NBER WORKING PAPER SERIES

SIGNALS, STEPS, AND SETBACKS: AGE DISCRIMINATION IN A LADDERED LABOR MARKET

Shuaizhang Feng
Lars Lefgren
Jieyi Liu
Brennan Platt
Xiaoyu Xia

Working Paper 35709
http://www.nber.org/papers/w35709

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2026

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Shuaizhang Feng, Lars Lefgren, Jieyi Liu, Brennan Platt, and Xiaoyu Xia. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Signals, Steps, and Setbacks: Age Discrimination in a Laddered Labor Market
Shuaizhang Feng, Lars Lefgren, Jieyi Liu, Brennan Platt, and Xiaoyu Xia
NBER Working Paper No. 35709
September 2026
JEL No. J71

ABSTRACT

We study age discrimination in a labor market with a job ladder, where worker productivity is revealed over time. Using data from a large Chinese recruitment platform that links job postings, applications, and employer responses, we document sharp declines in applications and callbacks at posted age cutoffs. Discrimination operates in both directions: older applicants are disadvantaged in entry-level jobs, while younger applicants are less likely to be hired into managerial roles. A model of statistical discrimination with imperfect signal screening explains these patterns. Consistent with the model, older applicants to entry-level jobs are negatively selected relative to their peers.

Shuaizhang Feng
Jinan University
Department of Economics
shuaizhang.feng@gmail.com

Brennan Platt
Brigham Young University
Department of Economics
brennan_platt@byu.edu

Lars Lefgren
Brigham Young University
Department of Economics
and NBER
lars_lefgren@byu.edu

Xiaoyu Xia
University of Essex
Economics
xiaoyu.xia@essex.ac.uk

Jieyi Liu
Jinan University
Institute for Economic and Social Research
liujyjnu@163.com

1 Introduction

Age discrimination remains widespread in labor markets, even in countries such as the United States and those in Europe where it is legally prohibited (e.g., the Age Discrimination in Employment Act of 1967 and the Employment Equality Directive 2000/78/EC). A large body of experimental evidence demonstrates that older applicants receive significantly lower callback rates than younger candidates (Carlsson and Eriksson, 2019; Farber et al., 2016; Lahey, 2008; Neumark et al., 2016; Neumark, 2018; Neumark et al., 2019; Riach and Rich, 2010). Survey evidence confirms that age discrimination has meaningful effects on older workers' labor market outcomes (Johnson and Neumark, 1997; Stypinska and Turek, 2017). Yet the mechanisms driving this discrimination remain poorly understood.

A central puzzle is that older workers typically earn more, and many job postings require experience levels younger workers cannot meet—making the prevalence of discrimination against them in entry-level positions all the more surprising. While taste-based discrimination[1] almost certainly plays a role in some contexts—for example, when customer-facing firms believe clients prefer younger workers, or when managers prefer younger colleagues—it cannot by itself account for both discrimination in entry-level settings and preference for older workers in others without assuming arbitrary heterogeneity in employer preferences.

Alongside race, gender, and education, age provides an easily observable signal that employers may use when evaluating job applicants. Employers often hold negative stereotypes about older workers, viewing them as less capable in learning ability, adaptability, interpersonal skills, and physical strength (Burn et al., 2022; Ng and Feldman, 2012; Van Borm et al., 2021). However, models of statistical discrimination[2] appear to face the same challenge: explaining both the disadvantage of older applicants in entry-level jobs and their advantage in more advanced positions. Yet these patterns arise naturally in a simple job ladder framework, in which worker productivity is gradually revealed over time (Gibbons and Waldman, 1999; Waldman, 1984). Early in their careers, a worker's productivity is unobserved. Over time, as employees take on greater responsibility, their productivity is gradually revealed. Productive workers accumulate positive signals, qualifying them for demanding, higher-paying jobs where experience requirements ensure that applicants have demonstrated their type. By contrast, low-productivity workers fail to generate such signals and, by mid-career, have been revealed to be weaker than the average new entrant. Consequently, they fail to secure high-paying jobs that require proven

[1] Becker (1971) formalized the concept of taste-based discrimination, where interactions with certain groups generate negative utility for the discriminator.

[2] Statistical discrimination, in which employers rely on observable characteristics to predict performance, was first proposed by Phelps (1972) and Arrow et al. (1973).

performance. When these workers apply for entry-level positions, firms infer that they are disproportionately likely to be low productivity workers, leading to discrimination against older applicants.

This framework generates several testable implications. First, jobs can be categorized into early-career, entry-level positions and later-career, managerial positions. Younger workers disproportionately apply to entry-level jobs, while older workers disproportionately apply to managerial jobs. Second, because employers statistically infer productivity from the interaction of applicant age and job type, entry-level jobs are less likely to hire older applicants, whereas managerial jobs are less likely to hire younger ones. Third, managerial jobs, by design, pay more on average than entry-level jobs and typically require evidence of past success, such as promotions or strong performance in demanding roles. Finally, older applicants who nonetheless apply to entry-level positions are negatively selected relative to their peers; they are disproportionately individuals whose prior labor market experiences failed to generate positive signals of productivity.

