of success; consequently, older applicants avoid jobs that appear more ageist (Firoozi et al., forthcoming). This indicates the importance of considering job seekers' sorting and selection when measuring discrimination – an issue that correspondence studies cannot address.

In contrast to audit and correspondence studies, the traditional regression-based approach using established data suffers from spurious omitted variable bias when identifying discrimination,

[1] Before the enactment of the U.S. 1967 Age Discrimination in Employment Act (ADEA), nearly 60% of job advertisements included explicit age requirements (Neumark, 2018). Similarly, 40% of job postings in Spain, Greece, and the UK imposed upper age limits before the EU Framework Directive 2000/78 (Lahey, 2010; Drury, 1994; Bokum and Bartelings, 2009). This trend persists in countries where stating age requirements remains legal: 24-80% of job ads in China, Mexico and Japan specified age requirements (Helleseter et al., 2020; Zvedelikova, 2024).

[2] To match on applicants' work experience, some studies assign employment gaps or out-of-field experience on older applicants to guarantee the equal length of experience (Bendick Jr et al., 1997; Lahey, 2008), and some studies emphasize experience commensurate with age (Neumark et al., 2019).

because applicants may differ in productivity proxies that are not available to researchers but observed by recruiters. To ensure that "the researcher observes all that the decision-maker observes" is considered "a hopeless task" (Bertrand and Duflo, 2017).

Moreover, existing literature on age bias primarily focuses on populations aged 50 and above and compares the outcomes of two groups (i.e., young vs. old) that may be 20 years apart. The questions of when age discrimination begins and how it evolves across various age groups remain largely underexplored. It is unclear whether age bias gradually develops as workers transition from younger to older age groups or only gets significantly worsened when workers reach certain ages. This highlights the need to examine bias at a more granular level to estimate the potential impact of smaller age increments.

In this paper, we address these limitations and present robust evidence of age discrimination in recruitment, leveraging 14 million application records submitted to 451,692 job ads from a leading online recruitment website ("platform" hereafter) in China. This data set represents a 5% random sample of positions posted on the platform, from 2018 to 2021, and all applications submitted to these jobs within the 90-day window. It is broadly representative of China's white-collar labor market across various industry sectors, occupations, and seniority levels. Our analysis controls for job-ad fixed effects, resembling the real-world scenario wherein a recruiter selects qualified candidates upon receiving a stack of applications for a posted job. It reflects a more 'organic' job search process, where applicants may select job ads strategically to avoid the most discriminatory employers, and their selection is treated as given.

We alleviate the omitted variable bias in regression-based studies, by exploiting the platform's automated process that presents applicant information to recruiters in a standardized format. On this platform, recruiters are initially presented with an applicant's summary card, which outlines only basic information about the applicant – including age, and recruiters can choose whether to review the applicant's full resume (*FullReview*). Upon reviewing resumes, recruiters can then opt to contact selected applicants (*Callback*). The empirical feature of presenting summary card has two advantages: First, summary cards and resumes are generated by the platform, using information entered by job applicants into their personal profiles, so variables we observe almost reflect all major productivity components that have been made available to recruiters. This standardized practice also removes concerns that variations in resume quality affect job search success (Bertrand and Mullainathan, 2004). Second, when a complete resume is presented, as is typical in the hiring process, it is difficult to determine which information recruiters rely on. Anecdotal evidence suggests that recruiters spend only 6-10 seconds skimming resumes, potentially leading to biased attention allocation (Bartoš et al., 2016). In our context, presenting summary card better ensures us that age is a relevant factor considered by recruiters.

Our analysis focuses on applicants aged 25 to 55, representing the main body of China's urban labor force.[3] We start by presenting preliminary evidence from an estimation over a complete age profile, examining the likelihood of receiving *FullReview* and *Callback* at each applicant age from 26 to 55, using 25-year-old applicants as the reference group. We control for job-ad fixed effects and date fixed effects to account for variations related to the posted positions and variations relating to the seasonal trend. Without controlling for other individual characteristics or productivity proxies, this estimation captures the combined influence of both age bias and any potential benefits of accumulated work experience. Overall, we observe an inverted U-shaped pattern for *FullReview*: The chance of applicants' resumes receiving full reviews starts with an initially modest increase from ages 26 to 30, and then followed by a significantly sharper decline after age 30. As a result, a 38-year-old applicant is already viewed as less employable than a 25-year-old. Among those whose resumes are reviewed, the likelihood of receiving callbacks also decreases for applicants over 30, though the decline is less steep.

Next, we estimate age bias by comparing applicants who are only one year older than their reference counterparts across a span of closely matched cohorts. We sequentially contrast 30 pairs – from age 26 versus 25 to 55 versus 54 – to estimate age bias within narrowly defined age intervals. When examining *FullReview* or *Callback*, we control for job-ad fixed effects and date fixed effects, as well as variables presented on the summary card or resume correspondingly to account for productivity components visible to recruiters at each decision point.

Remarkably, even among these closely comparable applicants within the same age cohort, we consistently find age bias against the one-year older applicants when they are in their 30s and 40s, affecting both *FullReview* and *Callback*. This bias emerges subtly but statistically significantly when applicants reach early 30s and persist until they reach 50. For applicants aged 31-50, on

[3]During the analysis period, the statutory retirement age in China is 60 for men, 50 for women in blue-collar jobs and 55 for women in white-collar jobs. In 2019, individuals aged 25-59 constituted 82.1% of China's total labor force China Statistical Yearbook (2019), with a higher proportion in urban employment sectors.

average, there is a 0.51 percentage point (1.47 percent) decrease in receiving *FullReviews* with each additional year of age. Moreover, while age bias develops gradually affecting middle-aged workers, they experience extra age penalty when reaching 40 and 50: An applicant aged 40 (50) is 3.22 (4.70) percent less likely to be reviewed compared to a similar applicant aged 39 (49), suggesting that age bias near round-number ages is two to three times greater than at other ages. Among those whose resumes are reviewed, being just one year older leads to an additional 0.37 percentage point (1.66 percent) decrease in the likelihood of receiving *Callbacks*.

Our findings are robust to a number of alternative measurement assumptions and sample restrictions, and are further validated by two quasi-experimental analyses. In the first, we exploit the feature that only limited information is presented on the summary card to construct a sample of identical applicants, by keeping only those whose summary card information – excluding age – is identical within each job's application pool. In the second, we replicate our main results using worker fixed effect, leveraging the natural increase in an applicant's age to compare their job search success before and after the age change within a 30-day window. Results from both quasi-experiments are consistent with our main findings.

Our heterogeneity analysis examines the age-related penalty on the likelihood of resume views across various factors, including job and employer characteristics, labor market conditions, industry sectors, occupation types, and worker characteristics. The findings indicate that the age effect remains consistent across most groups, with some contexts where bias is more pronounced. Specifically, age bias against older applicants tends to be stronger in jobs that require lower education, less work experience, non-managerial roles, jobs that receive more applications, and those posted by employers with potentially younger workforce. Furthermore, female applicants encounter age-related bias earlier in their careers, though their magnitudes at most ages are similar to that experienced by men.

One major distinction of our paper from previous studies is that we find robust evidence of age bias beginning at ages earlier than what have been documented in the literature. Such '*middle ageism*' effect has been anecdotally reported to exist in the labor market (Kelly, 2023), but has rarely been studied in the literature. Determining the age at which workers are considered 'old' is challenging. One revealing perspective is to examine how age is factored into the assessment of skilled immigrants. Australia and Canada award maximum immigration points to applicants aged

25-33 and 18-35, respectively. In the US, Republican senators proposed the *RAISE Act* (Reforming American Immigration for Strong Employment) in 2017, prioritizing applicants aged 22-35 as the most demanded group. Another perspective can be gained by examining the implied age suggested by implicit age terms used in job ads. For instance, the case of *Kleber v. Carefusion Corp* indicates how employers use experience caps – such as "3 to 7 years (no more than 7 years)" – to screen out applicants above 30s.[4]

Workers may be more aware of when they begin to be seen as 'old'. A recent study by Firoozi et al. (forthcoming) find that applicants aged around 35 start to avoid applying for positions that use ageist language in job ads, suggesting that they start to discern age bias and already perceive themselves to be less preferred when faced with ageist-phrased job ads. In China, a recent workforce survey (Zhilian Zhaopin, 2023) reports that 60.2% of respondents express the need to address age discrimination, ranking it as the most severe form of discrimination in the workplace, and nearly 85% respondents acknowledge the "Curse of 35" in career trajectories.

Our study enriches the existing literature in four significant ways. Firstly, we extend the comprehensive literature on age discrimination centered on the population over 50 or nearing retirement (Neumark et al., 2019; Riach and Rich, 2010; Lahey, 2008; Bendick Jr et al., 1997), which often attribute ageism to health or physical issues, or a deficit in technological proficiency (Van Borm et al., 2021). Most existing studies compare callbacks between their defined 'old' and 'young' groups, leaving the dynamics of what happens in between unclear. Our findings of 'middle-ageism' suggest that age discrimination occurs gradually as workers age and commences earlier than expected, affecting workers in their 30s – a period traditionally regarded as prime age. This finding aligns with Farber et al. (2019) and Carlsson and Eriksson (2019), who examine age bias across a broader age range in correspondence studies; with Helleseter et al. (2020) and Zvedelikova (2024), who analyze age preferences in job advertisements and report that explicit age limits most frequently set the eligibility cutoff at 35; and relates to Fang and Qiu (2023), who find that Chinese workers reach their labor market peak around age 35. Our paper provides an additional empirical explanation that the challenge of job search beginning early 30s can constrain wage growth beyond this age.[5]

[4] In countries where setting age requirements in job ads is legal, jobs most frequently set the cutoff at 35 years old (Helleseter et al., 2020; Zvedelikova, 2024).

[5] The labor market peak comes much later in life in the US, when employees reaching 50s. Lagakos et al. (2018)

Secondly, this paper complements the existing literature of age discrimination based on correspondence studies (Lahey, 2008; Büsch et al., 2009; Riach and Rich, 2010; Ahmed et al., 2012; Carlsson and Eriksson, 2019; Neumark et al., 2019; Van Borm et al., 2021). Our paper provides the first evidence of age discrimination at the margin and quantifies the effect of being one year older across the 25 to 55 age range, even after accounting for job seekers' search strategies to avoid highly discriminatory employers. The estimated age bias persists across various industries, occupations, and seniority levels, reinforcing the external validity of previous research.

Thirdly, our research joins the extensive body of work on labor market discrimination, where a critical aspect is on distinguishing between statistical discrimination and taste-based discrimination (Becker, 1957; Arrow, 1972; Phelps, 1972; Guryan and Charles, 2013). Statistical discrimination involves making inferences about an individual's quality based on the average characteristics of their group. Plausibly, it becomes more difficult to differentiate the group average as the number of groups increases and the distinctions among individuals within each group become less pronounced. Our analysis to examine age bias when applicants are only one year older, together with the two quasi-experiments where applicants are mostly identical except for their age, enables us to minimize statistical discrimination to the greatest extent. This suggests that the observed age-related recruitment disparities are, to a large extent, caused by taste-based discrimination.

Lastly, our study is also relevant to gender discrimination in hiring, particularly concerning young women. Existing research indicates that age discrimination manifests differently by gender (Helleseter et al., 2020). Our findings suggest that age bias occurs earlier for female workers, suggesting that concerns about potential family responsibilities contribute to middle-ageism on female workers (Petit, 2007; Kleven et al., 2019; He et al., 2023).

The rest of the paper proceeds as follows. Section 2 describes the empirical background and data. Section 3 presents the preliminary evidence of the complete age profile. Section 4 presents our main analysis when applicants age by one year. Section 5 provides robustness checks and Section 6 explores potential explanations and heterogeneity. Section 7 concludes.

compares life cycle wage in developed and developing countries and find flatter wage profiles in poor countries, attributed to the lower human capital accumulation and larger job search frictions in the developing world.

2 Background and Data

2.1 The Platform

The platform under study is *Liepin.com*, one of the leading internet recruiting platforms in China, holding nearly one fifth of the market share during the analysis period from 2018 to 2021. It operates with a standard online job platform design: Recruiters post job ads, job seekers submit online applications to the posted ads, then recruiters review and process received applications in the platform's online environment. Each component is detailed further below.

Recruiters post job ads: Employers subscribe and incur fees for posting job advertisements and using recruitment tools. Job advertisements detail the job's characteristics, such as the offered wage, the industry, occupation, and location, along with requirements for the target candidate, including their levels of education and work experience. A job posting remains active for up to 90 days unless closed earlier by the recruiter. For the firms posting jobs, we observe their firm size, ownership type, as well as the hiring agent associated with posting the job and managing the subsequent processing.[6]

Job seekers submit applications: Job seekers create their profiles for free by filling out their information including their name, age, gender, current wage, marital status, working experience, educational background, employment status, job title, and their current industry and occupation. Job seekers can browse advertisements and apply to jobs of interest. Following an application, the platform recommends similar positions, allowing job seekers to apply to recommended roles with a single confirmation using the "Batch Apply" option – batch applications cannot be identified by recruiters.[7]

Recruiters review applications: On the recruiter's portal, applications first appear as a series of summary cards, sorted by application time with the newest applications first, and recruiters can scroll down to progressively load additional summary cards for review. The summary card only presents limited information: the applicant's name, gender, age, location, educational level, years of experience, job title and the current employer. A recruiter can click on the summary card to

[6]Firm size is categorized based on the number of employees: fewer than 100, 100 - 1000, 1000 - 5000, and over 5000 employees. Firm ownership is categorized as: foreign-invested, private-owned, state-owned, and publicly-listed enterprises.

[7]This feature is also used on other job search websites to simplify and streamline the application process, such as LinkedIn's "Easy Apply", and Indeed's "Indeed Apply".

review a job applicant's complete resume, defined as *FullReview* response in our analysis.

Recruiters give callbacks: Although we cannot observe employers' recruitment decisions, we can infer the interest of hiring agents in specific candidates based on their actions on the platform. After the full resume is viewed, the recruiter can "Target" or "Save" the application. Recruiters often "Target" an applicant if the applicant is considered a suitable candidate for the posted job. "Save" implies that the recruiter has downloaded or saved the applicant's resume. If an application is either targeted or saved by the recruiter, we define it as a *Callback* response in the paper.

2.2 Main Statistics

Our master dataset represents a 5% random sample of job ads posted on the platform from 2018 to 2021, and all linked applications submitted to these job ads during the 90-day window before the ads expire. We exclude 3.54% of job ads that specify age limits.[8] The main analysis sample consists of 14,032,842 applications from 3,095,890 job seekers to 451,692 job ads, providing us with a snapshot of job vacancies and applications on this platform during the analysis period. As *Liepin.com* primarily focuses on the high-education, high-salaried market in major cities, our sample of job openings and applications on the platform is broadly more representative of China's labor market in developed urban areas, which make up a substantial segment of the online recruitment market.

Table 1 presents descriptive statistics on job advertisements. The average posted annual wage is 213,480 yuan ($\approx$ \$32,843 USD). Nearly half of the job listings are in the four Tier-1 cities in China – Beijing, Shanghai, Guangzhou, and Shenzhen. On average, a job posting receives 32 applications during the 90-day window. About 32% of job ads require applicants to have some postsecondary education, 58% require a Bachelor's degree, and 3% require a Master's or higher degree. For experience levels, 16%, 22%, 31%, 24%, and 7% of jobs ads require no experience, 1-2 years, 3 years, 4-5 years, and more than 5 years of experience, respectively. Nearly half of the jobs are listed by firms with over 1,000 employees, and most of them are from private-owned enterprises.

[8]Stating age requirement in job ads is not encouraged but still legal in China. Stating gender requirement in job ads was banned in China in March 2021, following the *Provisions on the Administration of Online Recruitment Services*. Before that, it was common for recruiters to set gender and age requests in job ads (Kuhn et al., 2020). On this platform, requirements on gender and age are only very rarely stated – 0.77% of job ads specify gender requirements, and 3.54% of job ads specify age requirements. Job ads indicating age requirements are excluded from the analysis.

Internet, real estate and automobile/manufacturing are the three most common industry sectors posting job ads on the platform.

Table 2 summarizes statistics of 14,032,842 applications in the main sample. This includes all applications submitted to the job ads, restricting to applicants aged between 25 and 55; applicants excluded from this age range only account for 4.9% of the data. We exclude applications above 55, because the normal pension age in China is 60 for men, 50 and 55 for blue-collar women and white-collar women. On the other hand, we keep applicants aged 25 and older, as they plausibly satisfy the most commonly stated requirements for education and experience, such as holding a Bachelor's degree and having three years of experience.

Around 63% of applications are sent by male applicants; and the average applicant age is 34. Applicants on this platform are highly educated: 61% of applications represent job seekers with a Bachelor's degree and 26% indicate a Master's or higher degree. About 17% of applicants are graduates from elite universities in China, known as Project 985 or Project 211 institutions. Almost 80% of applicants are admitted through the Unified-Exam Track and these degrees are better recognized. On average, applicants have 11 years of experience and a current wage of 266,220 yuan ($\approx$ \$40,957 USD). 34% of applicants are self-reported as unemployed, and 60% of them are in the four Tier-1 cities.

In terms of application behavior, about 52% of applications are submitted within four weeks after a job is initially posted. More than one quarter of applications are made using the "Batch Apply" function. We also create indicators to measure the match between the applicant and the job, assigning a value of 1 if the applicant's specified attributes meet the job's requirements or characteristics, and 0 otherwise. About 89% and 94% of applicants meet the job's education and experience requirements; and more than 70% of applications are submitted to jobs in the same city. In addition, we observe alignment between applicants' current job and their applied jobs: Almost 40% of applications are submitted for jobs in the same main industry, while 60% are for jobs in the same main occupation.[9]

As discussed above, our data includes two variables to measure application outcomes: *FullReview* and *Callback*. About 37% applications have their summary cards clicked in, leading

[9]The job's industry and occupation is categorized and encoded by the platform, which includes 12 (50) main-level (sub-level) industry categories and 22 (111) main-level (sub-level) occupation categories.

to a complete resume review by recruiters, and 9% of the applications receive a *Callback*.

3 Preliminary Evidence

We begin our analysis by examining how recruitment outcomes vary with applicant age over the complete age profile for all applicants from 25 to 55 years old. We estimate the following specification:

$$Y_{ijt} = \alpha + \sum_{k=26}^{k=55} \beta_k Age_i + \psi_j + \eta_t + \epsilon_{ijt}, \tag{1}$$

where Y_{ijt} is the recruitment outcome, i.e. *FullReview* (or *Callback*). It takes a value of 1 if the applicant i's application to job ad j on day t has been fully reviewed (or has elicited a callback after it is reviewed), and 0 otherwise. We regress recruitment outcomes on a full set of age dummies Age_i, and β_k captures the effect of each age from 26 to 55 on recruitment outcomes, relative to the reference group of applicants aged 25. Specification (1) estimates the raw disparities in hiring outcomes based on applicant age, without controlling for individual characteristics or human capital attributes. The estimated effect reflects the combined age effect we may observe in the real workplace, allowing for the possibility that work experience mitigates age bias against older workers.

Leveraging the large data set, we add job ad fixed effects ψ_j to control for the potential characteristics inherent to a posted job that may correlate with applicant age. This specification accounts for the differences in age compositions of applicant pools and recruiting criteria that may relate to applicant age for distinct jobs posted. Date fixed effects η_t are added to control for the time trend of application and recruitment over the analysis period.[10]

Our estimated results from Specification (1) are plotted in Figure 1, where *FullReview* is plotted with blue circles and *Callback* is plotted with green diamonds. Since these curves are estimated without accounting for productivity-related attributes, the observed age effects reflect both the influence of workers' accumulated experience and the potential impact of their age. Overall,

[10] In Section 5 and Appendix A.1.2, we conduct a quasi-experimental analysis controlling for individual fixed effects and our main effects remain similar.

FullReview curve initially shows a slight upward trend, indicating a quite modest increase over a brief age range from 26 to 30, which is then immediately followed by a predominant downward trajectory with a pronounced decline. The peak at age 30 suggests that, as workers move into their 30s, the negative effects of age bias start to outweigh the potential benefits of accumulated experience. Consequnetly, compared to 25-year-old applicants, individuals aged 38 and older are perceived as less employable, becoming significantly less likely to have their resumes reviewed by recruiters. For a 40-year-old applicant, the likelihood of receiving *FullReview* is 2.44 percentage points lower than that of a 25-year-old, representing a 6.59% decrease relative to the average viewing probability for 25-year-old applicants (37%). For a 50-year-old applicant, the likelihood drops by 8.12 percentage points, corresponding to a 21.92% reduction.

Next, we estimate and plot the age effects on *Callback*, focusing specifically on the subset of applications where the resume has already been reviewed.[11] The overall pattern of *Callback* remains similar: compared to 25-year-old applicants, whose average callback rate is 24.14% if their resume has been reviewed, the callback rate decreases by 4.21 percentage points (17.44%) for 40-year-old workers and by 9.13 percentage points (36.54%) for 50-year-old workers. The larger coefficient estimates for *Callback* and the greater magnitudes – amplified by the smaller reviewed sample – suggest that a significant age-related bias still remains even after resumes are fully reviewed.

4 Empirical Analysis

4.1 Main Specification and Results

Notably, the estimated age profile in Section 3 is clearly indicative of an age-based disparity in the recruitment process, but one major concern is that 25-year-old applicants may not serve as an ideal reference for comparing older applicants up to age 55. When there is a substantial age difference, educational background, work experience, and family responsibilities can vary significantly across age cohorts. To make workers with age differentials more comparable, we

[11]To isolate the impact on *Callback* from the effects of *FullReview*, we focus on the subset of applications where the resume has already been reviewed. In Figure 3, we estimate the effects using the full sample, without conditioning on whether an application has been fully reviewed.

estimate the following specification for 30 pairs of applicants separately:

$$Y_{ijt} = \alpha_r + \beta_r Old_{i,r} + X_{ijt,m}\rho_m + \psi_j + \eta_t + \epsilon_{ijt}, \tag{2}$$

where r denotes the omitted reference age group in each pair, and $Old_{i,r}$ is an indicator that applicant i is the older individual in the age pair corresponding to r (i.e. age is $r+1$). The 30 pairs of applications compare applicants aged 26 to 25 through 55 to 54, where the "older" applicants are only one year older than the reference group r. Thus, specification (2) identifies age-based bias among highly comparable individuals within a given job posting and enables us to assess the varying degrees of age bias across different age ranges.

$X_{ijt,m}$ represents a vector of variables capturing applicant characteristics, depending on the amount of information presented to the recruiter, where $m \in \{b, f\}$. When recruiters are to decide whether to fully review an application, only basic information on the summary card is displayed to the recruiter (i.e. $m = b$), which includes the applicant's gender, age, location, educational level, years of experience, job title and the current employer. The brevity of summary card information enables us to capture nearly all attributes that account for recruiters' *FullReview* decisions. $X_{ijt,b}$ includes a gender dummy, educational level dummies, a quadratic term in years of experience, and match indicators for location, industry, and occupation alignment between the applicant and the job.[12]

Callbacks are made only if applicants have received full reviews previously, whose complete resume is presented to recruiters (i.e. $m = f$). Thus, in addition to the variables included earlier, $X_{ijt,f}$ incorporates indicators for elite education, industry and occupation match for the applicant's last two jobs, and their current wage. We also include days elapsed between posting and application, and a vector of applicants' search status dummies (unemployed searching, employed and actively searching, employed and casually browsing, and employed and browsing with no intent to switch jobs) to account for potential behavioral impacts.

The estimation results on *FullReview* and *Callback* are plotted in Figure 2(a), (b). The patterns in (a) and (b) are similar: Despite the positive age effect observed before age 28, among applicants

[12] We do not observe applicants' current employer, so we include an industry match indicator to capture the main message that can be implied. In Section 5 and Appendix A.2, we have conducted several robustness practices and our results are similar to the main results.

aged 30 to 50, we consistently detect age bias against individuals who are just one year older, even within this closely comparable group, affecting both *FullReview* and *Callback* outcomes. For applicants aged 31 to 50, each additional year of age is associated with an average decrease of 0.51 percentage points in the likelihood of receiving a *FullReview*. Given that 35% of applications in this age range receive a *FullReview* on average, this decrease translates to a 1.47% reduction in the likelihood of receiving reviews as applicants age by one year. The average age effect estimated for *Callback* is slightly smaller – 0.37 percentage points – representing a 1.66% decline in the likelihood of receiving callbacks with each additional year of age. The magnitude is amplified because the coefficients on *Callback* are calculated from a subset of applicants whose resumes have already been reviewed previously.

We highlight three interesting patterns from the above analysis. First, on the *FullReview* curve, the estimates for ages 31 to 39 and 41 to 49 show a relatively steady negative effect, while the estimates at ages 40 and 50 exhibit a sharp and more pronounced decline compared to the preceding ages. This sudden drop suggests that recruiters may use round numbers as thresholds to filter out older applicants when deciding whose resumes to review. An applicant aged 40 is 1.19 percentage points less likely to be reviewed than a comparable applicant aged 39, whose likelihood of being reviewed is 37%, representing a 3.22% decline. Similarly, for a 50-year-old applicant, the 1.41 percentage points decrease corresponds to a 4.70% reduction compared to a 49-year-old applicant. These estimates indicate that age bias near round-number ages is double or even triple the magnitude of age bias observed at other ages.

Second, as shown in Figure 2, there is a reduction in the coefficient estimates from *FullReview* to *Callback*. Related to the point discussed previously, the sudden drops at round numbers are only found on the *FullReview* curve but not on the *Callback* curve. On temping explanation is that age bias in the review process is mostly driven by statistical discrimination. According to statistical discrimination theory (Aigner and Cain, 1977), providing more information about a candidate's productivity can reduce employers' reliance on group characteristics as a selection criterion. When recruiters have access to complete resumes rather than just summary cards, they are likely to have an enhanced assessment of candidates' competencies, leading to smaller disparities in callback rates compared to the differences observed in reviews. On the other hand, however, these two curves cannot be directly contrasted, because callbacks are conditional on applications having already undergone reviews. In other words, it is possible that smaller disparities are detected in the callback process because the most discriminated applicants are already excluded in the previous review process. An ideal experiment design to measure the extent of statistical discrimination would involve presenting only the summary cards to the control group of recruiters, while providing the treatment group with complete resume.

Third, to align our results with previous literature focused on callback outcomes – where the summary card stage is not present in correspondence studies – we re-estimate the age effects in specifications (1) and (2) for *Unconditional Callback*, using the full sample of applications regardless of whether they had been reviewed previously. Similar patterns are observed in Figure 3: As shown in panel (a), the complete age profile for *Unconditional Callback* is flatter but still shows a decline starting at age 30, with applicants in their late 30s receiving fewer callbacks compared to a 25-year-old applying for the same job. The coefficients plotted in panel (b) align with our primary findings: applicants in their mid-30s experience a decline in hiring success with each additional year of age. For applicants aged 31 to 50, each additional year reduces their unconditional callback success by 0.23 percentage points, reflecting a 2.47% decline.

Our estimates indicate that, cumulatively, a 50-year-old applicant is approximately 50% less likely to receive a callback compared to a 30-year-old applicant. This substantial effect is partly driven by the relatively low callback rate in our empirical context (raw callback rate is 9.33% in the entire application sample), as the high volume of applications may amplify the magnitude of the estimations. Compared to the existing literature, our estimate is relatively on the high side for correspondence studies based on the US labor market but falls around the middle for studies conducted in European labor markets. For similar age ranges, Neumark et al. (2019) and Farber et al. (2019) report discrimination gaps of approximately 15-18% in the US labor market, while Lahey (2008) identifies a larger gap of about 40% for older women. In contrast, larger age discrimination gaps have been reported in studies focusing on European contexts (Riach and Rich, 2006, 2010; Ahmed et al., 2012; Carlsson and Eriksson, 2019).

As callbacks are made conditional on applications having been previously reviewed, our paper hereafter focuses on *FullReviews*, where the applicant's age first appears, and when we observe nearly the same information as recruiters. Nevertheless, the consistent pattern in Figures 1, 2 and 3 indicate that a pronounced age-based bias still persists even after applicants' full resume has been reviewed.

4.2 Gender

As discussed above, age presents a common challenge for workers in their 30s and beyond, as older applicants often face reduced opportunities for their resumes to be reviewed or to receive callbacks. However, the effect of age can differ between men and women. Previous research has documented a pronounced age-gender interaction, showing that aging tends to have more pronounced negative effects on women than on men (Carlsson and Eriksson, 2019; Cortés and Pan, 2023; Lahey and Oxley, 2018), or that it affects women and men differently across various age ranges (Helleseter et al., 2020). In particular, women may encounter a dual penalty linked to both age and parenthood during middle age, further compounding their disadvantages in the labor market. To examine these gender-specific differences in age-related hiring outcomes, we disaggregate our main analysis on *FullReviews* by gender.

Figure 4 demonstrates that women experience an age penalty earlier than their male counterparts. Panel (a) shows that younger women (below their mid-30s) experience smaller positive effects in receiving full resume views. Without accounting for human capital attributes, these estimates capture both the potential benefits of experience accumulation and the effects of age. The smaller age premium observed for younger women suggest that female workers either face greater age bias even at a young age, or their human capital is valued less than that of their male counterparts. As shown in panel (b), the negative impact of aging on the likelihood of receiving a full review begins around ages 29-31 for females and 31-33 for males. Beyond these points, age penalties tend to be slightly larger for women, particularly in their mid-to-late 30s. This pattern may be linked to family responsibilities and the "child penalty" (Petit, 2007; He et al., 2023), which disproportionately affects women of childbearing age. Supporting this hypothesis, our estimates show that beyond the typical fertility ages in China (above 40), the gap in age-related penalties between male and female workers largely disappears.

5 Robustness

In this section, we conduct and discuss robustness checks on our main results regarding the one-year age effect on *FullReview*, with detailed regression results provided in Appendix A. First, there may be concerns that the productivity proxies controlled for in our analysis do not fully capture variations in applicants' perceived productivity. To address this, we conduct two quasi-experimental analyses.

In the first analysis, we leverage the limited information available on the summary card to construct a sample of applicants whose summary card details -except for age- are identical within each job posting, and we re-estimate the effect of one-year age increase on *FullReviews*. As shown in Appendix A.1.1, even when comparing among applicants with identical summary card information, those who are just one year older are less likely to be reviewed after age 30. The magnitude of this effect is similar to that observed in Figure 2(a). Moreover, the round-number effect becomes even more pronounced. Applicants turning 40 (50) face a penalty more than three times greater compared to the transition from age 38 to 39 (48 to 49).

In the second quasi-experimental design, we utilize the natural age progression of the same applicant to analyze job search outcomes, using worker fixed effects. Since an applicant's exact birth date is not directly observable, we first identify the initial instance in our data when applicants become one year older, and then include all applications submitted within a 30-day window surrounding this date to estimate the impact of being one-year older on job search success. We replace the job-ad and date fixed effects used in the main analysis with individual and year fixed effects. Results in Appendix A.1.2(a) show that turning one year older decreases the likelihood of an application being viewed compared to those submitted by the same applicant within the previous 31-60 days, when the summary card displayed the applicant as one year younger. Moreover, the negative age effect starts earlier and persists across all ages.[13]

Second, a potential issue in our data is the absence of information about an applicant's current employer on the summary card and resume, while recruiters can access this information and

[13]The estimated magnitudes are also larger in this specification, likely due to sample selection bias rather than differences in the estimation approach. Individual fixed effects exclude observations where an applicant is only observed at a single age within each two-year pair. This leads to the disproportionate exclusion of older applicants, who typically apply less frequently. As a result, the sample becomes skewed toward younger applicants and applicants who apply more often. When we replicate specification (2) (i.e. using job and date fixed effects) on the same sample in Appendix A.1.2(b), the patterns and magnitudes remain highly similar.

may use it to evaluate the applicant's productivity or human capital accumulation. In the main specification, we have controlled for the industry and occupation match indicators, which is the most important information can be inferred from an applicant's current employer, and we believe this should have largely alleviated the concern. In Appendix A.2, we have conducted two additional empirical tests to further assess the robustness of our results. First, when only summary card is presented, recruiters may use an applicant's current employer to infer the applicant's human capital attributes such as working history, current wage level, and actual intention to depart from the current employer. To test whether such inferred information may affect the review outcomes, we include these additional variables in the *FullReview* regression – before they actually become available to recruiters – to re-examine our review findings. As shown in Appendix A.2.1, the results closely align with our main findings. This suggests that these attributes are either rarely or cannot be accurately inferred, or that they do not significantly influence the likelihood of a resume being reviewed beyond the information already available on the summary card. In the second robustness test presented in Appendix A.2.2, we only keep applications if the applicant's current employer is *not* in the same industry as the recruiter, i.e. industry match indicator = 0. Presumably, when a manufacturing firm is hiring and has received an applicant currently working in the construction sector, the recruiter's knowledge about the applicant's current employer is likely limited – mirroring our scenario where the current employer is not observed. Again, the results are highly consistent with our main findings, indicating that our analysis is not sensitive to not observing an applicant's current employer.

Third, younger and older applicants may have different job search strategies, targeting positions that maximize their chances of success. Moreover, recruiters may also use ageist language in job postings to deter older applicants from applying (Firoozi et al., forthcoming). We are not too worried about this selection issue, because our analysis aims at identifying age bias in recruiters' evaluation of applicants, taking applicants' job search strategy as given. Hence, allowing optimal search strategy – rather than mandating younger and older applicants submit to the same jobs, as is common in correspondence studies– enables us to detect age bias *at the margin*, making our analysis more reflective of real-world job markets. Nevertheless, in Appendix A.3, we perform a robustness exercise to test if the observed age bias is sensitive to applicants' endogenous job search strategy. To do so, we focus on Batch applications, which are submitted to multiple job positions that are recommended by the platform with a single click. Since applicants cannot thoroughly review the detailed requirements and descriptions for these jobs when using the Batch function, these applications are less prone to job seeker's selection concern. We find very similar results, reassuring us that our findings are not driven by applicants' job search strategy.