Our analysis draws on data from a leading Chinese recruitment platform, covering one million job postings and application records for 15,000 job seekers. The Chinese setting is informative because age restrictions are widespread and often explicit, letting us verify that firms' hiring behavior aligns with declared restrictions. Remarkably, firms without explicit restrictions behave in strikingly similar ways.

This paper makes three contributions to the literature on age discrimination. First, we provide new evidence on age discrimination among college-educated workers in China, using unique data linking job postings, applicant behavior, and employer responses.[3] Second, we move beyond the low-skill contexts emphasized in most prior studies (Lahey, 2008; Farber et al., 2016; Neumark et al., 2016), documenting systematic hiring patterns across a range of high-skill and managerial jobs. Third, and most importantly, we develop and test a simple job ladder framework that explains why older applicants face disadvantage in entry-level hiring yet advantage in more advanced roles. Consistent with the model, we find that older applicants to entry-level jobs are negatively selected relative to their peers, while those applying to higher-level positions are positively selected.

Prior work documents age discrimination broadly but leaves its mechanisms underexplored. Audit studies consistently find lower callback rates for older applicants across a range of settings (Lahey, 2008; Farber et al., 2016; Neumark et al., 2016; Carlsson and Eriksson, 2019), with most focusing on low- and middle-skill occupations. Studies examining mechanisms find that ageist language in job ads predicts discrimination (Burn et al., 2022) and that older applicants fare worst when jobs require technological adaptability or

[3] Helleseter et al. (2020) and Luo et al. (2025) both examine age discrimination in job postings in China but do not have access to applicant behavior across postings.

trainability (Van Borm et al., 2021). In China specifically, Luo et al. (2025) and Helleseter et al. (2020) document age restrictions in postings, and Lin (2025) uses a regression discontinuity design to show that dismissal rates rise sharply at age 35. Unlike these studies, we observe both application behavior and employer responses, and our empirical tests are guided by a theoretical framework that generates bidirectional discrimination as an equilibrium outcome.

The remainder of the paper presents the theoretical framework, describes the data, reports the results, and concludes.

2 Theoretical Framework

Job ladder theories emphasize that career progression occurs through hierarchical levels, where promotions signal skills and productivity (Gibbons and Waldman, 1999; Waldman, 1984). Rapid advancement conveys high potential, while slow progress suggests limited ability (DiPrete and Soule, 1988). For new entrants, observable credentials such as education and training act as initial signals, whereas for experienced workers, prior job levels and promotions provide the key indicators of competence (Cassidy et al., 2016).

Signaling theory provides a natural lens for understanding how these cues function. When employers cannot directly observe productivity, they rely on signals such as career trajectory to form expectations about future performance (Spence, 1978). Deviations from expected progression are especially informative: workers who rise quickly convey strong ability, while those who remain at lower levels may be perceived as less capable.

Because career progression is closely correlated with age, a worker's age often serves as a proxy for their expected position on the job ladder (Lazear, 1992; Gibbons and Waldman, 1999). When individuals appear at unexpected rungs—for example, older workers applying for jobs typically held by much younger workers—employers interpret these deviations as signals of poor prior performance or low ability. This constitutes statistical discrimination, as such workers are negatively selected and face lower callback rates. Our empirical tests follow the tradition of Altonji and Pierret (2001), Pinkston (2006), and Arcidiacono et al. (2010), who derive and test empirical implications of statistical discrimination models.

We formalize this intuition with a model of applicant screening in which productivity signals arrive stochastically over a worker's career. The model generates a two-rung job ladder. Entry-level positions form the lower rung, accessible to all workers early in their careers; managerial positions form the upper rung, accessible only to workers who have demonstrated sufficient ability through a career signal. The key mechanism driving age discrimination is the differential rate at which high- and low-productivity workers accumulate career signals, combined with firms' imperfect ability to screen those

signals. This produces declining average quality among older entry-level applicants and a hump-shaped productivity profile among screened managerial applicants—giving rise to statistical discrimination at both ends of the ladder. In a richer model with multiple signal thresholds, this two-rung structure generalizes naturally to a k-rung ladder with qualitatively identical dynamics at each transition; the proofs, model figures, and extensions are provided in Appendix 1.

2.1 Model Setup

Consider a job market with two types of workers, with constant marginal productivity of either p_h or $p_\ell < p_h$, and let ξ denote the ratio of low- to high-types among new entrants. Over time, workers obtain a discrete signal of ability—such as a commendation, an important assignment, or a professional certification—arriving at Poisson rate α for high-productivity workers and $\beta < \alpha$ for low-productivity workers.[4] Two positions exist: entry-level (N) and managerial (M), with exogenous wages w_N and w_M.[5] We assume

$$p_h > w_M > h(0)\,p_h + \big(1 - h(0)\big)\,p_\ell > w_N > p_\ell, \tag{1}$$

where $h(0) = \frac{1}{1+\xi}$ is the fraction of high types among new entrants, ensuring that new workers qualify for entry-level jobs but the average worker cannot reach managerial positions without a signal. Employers observe applicant age but screen their signal only imprecisely: with probability $\tau \in (1/2, 1)$ they correctly assess whether a signal is genuine. Since application is costless, all workers apply to managerial jobs and embellish their resumes if they lack a true signal.