Another concern is that older applicants may submit their applications later in time, so they receive fewer reviews or callbacks because the posted vacancy has already been filled. To account for this possibility, we control for job-ad by elapsed duration fixed effects, allowing the job-ad identifier to interact with six time dummies indicating when the application was submitted: 1, 2, 3, or 4 weeks, 5-8 weeks, and 9-13 weeks after the job posting date. Relatedly, this practice also helps to address the concern that the recruiter's perception of an applicant's age may be relatively influenced by the other applications received during the same time period. The results, presented in Appendix A.4, are again quite similar to our main findings.

Lastly, there is a potential concern that the pandemic period might have influenced our results, as job vacancies and job search behaviors could have changed during this time. To address this, we split our sample into two periods: the pre-COVID period, covering all jobs posted from 2018 to 2019, and the COVID period, covering jobs posted in 2020 and 2021. As shown in Appendix A.5, our main results remain consistent in both periods.

6 Potential Mechanisms and Heterogeneity

In this section, we explore potential explanations for the age bias observed in our analysis. We first consider factors documented in previous literature regarding the reasons for age discrimination. The most predominant explanation for age discrimination is concerns about the physical well-being and health conditions associated with older individuals, such as stereotypes about their diminished physical capabilities, poorer health outcomes, and reduced speed (Kroon et al., 2018; Burn et al., 2019). Secondly, older workers often have or are believed to have lower education levels or less up-to-date skill training compared to younger cohorts, further contributing to age discrimination.

However, these factors can hardly apply to our analysis, as we focus on hiring bias within each 2-year cohort, where there is plausibly no significant (perceived) difference in physical health conditions and educational background. While it is challenging to pinpoint the exact causes of age discrimination, we investigate potential mechanisms associated with the estimated age bias. Specifically, we explore how the effect of age on the likelihood of an application being viewed varies with job requirements, labor market conditions, industry sectors, occupation types, employer characteristics, and worker characteristics. To do so, we analyze the age effect on *FullReview* across various subsamples, excluding individuals over 50 due to insufficient sample size.

The following part of this section focuses on two key aspects: First, we examine the evolution of the age effect over time, as estimated in our main analysis, across various heterogeneous subgroups. Detailed results for each subsample are provided in Appendix B. Second, to obtain an overall estimate, we calculate an average effect of age using a meta-analysis approach, treating each regression estimated from the 2-year cohort as an individual study.[14] The overall age effects are presented in Figure 5.

6.1 Job Requirements

Experience Requirement: From the perspective of human capital accumulation, the evidence of older workers to be less favored in the labor market suggests that their work experience is undervalued. This may reflect various degrees of age discrimination based on the job's required experience level. We analyze this by categorizing job listings into four groups based on experience requirements: up to 1 year, 1 to 3 years (inclusive), 3 to 5 years (inclusive), and more than 5 years.

Our findings in Appendix B.1.1 show that jobs requiring longer work experience tend to exhibit age discrimination at a later stage. Specifically, jobs require no more than 1 year or 1-3 years of experience show a higher degree of age bias, with applicants in their late 20s start to experience significant age bias. For jobs that require 3-5 years of experience, we observe a positive age effect among applicants in their late 20s, and age begins to be perceived as a disadvantage once applicants reach their mid-30s. A similar pattern is observed for jobs requiring more than 5 years of experience, but the negative age effect is delayed until after the age of 35.

Overall, as shown in Figure 5(a), for all applicants aged 26 to 50, average age penalties are

[14]The average effects of being one year older on *FullReview* are derived using a meta-analysis approach (Glass, 1976; Borenstein et al., 2021), which calculates a weighted average of the estimated effect of age across two-year cohorts within the age range of 26 to 50 in Figure 2.

more pronounced in jobs requiring less experience, while positions demanding greater experience exhibit a weaker effect.

Management Role: Jobs hiring for managerial positions may exhibit less age bias, as employers are more likely to value work experience for these roles. To examine this, we distinguish between management roles – defined as positions involving supervision of subordinates – and non-management roles in Appendix B.1.2. Although estimates for management roles are noisier due to a smaller sample, the results still clearly indicate that age bias in these roles tends to be postponed until applicants reach their late 30s.

Education Requirement: Next, we disaggregated the main results by the education levels specified in job postings. As detailed in Appendix B.1.3, the negative impact of age on review probability is primarily concentrated in positions requiring an associate's or bachelor's degree. In contrast, for jobs requiring more advanced degrees (Master's, PhD, or MBA), the estimated age effects are mostly smaller and not significant; the insignificant estimates may also be due to the smaller number of observations for these higher-education postings.

Skill Requirement: Another possible explanation for older workers being disadvantaged compared to younger ones could be their lack of familiarity with new technology and skills, such as computer proficiency (Burn et al., 2022). However, this factor alone cannot fully explain the age discrimination observed in our study, as a substantial proportion (over 50%) of the workers are employed in technology and internet-related positions. Nonetheless, it is important to identify specific skills associated with more pronounced age discrimination. We classify key skills from sub-level occupation codes using the skills framework developed by Deming (2017), which categorizes skills into three primary groups: routine skills, non-routine skills, and interpersonal skills. For example, for Data Analyst, non-routine skills emerge as the predominant requirement. While the overall effects remain similar, as shown in Figure 5, the detailed trends in Appendix B.1.4 suggest that jobs demanding higher levels of interpersonal skills exhibit age bias the earliest, beginning in the early 30s. In contrast, positions requiring more routine and non-routine skills tend to show age bias a bit later, when applicants reach their mid-30s. This pattern may be caused by customer-based discrimination against older workers or employers' perception that older applicants have weaker social skills (Richardson et al., 2013).

6.2 Labor Market Conditions

Labor Market Tightness: Previous research suggests that labor market conditions significantly influence employers' hiring strategies, with discrimination often varying based on labor market tightness (Dahl and Knepper, 2023). To examine this, we measure market tightness using the ratio of job openings to job seekers within each city-industry-occupation cell and categorize markets into tight, moderate, and slack based on tertiles. While the overall effects remain similar in Figure 5, the detailed patterns in Appendix B.2.1 indicate that age bias occurs later in tighter labor markets than in slacker ones (34 versus 31), consistent with prior research showing that discrimination tends to be lower in markets where hiring is more challenging.

Job-level Competitiveness: Similarly, when the applicant pool is large, employers may be more selective and exhibit increased discrimination. We measure job competition by the total number of applications each job receives, categorizing our sample into three groups: highly competitive jobs (more than 50 applications), moderately competitive jobs (30-50 applications), and less competitive jobs (fewer than 30 applications). As shown in Appendix B.2.2, while standard errors are larger for jobs with fewer applications – making their estimates statistically insignificant –the coefficient estimates for ages 30-40 are in fact quite similar among the three groups.

6.3 Employer Characteristics

Coworker Age: Another aspect of discrimination may come from coworker compositions. In particular, firms with a younger workforce may prefer hiring younger employees due to the prevalence of younger peers. To examine this, we classify younger and older firms based on the average age of applicants applying to *other* positions the firm has posted. Our findings in Appendix B.3.1 suggests that firms with younger workforces are more likely to favor younger candidates, with the likelihood of being considered declining for applicants as early as their late 20s.

Firm Size, Firm Ownership, and Firm Location: In Appendix B.3.2 to B.3.4, we examine the relationship between age bias and employer's firm size, ownership type, and geographic location. Overall, age-related bias is consistently present across different types of employers.

Regarding firm size, smaller firms (fewer than 100 employees) exhibit less age bias, which may be due to receiving fewer applications. Meanwhile, firms with 100-1,000, 1,000-5,000, and more than 5,000 employees display similar levels of age discrimination.

In terms of firm ownership, the estimates remain relatively consistent across all age groups. However, age bias appears to emerge slightly later in foreign-invested firms and publicly listed enterprises, compared to privately owned and state-owned enterprises.

Geographic location also plays a role. In tier-1 cities – Beijing, Shanghai, Guangzhou, and Shenzhen – age-related penalties in resume reviews begin as early as applicants enter their early 30s. In contrast, in other cities, the penalty typically emerges after age 35.

6.4 Industry Sector and Occupation Type

Appendix B.4.1 illustrate the average effects of age from 26 to 50 across 12 main industry categories, based on the platform's industry classification. Overall, we observe a consistent negative effect of aging on *FullReviews*, with the impact being particularly pronounced in the real estate and services industries, consistent with previous findings that jobs requiring higher interpersonal skills tend to exhibit greater age bias in hiring. Appendix B.4.2 presents the average effects across the most represented 10 occupations in our data. Most occupations exhibit negative age effects, and the impact is especially strong in occupations related to retail, HR/admin, and financial services.

6.5 Worker Characteristics

Education Level: We further examine whether specific worker attributes can mitigate potential age-related bias. Since educational attainment is considered a key signal of a worker's productivity, we first assess whether holding a Master's or higher degree can alleviate age-related bias. As shown in Figure B.5.1, having an advanced degree slightly delays the onset of age-related bias, shifting it from the early 30s to around 35. However, beyond age 35, the effects converge. The overall findings suggest that pursuing graduate education does not significantly reduce the likelihood of experiencing age-related bias.

Next, we examine how additional worker attributes influence age bias by considering information available once an applicant's resume has been fully reviewed, including education recognition, employment status, and current wage. Accordingly, we assess their effects on

Callbacks.

Elite Institution: In addition to education level, graduating from an elite institution may also influence the extent of age bias a worker experiences. This is particularly relevant in our context, where elite education is highly valued in Asian cultures and is believed to enhance perceived human capital through stronger networking effects. Elite-educated workers are defined as those graduating from Project-985 or Project-211 universities. As shown in Figure B.5.2, while most age bias estimates for elite-educated applicants are statistically insignificant, their magnitudes are similar to those for non-elite-educated applicants. This suggests that graduating from an elite institution does not effectively mitigate potential age-related bias.

Employment Status: When an applicant's resume is reviewed, their employment status is visible to recruiters. Unemployment can influence age bias in both directions. On one hand, it may exacerbate age bias, as prolonged unemployment could signal human capital depreciation. On the other hand, it might suggest a lower outside option, potentially making the applicant more attractive to employers and increasing the likelihood of callbacks. As shown in Appendix B.5.3, unemployed and employed workers experience a similar level of age bias.

Worker Wage: Lastly, workers' wages can have a mixed impact on age bias. Higher wages can signal greater labor market value and higher productivity, potentially reducing age bias if recruiters prioritize highly productive candidates. However, they may also intensify age bias if employers perceive higher wages as a proxy for greater reservation wages and increased hiring costs. As a result, hiring younger, lower-cost workers may be more financially appealing. To test this hypothesis, we examine age discrimination among higher- and lower-wage workers. A worker is classified as high-wage if their current salary exceeds the median of the applied job's posted wage range, with approximately 40% of applicants falling into this category. The findings in Appendix B.5.4 suggest that higher wages may accelerate the onset of age bias – from age 36 to as early as age 30. While the estimates are not statistically significant, we observe a consistent pattern of negative coefficients for high-wage workers beginning at age 30, compared to negative effects emerging at age 36 for lower-wage workers.

7 Conclusion

This paper studies age-based hiring bias through an extensive analysis of 14 million application records to over 451,692 job advertisements on a leading Chinese online job platform. The findings provide robust evidence of age discrimination *at the margin*, where workers self-select into jobs they apply for and incur an additional penalty with each year of aging. To the best of our knowledge, this is the first study to measure "middle-ageism" in the labor market. Our empirical findings reveal a striking reality: Applicants entering 30s start to experience a subtle yet consistent decline in hiring success rates, with this trend intensifying significantly for those aged 35 to 50. Termed "middle-ageism", this phenomenon highlights the systemic barriers faced by mid-career applicants and underscores a broader issue of ageism that permeates the labor market. Moreover, this bias does not affect all groups equally, as female applicants bear a disproportionately greater burden of discrimination compared to their male counterparts.

While discussions on age discrimination often focus on issues related to retirement benefits or social security, our findings reveal that the marginal effects of aging occur much earlier than these regulatory or social benefits thresholds suggest. Workers may face diminishing opportunities even during their prime working years, a phenomenon not adequately explained by factors in previous literature such as declining proficiency with new technologies or decreased physical capacity. Notably, in slack labor markets or highly competitive job roles, there is evidence that employers may use age as an initial screening criterion, filtering out older applicants based on even a one-year age difference. This occurs at the very first glance of an application, often without a thorough review of the applicant's full resume.

Our paper introduces several refinements to prior studies. Firstly, by acknowledging applicants' strategic approaches to job searching, this paper offers nuanced evidence of marginal age discrimination, thereby enriching the literature on discrimination with insights that extend beyond the conventional scope of average discrimination typically highlighted in correspondence studies. Secondly, by leveraging the job platform's standardized process for presenting applicant information, we ensure that the human capital attributes we observe closely match the information available to recruiters, thereby significantly reducing omitted variable bias. Moreover, our paper expands the focus of correspondence studies due to constraints in generating and sending fictitious studies. By broadening the spectrum of investigation to include a wider range of industry sectors, professions and seniority levels, our research provides a more comprehensive understanding of the prevalence and varying degrees of age-based discrimination across labor markets.

Meanwhile, we acknowledge the limitations in our study that future research could address. First, a primary question in the discrimination literature is to distinguish statistical discrimination from taste-based discrimination. While this remains a complex challenge, future papers may consider manipulating the amount of information disclosed to recruiters to help address this question. Second, our focus on those aged 25-55 offers only limited explanatory power for applicants over 50, leaving room for further investigation into this age group. In addition, due to the absence of data on characteristics of recruiting agents, we are not able to examine whether the observed discrimination stems from the personal biases of recruiters or from the discriminatory policies of the recruiting enterprises. Future research aimed at pinpointing the specific sources of discrimination could yield valuable policy implications.

Figures and Tables

Figure 1: Hiring Outcomes Across the Complete Age Profile

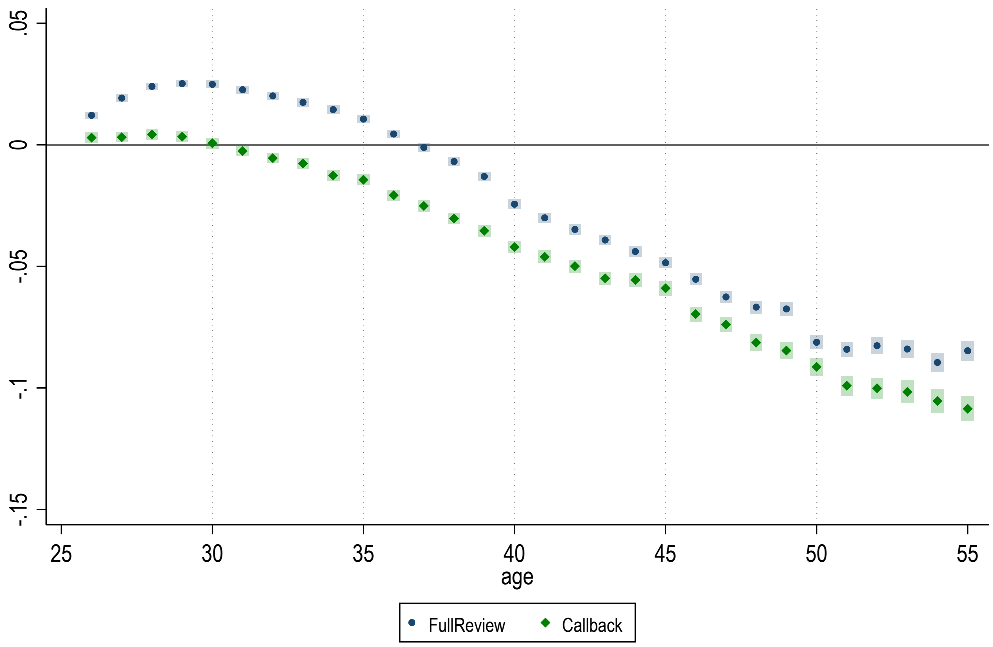

Notes: This figure presents recruitment outcomes *FullReview* and *Callback* for applicants of different ages relative to the reference group of 25-year-olds, estimated using Specification (1). The plotted coefficients for age dummies from 25 to 55 include 95% confidence intervals. The regression controls for job-ad and application date fixed effects, with robust standard errors clustered at the job-ad level.

Figure 2: Hiring Outcomes by One-Year Age Increments

(a) FullReview

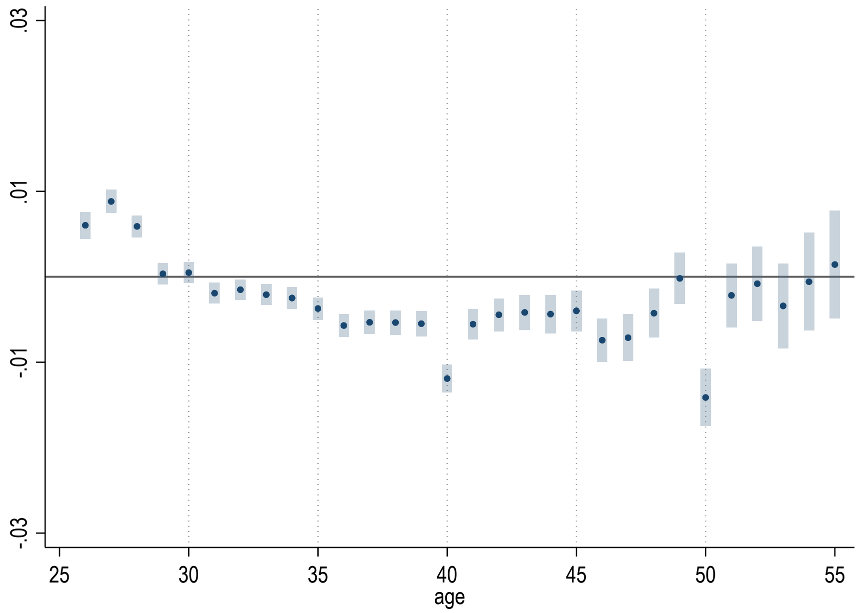

(b) Callback

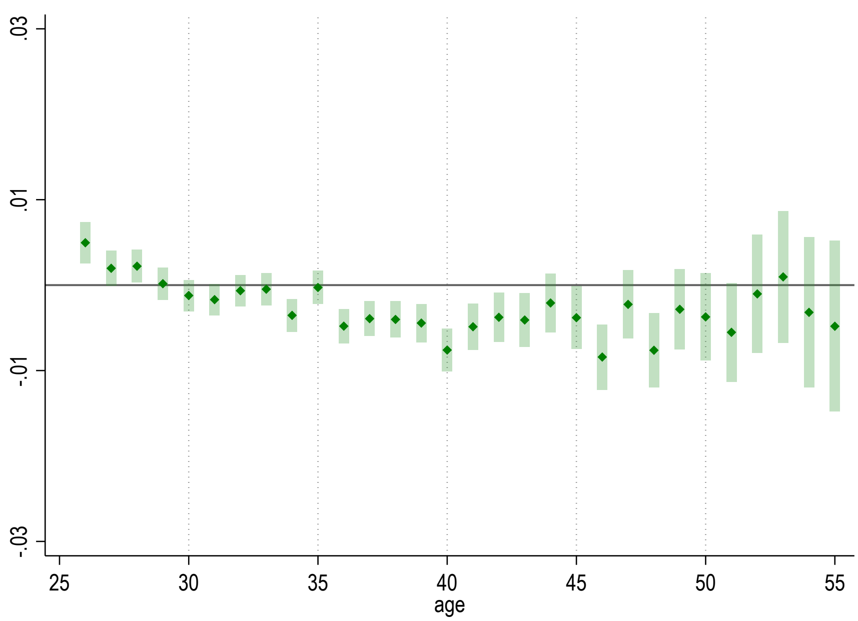

Notes: These figures presents the marginal effect of a one-year age difference on hiring outcomes, estimated using Specification (2). Results on *FullReview* and Callback are presented in (a) and (b), correspondingly. Each point represents a separate regression comparing applicants aged $r + 1$ to the reference group r, ranging from applicants aged 26 vs 25 to those aged 55 vs 54. The regression controls for job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.

Figure 3: Unconditional Callbacks by Age

(a) Unconditional Callbacks Across the Complete Age Profile

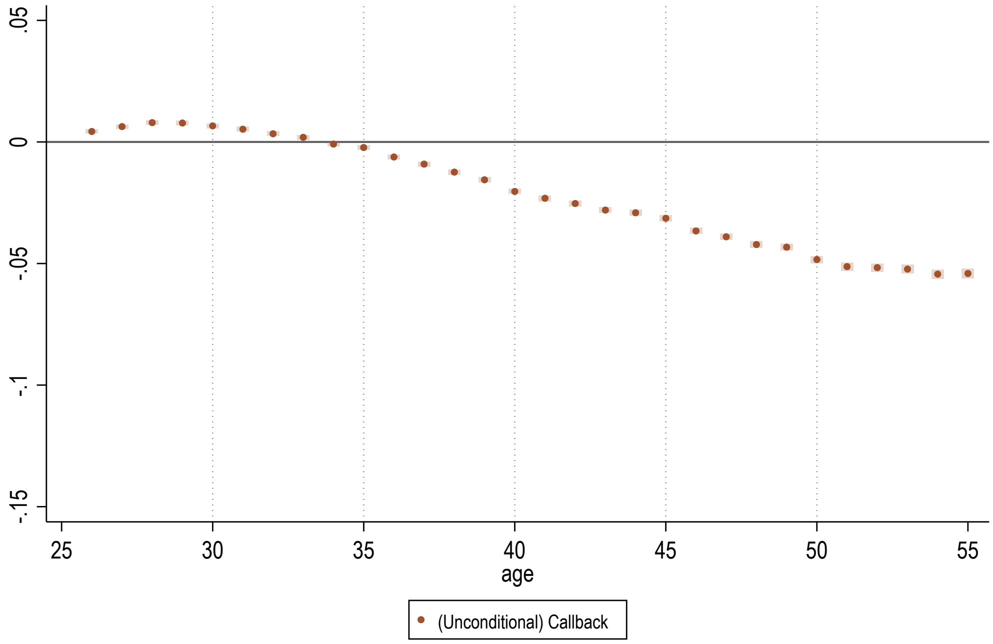

(b) Unconditional Callbacks by One-Year Age Increments

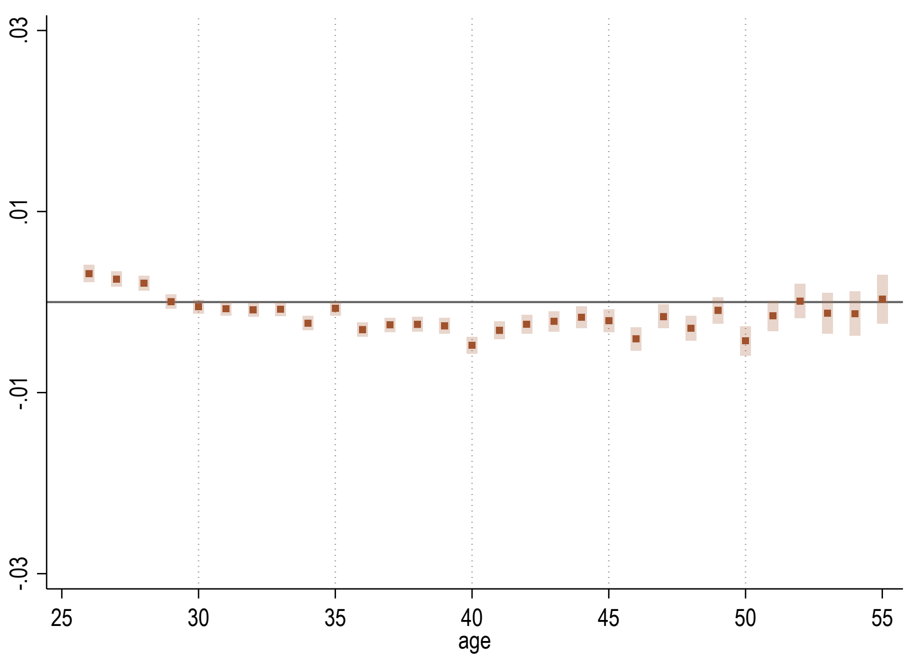

Notes: Figure 3(a) and (b) replicate results in Figure 1 and Figure, respectively, using the full sample of applications, irrespective of whether the applications had been reviewed previously. Regression includes job-ad and date fixed effects; robust standard errors are clustered at job-ad level.

Figure 4: *FullReview* by Age, by Gender

(a) *FullReview* Across the Complete Age Profile, by Gender

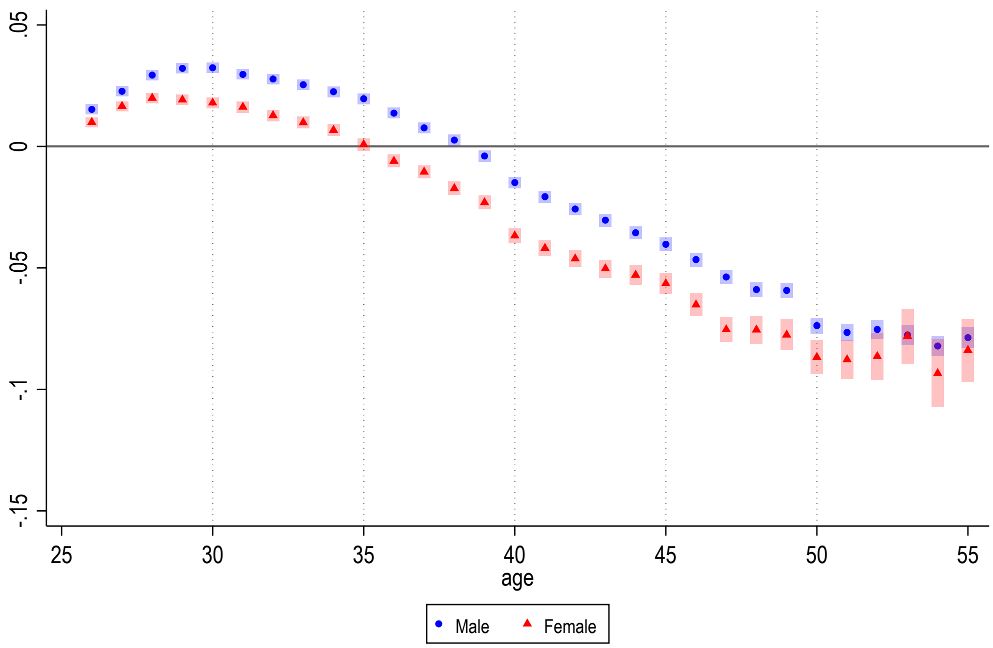

(b) *FullReview* by One-Year Age Increments, by Gender

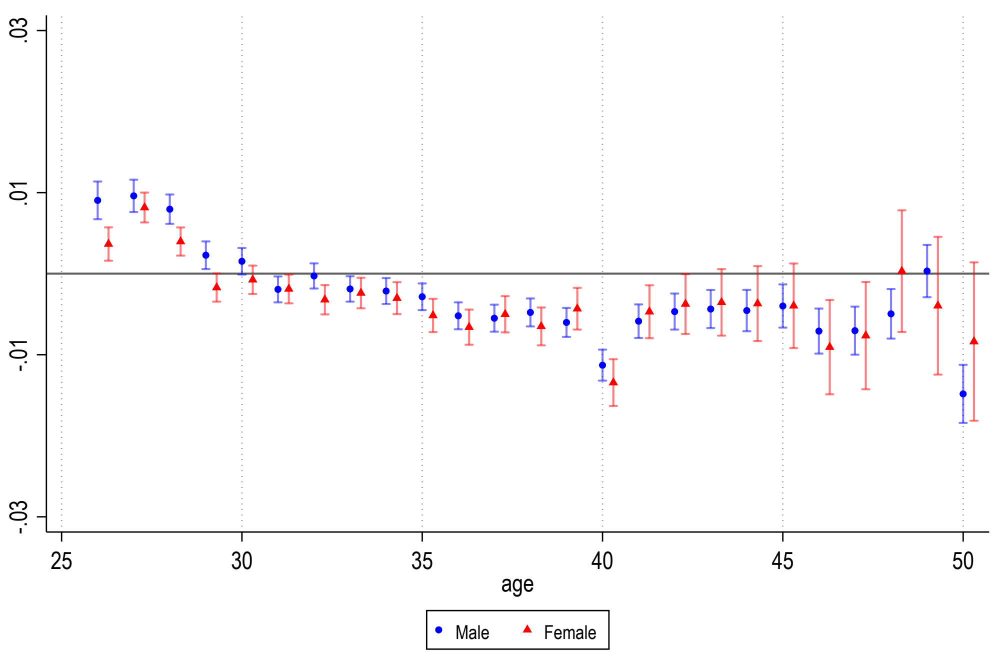

Notes: Figure 4(a) and 4(b) replicate the results from Figure 1 and Figure 2(a), respectively, separately for male and female applicants. The regressions include job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.

Figure 5: Heterogeneity by Job's and Employer's Characteristics, and Labor Market Conditions

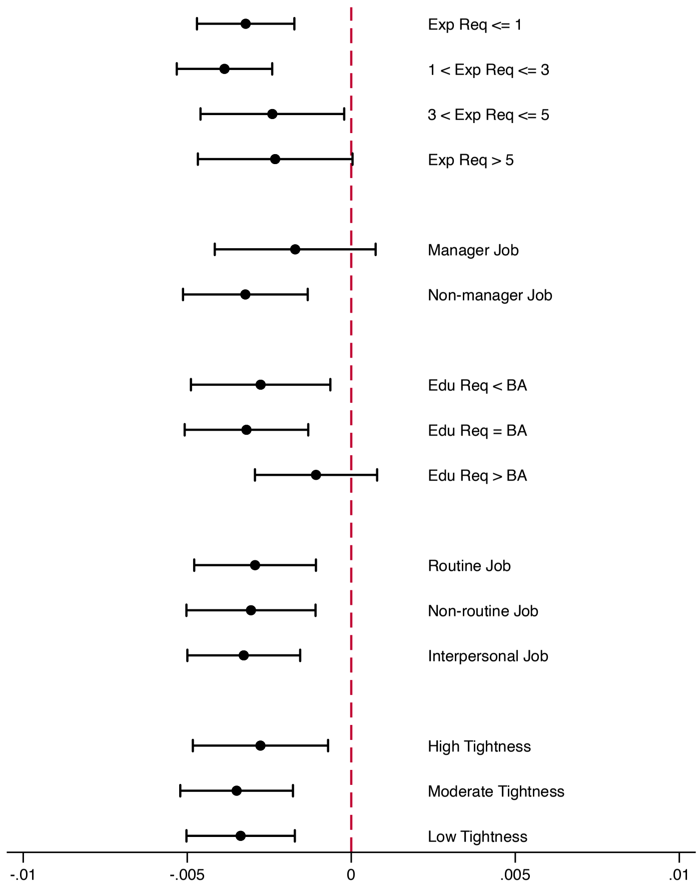

(a)

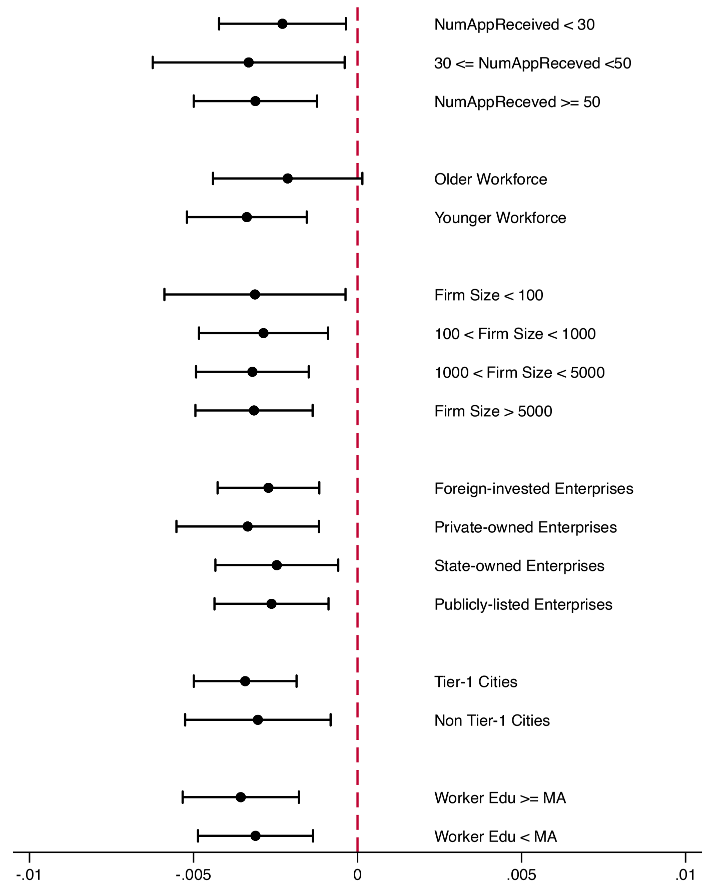

(b)

Notes: This figure plots the average coefficient estimates for applicants as they age by one year, for those aged 26-50. Estimates are calculated using a meta-analysis approach, treating each coefficient estimate obtained from (2) as an individual study. Standard errors are computed using the bootstrap method. Detailed coefficient estimates for each age are presented in Appendix B.

Table 1: Summary Statistics: Job ad sample (N = 451,692)

	Mean (S.D.)
Posted job ad characteristics	
Posted wage (annual, in 1,000 RMB)	213.48 (167.58)
Job in Tier-1 cities	49.92%
Number of applications received	32.28 (86.43)
Education requirement	
Below Bachelor's degree	38.53%
Bachelor's degree	58.43%
Master's or higher degree	3.04%
Experience requirement	
Up to 1 year	32.94%
1-3 years of experience	36.67%
4-5 years of experience	23.77%
$>$ 5 years of experience	6.58%
Years of experience requested	3.00 (2.40)
Employer characteristics	
Size of enterprise	
$<$ 100	14.41%
100 - 999	38.92%
1000 - 5000	21.92%
$>$ 5000	24.74%
Ownership of enterprise	
Private-owned enterprises (POE)	70.63%
Foreign-invested enterprises (FIE)	15.79%
State-owned enterprises (SOE)	6.23%
Publicly-listed enterprises	7.35%
Top three industries	
Internet, games, and software	29.19%
Real Estate, construction, property management	13.57%
Automobiles, machinery, manufacturing	8.68%

Notes: This table summarizes all job ads included in the main analysis, excluding the 3.54% of postings that explicitly impose age limits. The posted wage is calculated as the average of the annual salary range advertised in the job listing. Tier-1 cities include Beijing, Shanghai, Guangzhou, and Shenzhen. Industries are categorized based on the platform's primary industry classification, which consists of 12 main industry categories.

Table 2: Summary Statistics: Job application sample (N = 14,032,842)

	Mean (S.D.)
Applicant characteristics	
Male	62.85%
Age	34.38 (6.47)
Education level	
Some postsecondary education	13.02%
Bachelor's degree	60.72%
Master's or higher degree	26.26%
Education recognition	
Project 985/211 institutions	17.23%
Unified-Exam Track	78.49%
Years of experience	11.29 (7.27)
Current wage (annual, in 1,000 RMB)	266.22 (236.45)
Unemployed	33.61%
Applicant in Tier-1 cities	60.25%
Application behaviors and match	
Days elapsed since job is posted	
Week 1	17.73%
Week 2	14.97%
Week 3	11.17%
Week 4	8.69%
Weeks 5-8	24.66%
Weeks 9-13	22.78%
Batch application	26.94%
Education match	88.78%
Experience match	93.70%
Province and city match	71.32%
Main-level industry match	39.93%
Main-level occupation match	59.96%
Application results	
FullReview	36.72%
Callback	9.33%

Notes: This table summarizes main statistics from all applications included in the main analysis. Batch applications represent applications submitted to multiple recommended ads in one click. Match indicators are binary variables, taking a value of 1 if applicant's characteristics are aligned with the job requirements.

References

Ahmed, A. M., L. Andersson, and M. Hammarstedt (2012). Does age matter for employability? a field experiment on ageism in the swedish labour market. *Applied Economics Letters 19*(4), 403–406.

Aigner, D. J. and G. G. Cain (1977). Statistical theories of discrimination in labor markets. *Ilr Review 30*(2), 175–187.

Arrow, K. J. (1972). Some mathematical models of race discrimination in the labor market. *Racial discrimination in economic life*, 187–204.

Bartoš, V., M. Bauer, J. Chytilová, and F. Matějka (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review 106*(6), 1437–1475.

Becker, G. S. (1957). *The economics of discrimination*. University of Chicago press.

Bendick Jr, M., C. W. Jackson, and J. H. Romero (1997). Employment discrimination against older workers: An experimental study of hiring practices. *Journal of Aging & Social Policy 8*(4), 25–46.

Bertrand, M. and E. Duflo (2017). Field experiments on discrimination. *Handbook of economic field experiments 1*, 309–393.

Bertrand, M. and S. Mullainathan (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American economic review 94*(4), 991–1013.

Bokum, N. and P. Bartelings (2009). *Age discrimination law in Europe*, Volume 2. Kluwer Law International BV.

Borenstein, M., L. V. Hedges, J. P. Higgins, and H. R. Rothstein (2021). *Introduction to meta-analysis*. John Wiley & Sons.

Burn, I., P. Button, L. F. M. Corella, and D. Neumark (2019). Older workers need not apply? ageist language in job ads and age discrimination in hiring. Technical report, National Bureau of Economic Research.

Burn, I., P. Button, L. M. Corella, and D. Neumark (2022). Does ageist language in job ads predict age discrimination in hiring? *Journal of Labor Economics 40*(3), 613–667.

Büsch, V., S.-A. Dahl, and D. A. Dittrich (2009). An empirical study of age discrimination in norway and germany. *Applied Economics 41*(5), 633–651.

Carlsson, M. and S. Eriksson (2019). Age discrimination in hiring decisions: Evidence from a field experiment in the labor market. *Labour Economics 59*, 173–183.

China Statistical Yearbook (2019). Technical report, National Bureau of Statistics of China.

Cortés, P. and J. Pan (2023). Children and the remaining gender gaps in the labor market. *Journal of Economic Literature 61*(4), 1359–1409.

Dahl, G. B. and M. Knepper (2023). Age discrimination across the business cycle. *American Economic Journal: Economic Policy 15*(4), 75–112.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics 132*(4), 1593–1640.

Drury, E. (1994). Age discrimination against older workers in the european union. *The Geneva Papers on Risk and Insurance. Issues and Practice 19*(73), 496–502.

Fang, H. and X. Qiu (2023). âgolden agesâ: A tale of the labor markets in china and the united states. *Journal of Political Economy Macroeconomics 1*(4), 665–706.

Farber, H. S., C. M. Herbst, D. Silverman, and T. Von Wachter (2019). Whom do employers want? the role of recent employment and unemployment status and age. *Journal of Labor Economics 37*(2), 323–349.

Firoozi, D., I. Burn, D. Ladd, and D. Neumark (forthcoming). Help really wanted? the impact of age stereotypes in job ads on applications from older workers. *Journal of Labor Economics*.

Glass, G. V. (1976). Primary, secondary, and meta-analysis of research. *Educational researcher 5*(10), 3–8.

Guryan, J. and K. K. Charles (2013). Taste-based or statistical discrimination: the economics of discrimination returns to its roots. *The Economic Journal 123*(572), F417–F432.

He, H., S. X. Li, and Y. Han (2023). Labor market discrimination against family responsibilities: A correspondence study with policy change in china. *Journal of Labor Economics 41*(2), 361–387.

Heckman, J. J. (1998). Detecting discrimination. *Journal of economic perspectives 12*(2), 101–116.

Helleseter, M. D., P. Kuhn, and K. Shen (2020). The age twist in employers' gender requests evidence from four job boards. *Journal of Human Resources 55*(2), 428–469.

Kelly, J. (2023). Prior ageism allegations at google, facebook and ibm raise concerns about older workers being targeted for termination. *Forbes*, 187–258.

Kleven, H., C. Landais, and J. E. Søgaard (2019). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics 11*(4), 181–209.

Kroon, A. C., M. Van Selm, C. L. Ter Hoeven, and R. Vliegenthart (2018). Reliable and unproductive? stereotypes of older employees in corporate and news media. *Ageing & Society 38*(1), 166–191.

Kuhn, P., K. Shen, and S. Zhang (2020). Gender-targeted job ads in the recruitment process: Facts from a chinese job board. *Journal of Development Economics 147*, 102531.

Lagakos, D., B. Moll, T. Porzio, N. Qian, and T. Schoellman (2018). Life cycle wage growth across countries. *Journal of Political Economy 126*(2), 797–849.

Lahey, J. and D. R. Oxley (2018). Discrimination at the intersection of age, race, and gender: Evidence from a lab-in-the-field experiment. Technical report, National Bureau of Economic Research Cambridge, MA, USA.

Lahey, J. N. (2008). Age, women, and hiring an experimental study. *Journal of Human resources 43*(1), 30–56.

Lahey, J. N. (2010). International comparison of age discrimination laws. *Research on aging 32*(6), 679–697.

Neumark, D. (2018). Experimental research on labor market discrimination. *Journal of Economic Literature 56*(3), 799–866.

Neumark, D., I. Burn, and P. Button (2019). Is it harder for older workers to find jobs? new and improved evidence from a field experiment. *Journal of Political Economy 127*(2), 922–970.

Neumark, D., I. Burn, P. Button, and N. Chehras (2019). Do state laws protecting older workers from discrimination reduce age discrimination in hiring? evidence from a field experiment. *The Journal of Law and Economics 62*(2), 373–402.

Petit, P. (2007). The effects of age and family constraints on gender hiring discrimination: A field experiment in the french financial sector. *Labour Economics 14*(3), 371–391.

Phelps, E. S. (1972). The statistical theory of racism and sexism. *The american economic review 62*(4), 659–661.

Riach, P. A. and J. Rich (2006). An experimental investigation of age discrimination in the french labour market.

Riach, P. A. and J. Rich (2010). An experimental investigation of age discrimination in the english labor market. *Annals of Economics and Statistics/Annales d'économie et de Statistique*, 169–185.

Richardson, B., J. Webb, L. Webber, and K. Smith (2013). Age discrimination in the evaluation of job applicants. *Journal of Applied Social Psychology 43*(1), 35–44.

Rosen, B. and T. H. Jerdee (1976). The influence of age stereotypes on managerial decisions. *Journal of applied psychology 61*(4), 428.

Van Borm, H., I. Burn, and S. Baert (2021). What does a job candidate's age signal to employers? *Labour Economics 71*, 102003.

Zhilian Zhaopin (2023). 2023 q1 talent market hotspot briefing. Technical report.

Zvedelikova, M. (2024). Preference for young workers in mid-career recruiting using online ads for sales jobs: Evidence from japan. *The Journal of the Economics of Ageing 27*, 100479.

Appendices

A Robustness

A.1 Robustness: Quasi-experimental analyses

Figure A.1.1: Robustness – Restricting to Matched Sample

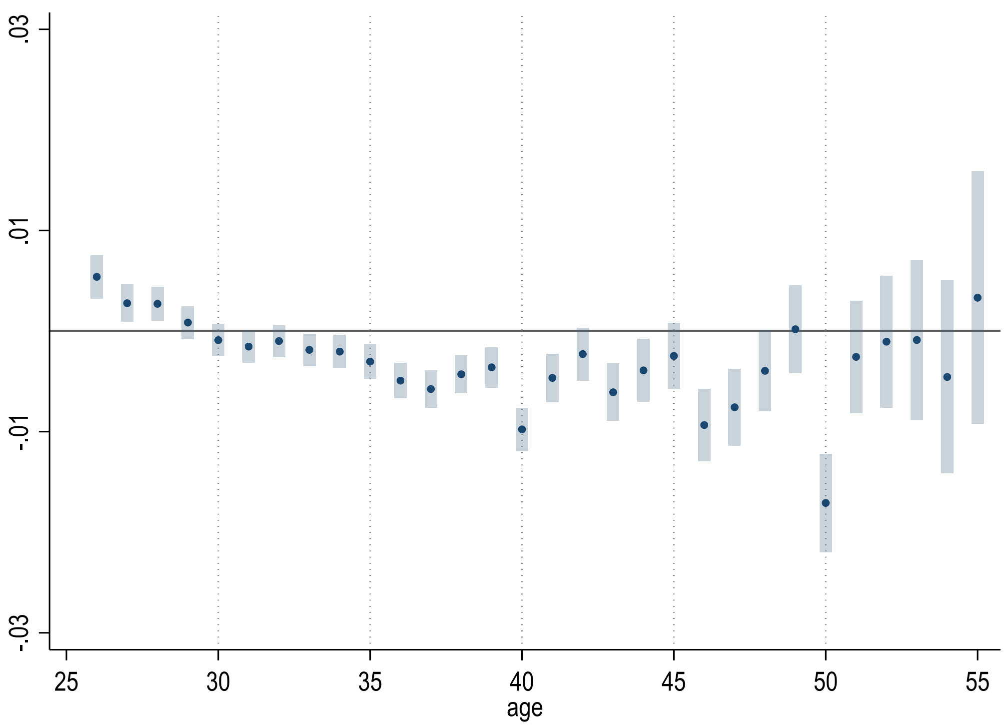

Notes: This figure plots results from replicating Figure 2(a) on *FullReview*, restricting to individuals whose basic information on the summary card – except for age – is identical within each job ad. Regression includes job-ad fixed effects and date fixed effects; robust standard errors are clustered at job-ad level.

Figure A.1.2: Robustness – 30 days around Age Change

(a) Restricting to Observations around Age Change (Individual FE)

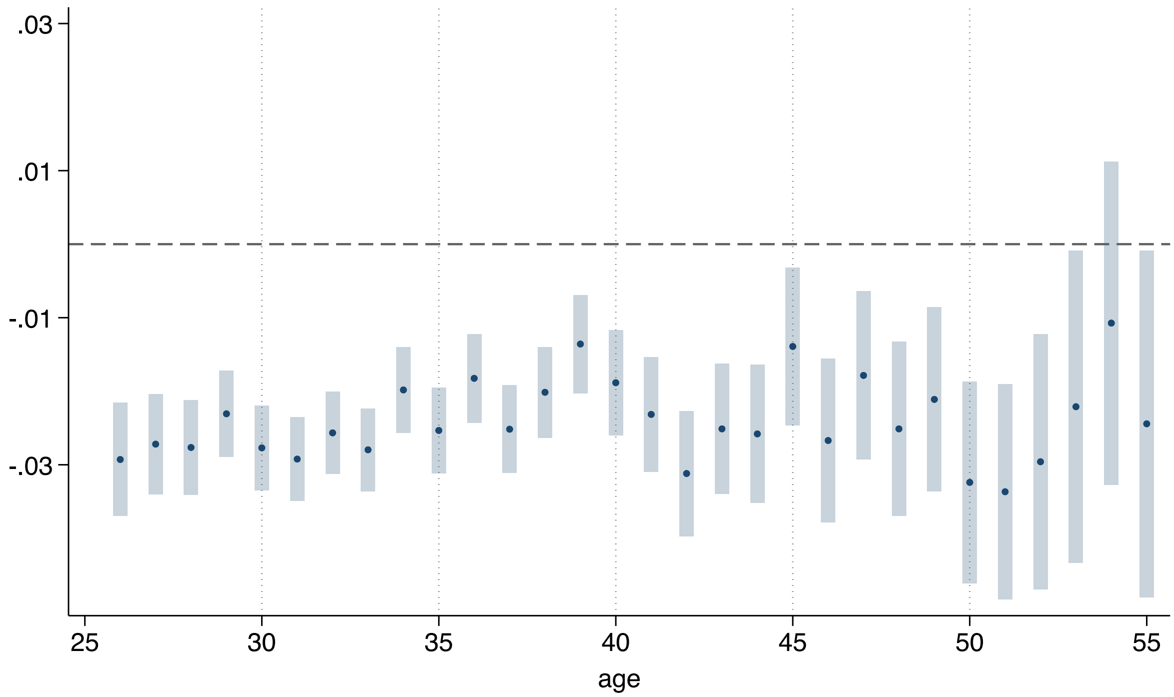

(b) Restricting to Observations around Age Change (Main specification)

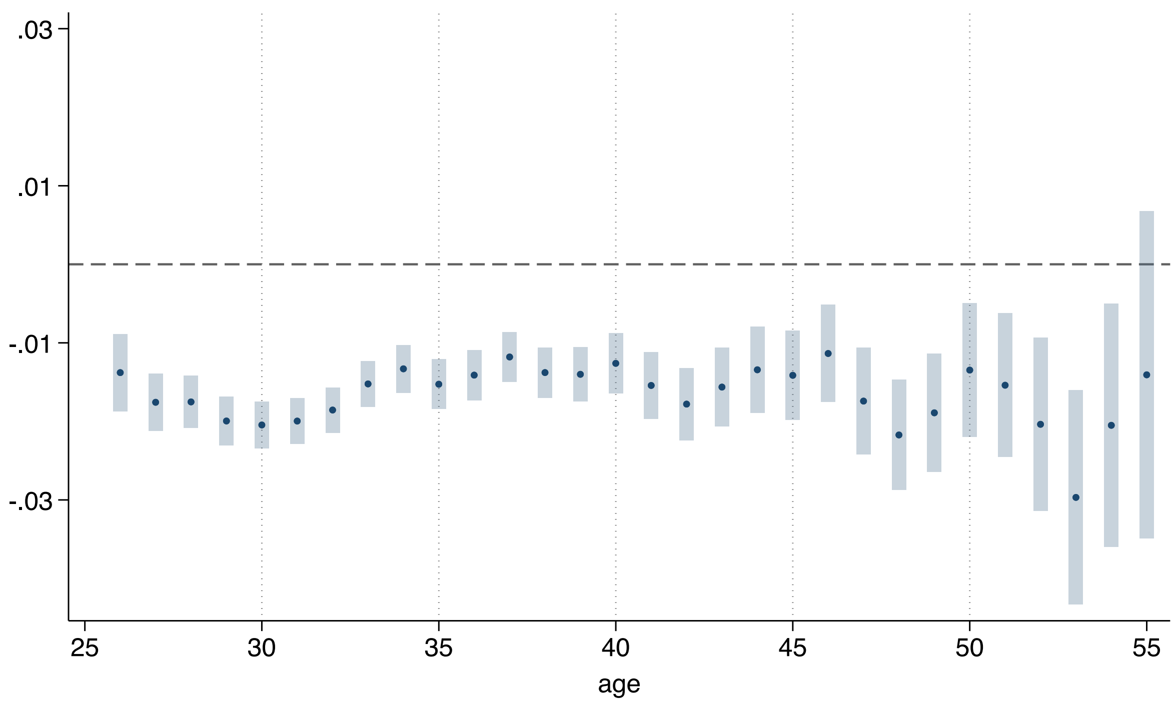

Notes: This figure plots results from replicating Figure 2(a) on *FullReview*, restricted to applications submitted within a 30-day window around an applicant's first instance of age change in the data set. In (a), we replace job-ad and date fixed effects by individual and year fixed effects. In (b), we follow the main specification in the paper, using job-ad and date fixed effects.

A.2 Robustness to Not Observing Current Employer

Figure A.2.1: Robustness – Inferred Information

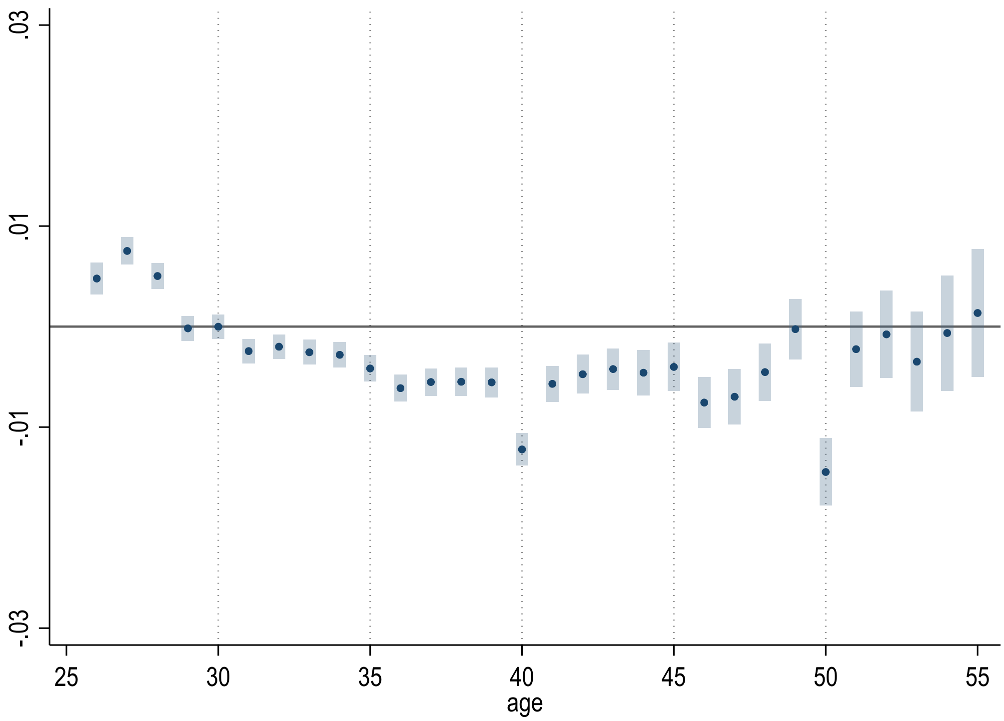

Notes: This figure plots results from replicating Figure 2(a) on *FullReview*, controlling for resume information that may be inferred from an applicant's current employer. Regression includes job-ad fixed effects and date fixed effects; robust standard errors are clustered at job-ad level.

Figure A.2.2: Robustness – Restrict to Cross-Industry Applications Only

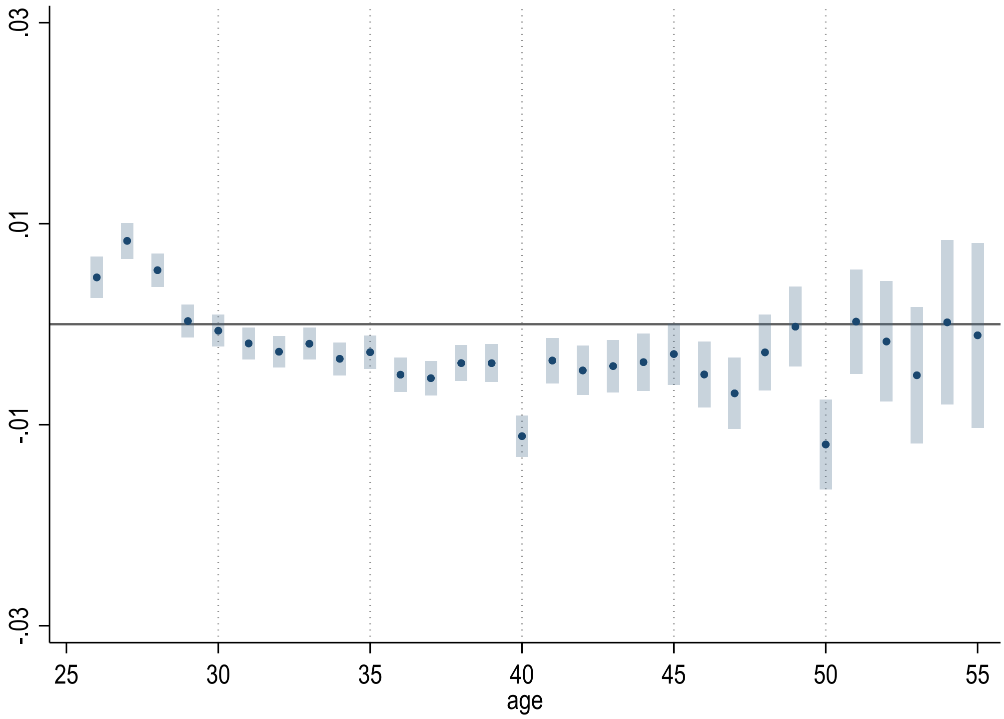

Notes: This figure plots results from replicating Figure 2(a) on cross-industry applications only (accounting for 60% of the full sample). Regression includes job-ad fixed effects and date fixed effects; robust standard errors are clustered at job-ad level.

A.3 Robustness to Applicants' Job Search Strategy

Figure A.3: Robustness – Restrict to Batch Applications Only

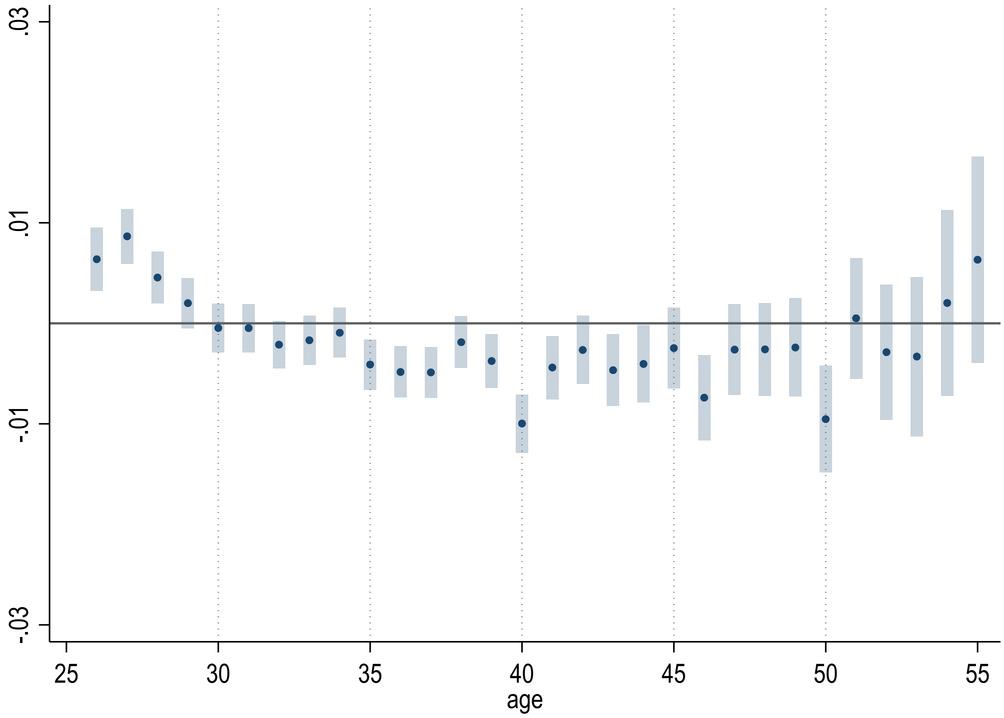

Notes: This figure plots results from replicating Figure 2(a) on batch applications only (accounting for 27% of the full sample). Regression includes job-ad fixed effects and date fixed effects; robust standard errors are clustered at job-ad level.

A.4 Robustness to Time of Application

Figure A.4: Robustness – Controlling for Job-ad by Elapsed Duration Fixed Effects

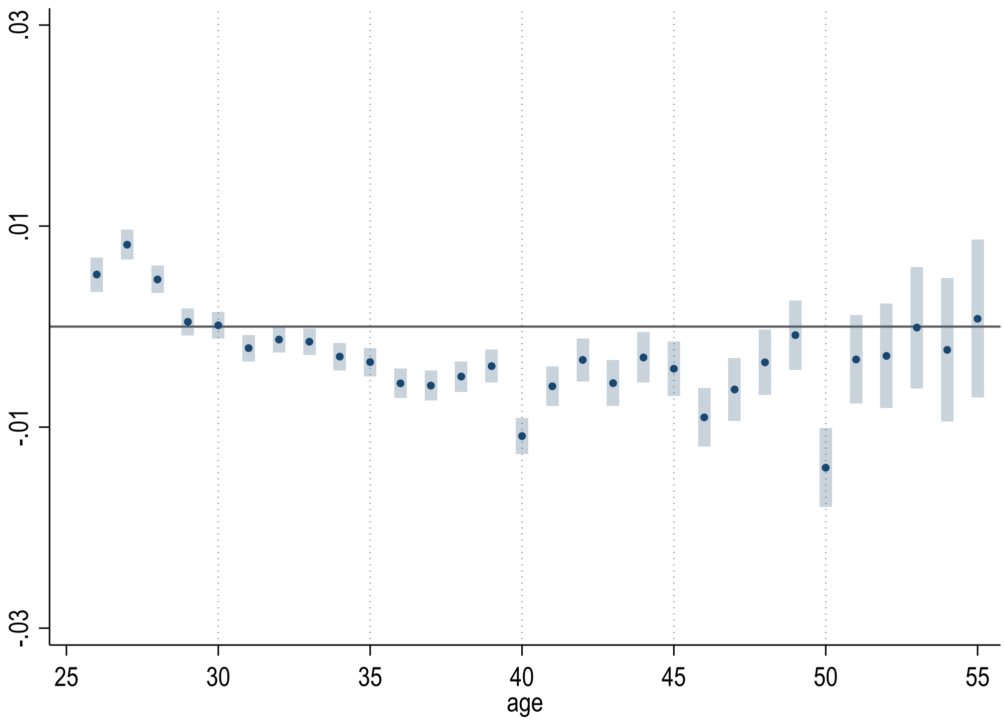

Notes: This figure plots results from replicating Figure 2(a), using job-ad by elapsed duration fixed effects. Elapsed duration is defined as six dummies identifying the application submission time: 1, 2, 3, 4 weeks, 5-8 weeks, and 9-13 weeks from the date a job-ad is posted. Robust standard errors are clustered at job-ad level.

A.5 Robustness to the Impact of COVID

Figure A.5: Robustness – Restrict to Pre-COVID and COVID Periods

(a) Pre-COVID Period (2018-2019)

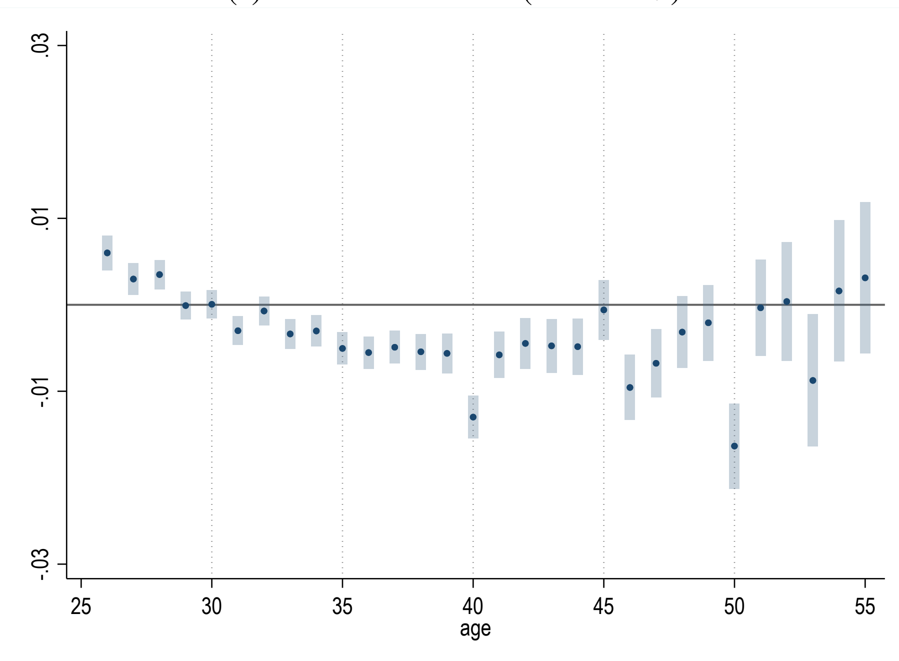

(b) COVID Period (2020-2021)

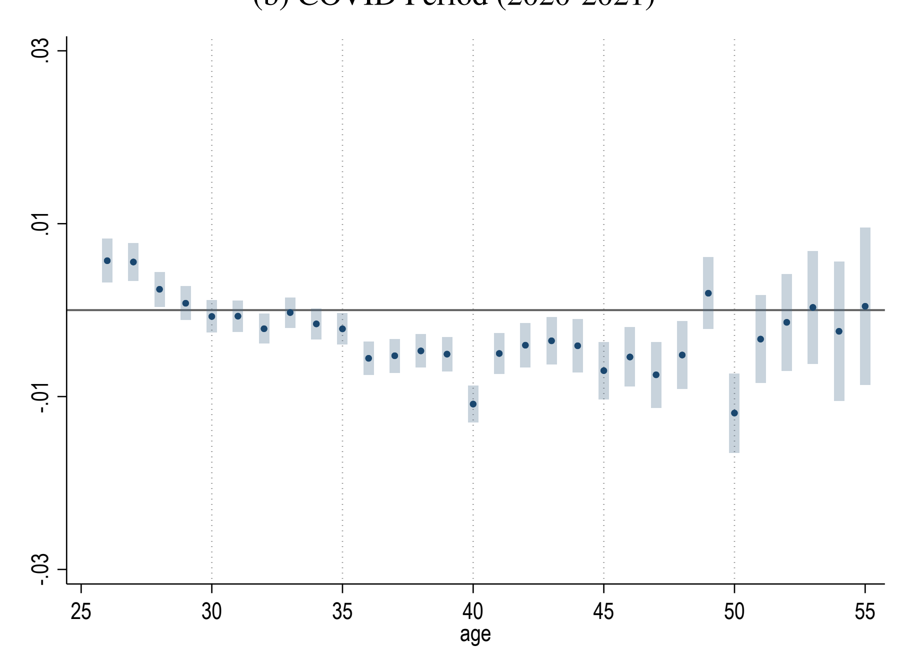

Notes: Figures (a) and (b) plot results from replicating Figure 2(a), restricting to pre-COVID period (job ads posted in 2018-2019) and COVID period (job ads posted in 2020-2021). Regression includes job-ad fixed effects and date fixed effects; robust standard errors are clustered at job-ad level.

B Heterogeneity, *FullReview* by One-Year Age Increments

B.1 Job Requirements

Figure B.1.1: *FullReview* by One-Year Age Increments, by Experience Requirements

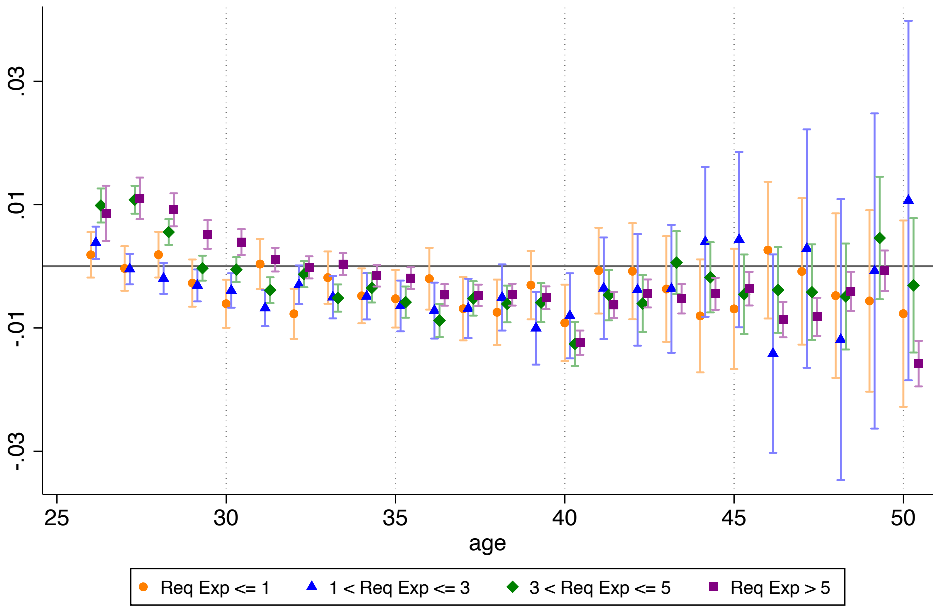

Notes: This figure replicates Figure 2(a) across subsamples based on the job's required years of work experience: less than 1 year, 1-3 years, 3-5 years, and more than 5 years. Regression includes job-ad and date fixed effects; robust standard errors are clustered at job-ad level.

Figure B.1.2: *FullReview* by One-Year Age Increments, by Management Role

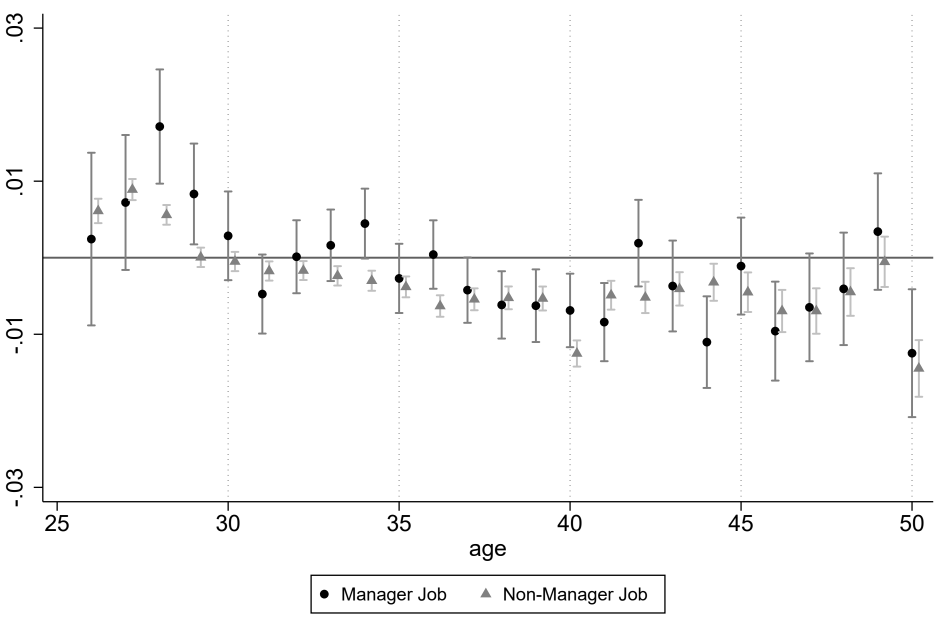

Notes: This figure replicates Figure 2(a), separately for managerial jobs and non-managerial jobs. Regression includes job-ad and date fixed effects; robust standard errors are clustered at job-ad level.

Figure B.1.3: *FullReview* by One-Year Age Increments, by Education Requirements

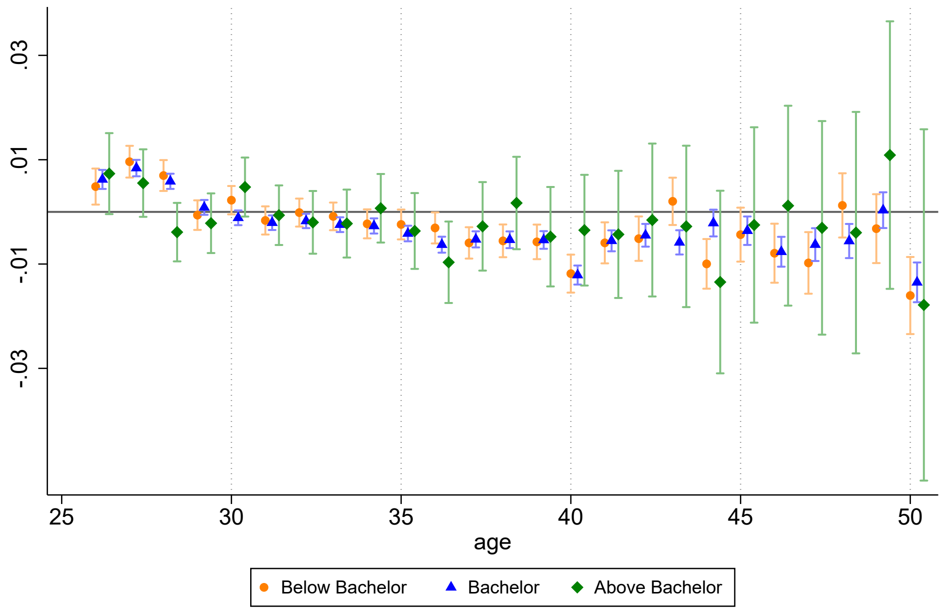

Notes: This figure replicates Figure 2(a) across subsamples based on the job's required education level: below Bachelor (associate or college), Bachelor, and above Bachelor (Master's, MBA and PhD). Regression includes job-ad and date fixed effects; robust standard errors are clustered at job-ad level.

Figure B.1.4: *FullReview* by One-Year Age Increments, by Job's Main Skill

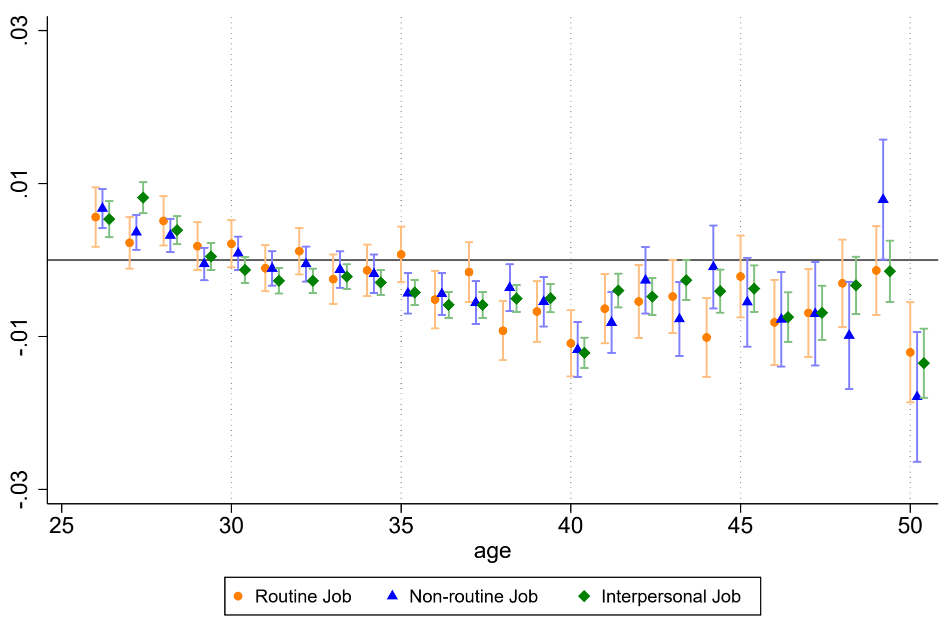

Notes: This figure replicates Figure 2(a) across subsamples based on the job's main skill: Routine skill, Non-Routine skill, and Interpersonal skill. Regression includes job-ad and date fixed effects; robust standard errors are clustered at job-ad level.

B.2 Labor Market Conditions

Figure B.2.1: *FullReview* by One-Year Age Increments, by Market Tightness

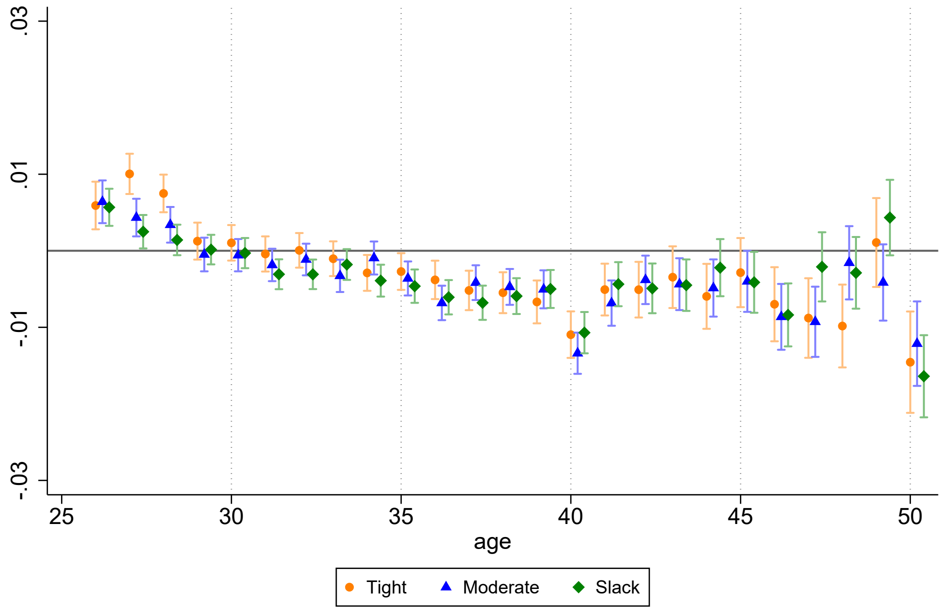

Notes: This figure replicates Figure 2(a) across subsamples based on local labor market tightness at the city-industry-occupation level. The regression includes job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.

Figure B.2.2: *FullReview* by One-Year Age Increments, by Job's Competition

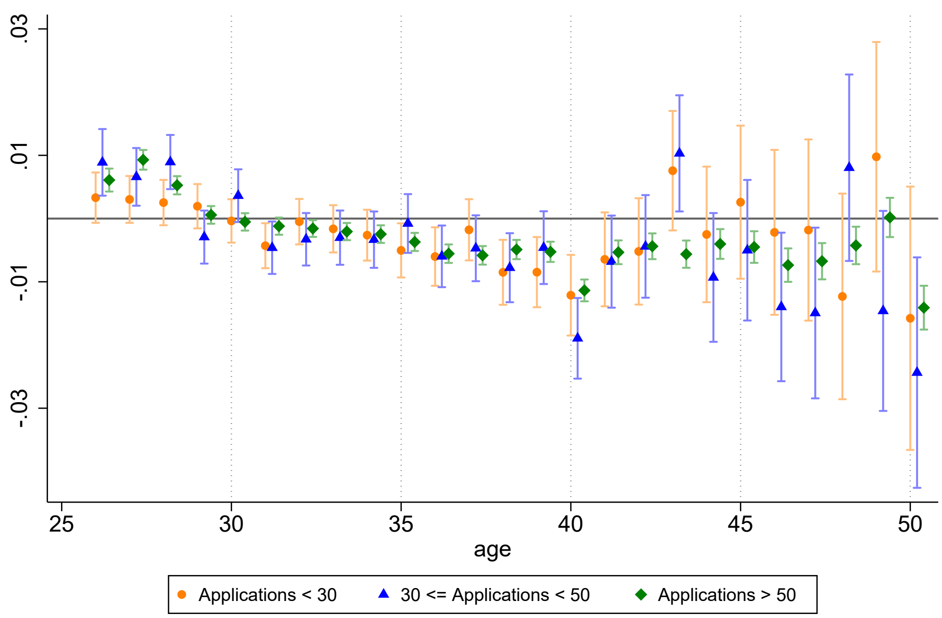

Notes: This figure replicates Figure 2(a) across subsamples based on the number of applications received: fewer than 30, 30-50, and more than 50. The regression includes job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.

B.3 Employer Characteristics

Figure B.3.1: *FullReview* by One-Year Age Increments, by Workforce Age Composition

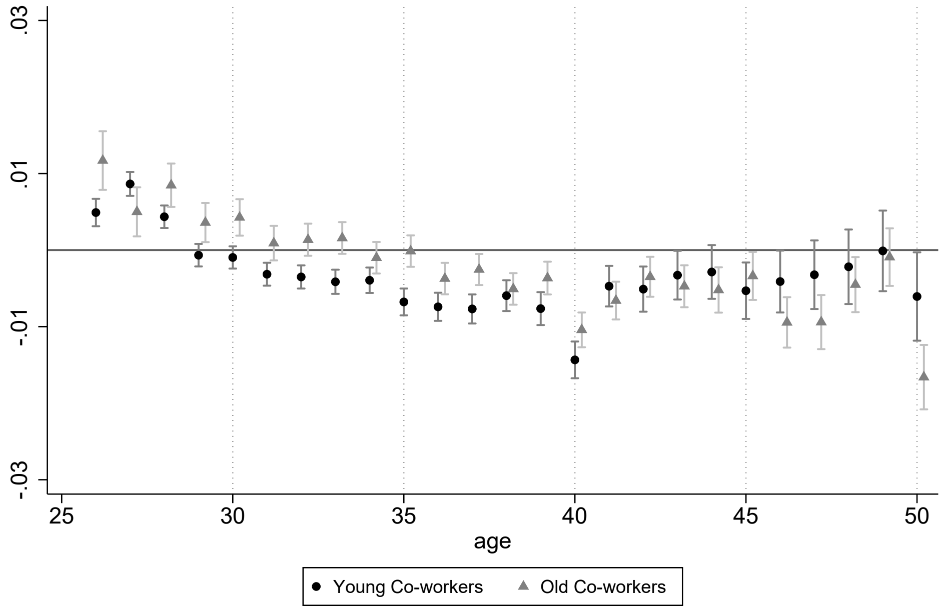

Notes: This figure replicates Figure 2(a) by cohort age groups based on the average age of applicants to the firm. The regression includes job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.

Figure B.3.2: *FullReview* by One-Year Age Increments, by Firm Size

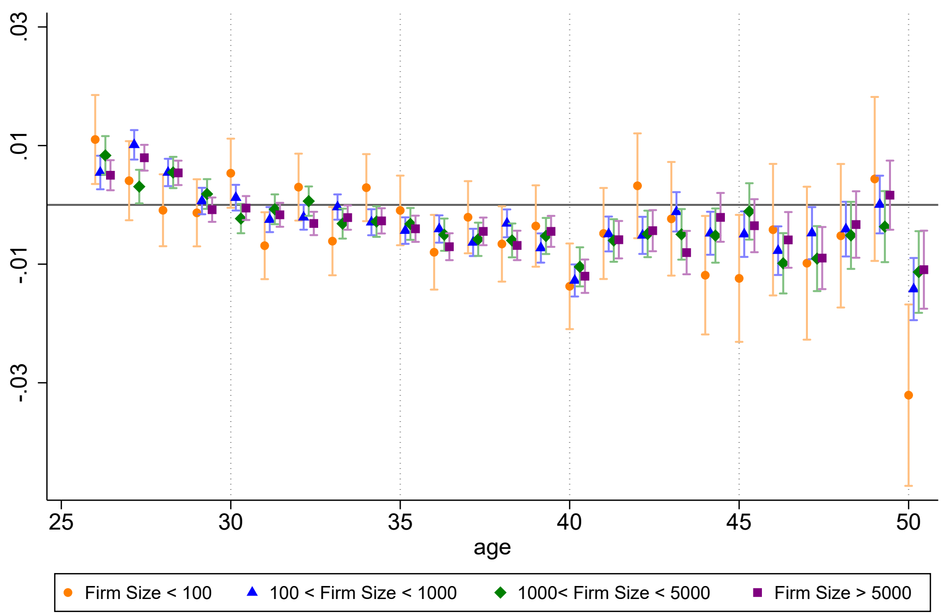

Notes: This figure replicates Figure 2(a) by subgroups based on the number of employees: less than 100 employees, 100-1000 employees, 1000-5000 employees, and more than 5000 employees. The regression includes job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.

Figure B.3.3: *FullReview* by One-Year Age Increments, by Firm Ownerership

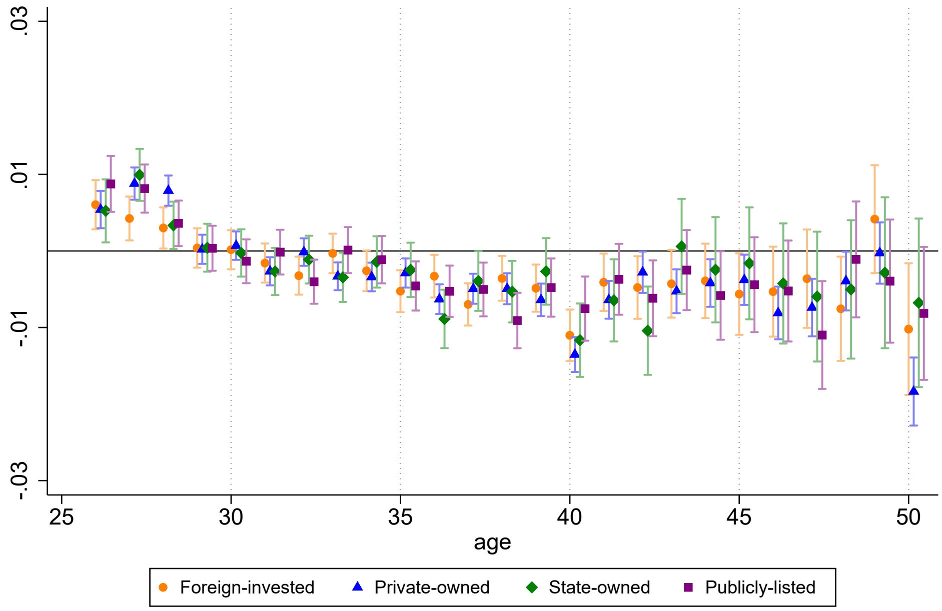

Notes: This figure replicates Figure 2(a) by subgroups based on firm ownership: foreign-invested enterprises, private-owned enterprises, state-owned enterprises, and publicly-listed firms. The regression includes job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.

Figure B.3.4: *FullReview* by One-Year Age Increments, by Firm Location

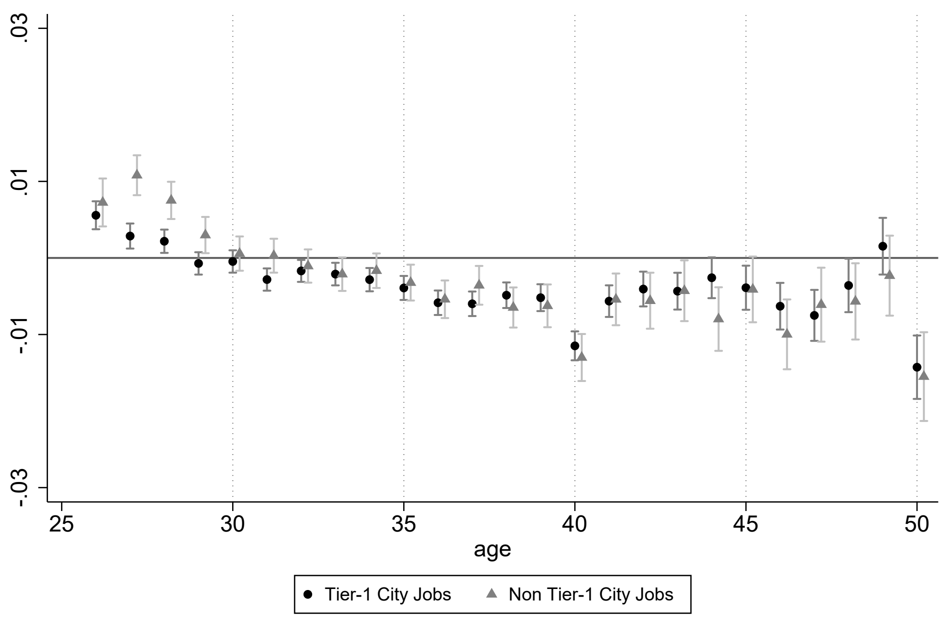

Notes: This figure replicates Figure 2(a) by subgroups based on firm location: in tier-1 cities Beijing, Shanghai, Guangzhou, and Shenzhen, and in non tier-1 cities. The regression includes job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.

B.4 Industry Sectors and Occupation Types

Figure B.4.1: *FullReview* by One-Year Age Increments, by Industry Sectors

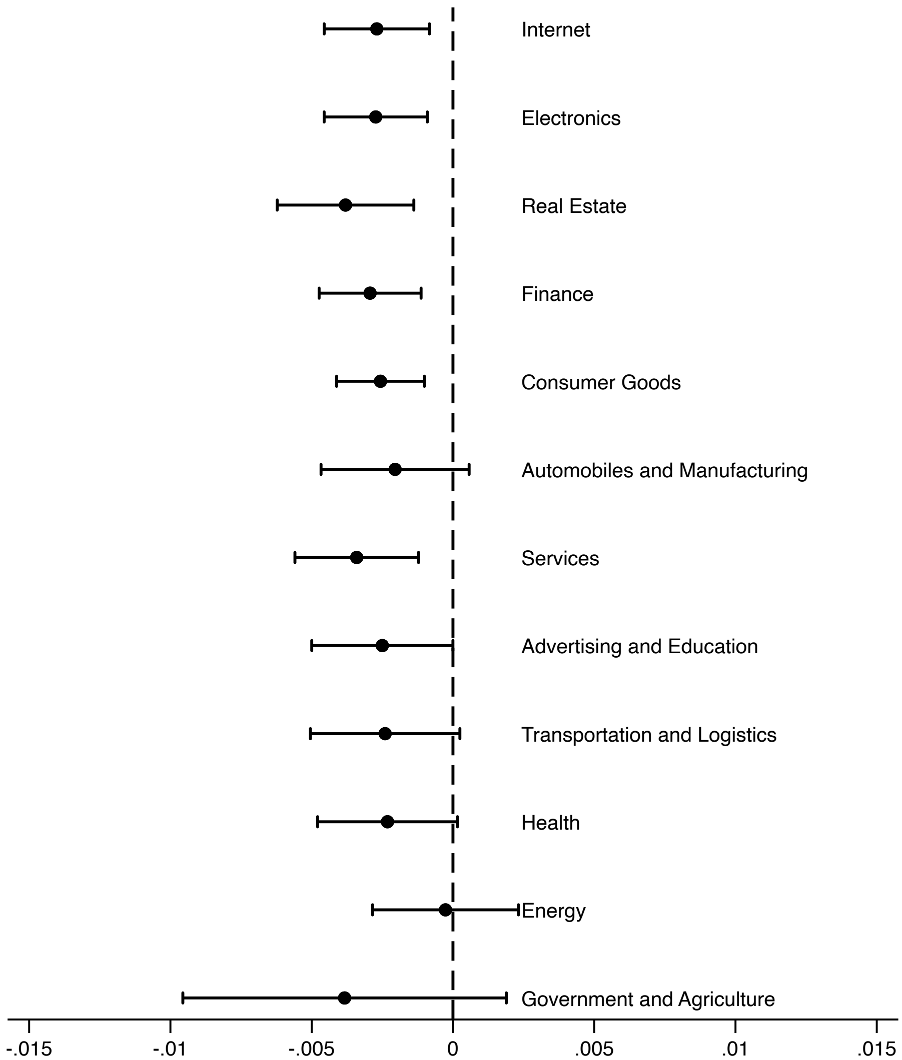

Notes: This figure plots the average coefficient estimates for applicants as they age by one year, for those aged 26-50, across 12 industry sectors. Estimates are calculated using the meta-analysis approach, treating each coefficient estimate obtained from (2) as an individual study. Standard errors are computed using the bootstrap method.

Figure B.4.2: *FullReview* by One-Year Age Increments, by Occupation Types

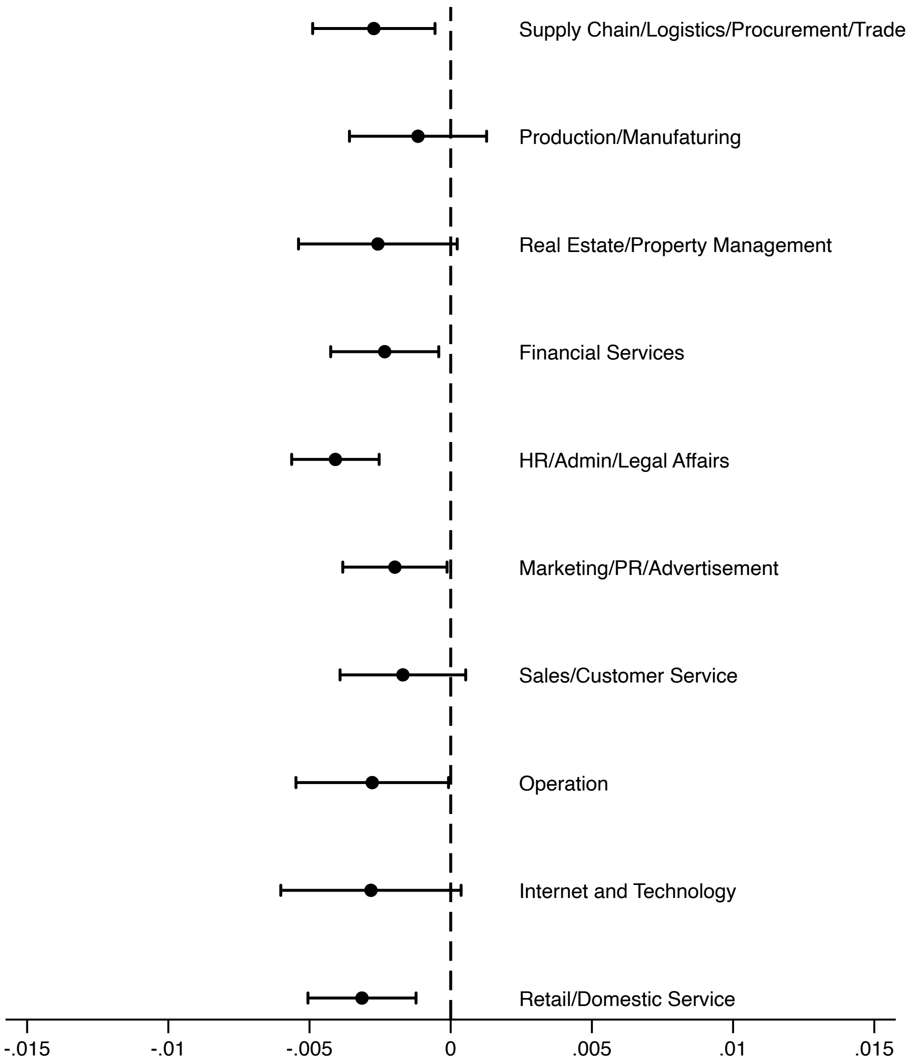

Notes: This figure plots the average coefficient estimates for applicants as they age by one year, for those aged 26-50, across the most represented 10 occupations in the data. Estimates are calculated using the meta-analysis approach, treating each coefficient estimate obtained from (2) as an individual study. Standard errors are computed using the bootstrap method.

B.5 Worker Characteristics

Figure B.5.1: *FullReview* by One-Year Age Increments, by Worker Education Level

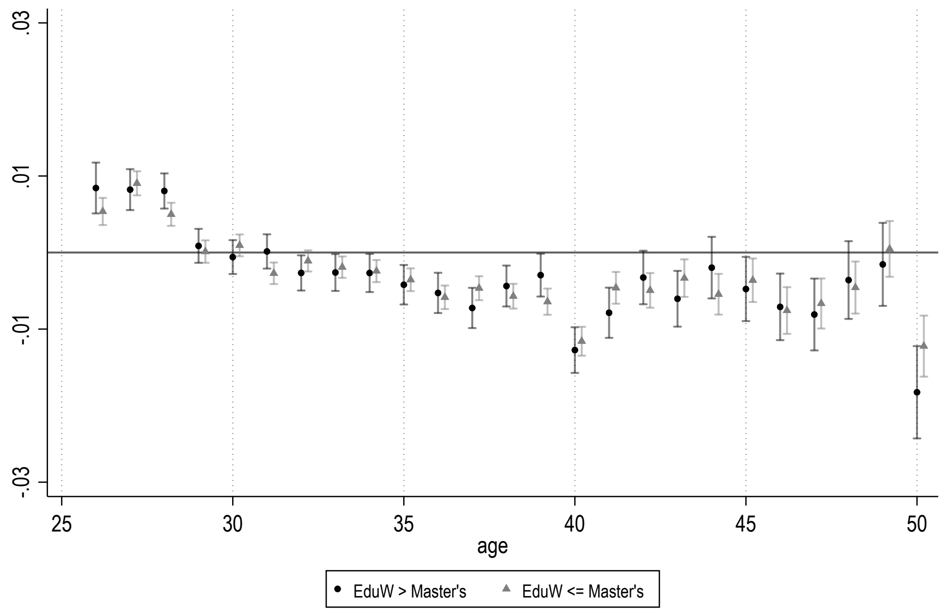

Notes: This figure replicates Figure 2(a) by subgroups based on worker educational level: Master's degree or higher versus below Master's degree. The regression includes job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.

Figure B.5.2: *Callback* by One-Year Age Increments, by Elite Education

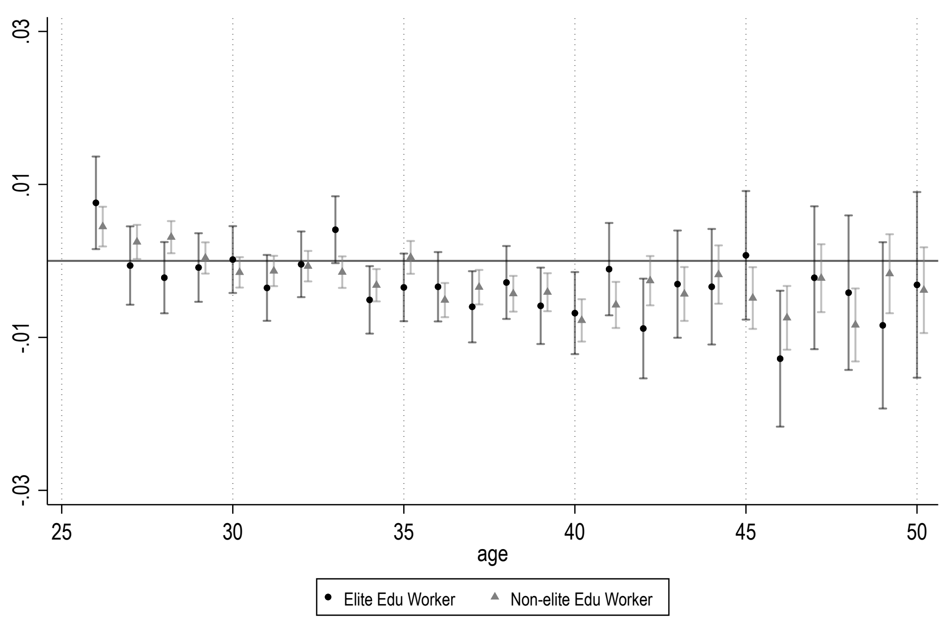

Notes: This figure replicates Figure 2(a) by subgroups based on worker's education recognition: elite institutes versus non-elite institutes. The regression includes job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.

Figure B.5.3: *Callback* by One-Year Age Increments, by Worker Employment Status

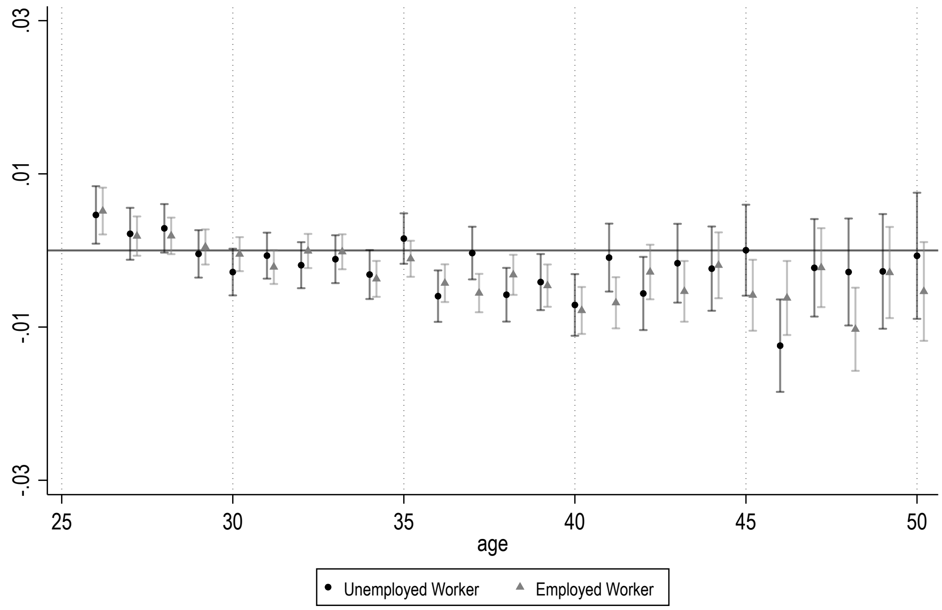

Notes: This figure replicates Figure 2(a), separately for employed and unemployed applicants. The regression includes job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.

Figure B.5.4: *Callback* by One-Year Age Increments, by Worker Wage

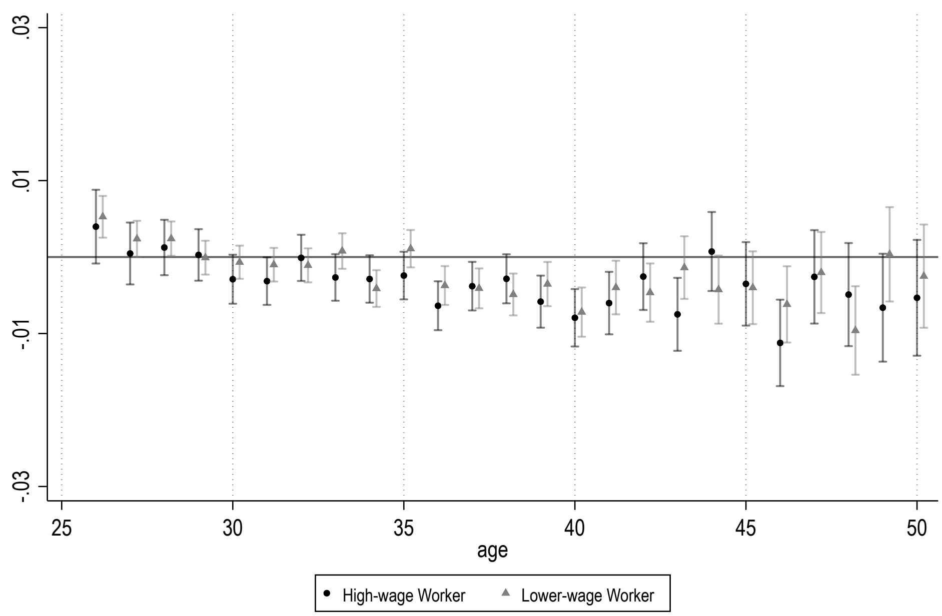

Notes: This figure replicates Figure 2(a), separately for workers whose current wage is above or below the job ad's posted wage. The regression includes job-ad and date fixed effects, with robust standard errors clustered at the job-ad level.