2.2 Entry-Level Positions

Among workers without a genuine signal,[6] the fraction who are high-type declines monotonically with age:

$$h(t) = \frac{1}{1 + \xi\, e^{(\alpha-\beta)t}}, \tag{2}$$

since high types accumulate signals faster ($\alpha > \beta$) and exit the entry-level pool more quickly. A firm hires an entry-level applicant of age t only if $h(t)p_h + (1-h(t))p_\ell \geq w_N$, yielding:

[4] The signal is binary, so the Poisson process describes the time to first receive a signal.

[5] Firms could pay more than these wages, due to competition over specific worker signals, without affecting our conclusions.

[6] By assumption, only these workers apply to entry-level positions. This can arise endogenously in a broader model that tracks workers' on-the-job search, as discussed in the appendix.

Proposition 1. *There is a unique age* $T_N \equiv \frac{1}{\alpha-\beta}\ln\left(\frac{p_h-w_N}{\xi(w_N-p_\ell)}\right)$ *such that employers hire entry-level applicants if and only if* $t \leq T_N$.

Firms rationally exclude older applicants not out of animus, but because adverse selection in the entry-level pool makes hiring unprofitable beyond T_N. An older worker who applies for an entry-level job is thereby revealed to have failed to accumulate the signal that would qualify them for a managerial role—their application itself is bad news about their type.

2.3 Managerial Positions

Since all workers apply to managerial positions, employers must screen. The average productivity of applicants who pass screening at age t is:

$$P(t) = p_\ell + (p_h - p_\ell)\,\frac{\tau + (1-2\tau)\,e^{-\alpha t}}{\tau(1+\xi) + (1-2\tau)\left(e^{-\alpha t} + \xi\, e^{-\beta t}\right)}. \tag{3}$$

Two forces shape $P(t)$: young workers' true signals are swamped by embellished resumes, depressing pool quality, while older workers eventually see low types catch up in signal acquisition, again diluting the pool. These dynamics are illustrated in Figures A1 and A2 in Appendix 1. These forces generate a hump-shaped productivity profile, formalized as:

Proposition 2. *If it is ever profitable to hire into managerial positions, there is a unique age range* $t \in [T_m, T_M]$ *such that employers consider and screen all managerial applicants. Younger applicants* ($t < T_m$) *and older applicants* ($t > T_M$) *are rejected without screening.*

The managerial wage must be sufficiently low (relative to high productivity) for hiring to be profitable, although this threshold can only be found numerically. Figure A2 gives an example in which hiring is profitable for some ages.

Together, Propositions 1 and 2 describe the full job ladder. The lower rung is reserved for workers who have not yet sorted themselves out—primarily the young—while the upper rung is reserved for those who have demonstrated their type, primarily the middle-aged. Critically, the model generates age discrimination in both directions simultaneously: entry-level firms exclude older applicants; managerial firms exclude the youngest. The lower bound T_m is not merely a theoretical artifact—it is directly observable in labor markets where firms post explicit age minimums or experience requirements that translate into effective age floors. We show below that 70 percent of job postings in our data carry an explicit or experience-implied minimum age, providing direct empirical grounding for this prediction.

An important implication of Proposition 2 is that workers can age out of managerial positions just as they age out of entry-level ones. The upper bound T_M exists because, as

workers grow older, low types gradually accumulate genuine signals at a rate sufficient to dilute the screened managerial pool below profitability. Without further ways to signal their quality, workers find themselves squeezed from both directions: no longer competitive for entry-level positions, yet increasingly disadvantaged in managerial hiring as well. This prediction is borne out in the data: as shown in Table 1, even senior management positions occasionally post explicit age maximums, and the implicit discrimination documented in Section 4.1 for firms without posted restrictions suggests that the same dynamic operates throughout the ladder. In a richer k-rung model, this logic applies at every transition: each rung has both an entry floor and an effective ceiling on age for a given signal history, and workers must sustain a career trajectory consistent with their age at every stage.

This bidirectional discrimination distinguishes statistical discrimination from simple taste-based accounts, and is visible in the empirical sorting and response-rate data of Section 4.2. The adverse selection driving these cutoffs is tested directly in Section 4.3.

The model's three testable predictions are: (1) wages rise with ladder position and experience requirements, managerial positions more frequently require positive signals of prior performance, and managerial hiring is bounded both below and above by age—reflecting both T_m and T_M—with explicit or experience-implied age minimums observable in a substantial share of postings; (2) workers sort applications toward jobs aligned with their career stage, with match success highest when age and job level align; and (3) older applicants to entry-level positions are adversely selected relative to both their managerial-job-seeking peers and the broader pool of same-age workers, with adverse selection growing with age.

3 Data and Descriptive Statistics

This study draws on two datasets from a leading Chinese online job platform. The first is a job postings dataset comprising a nationally representative sample of all job advertisements posted on the platform between August and September 2022. To focus on skilled positions, we restrict the sample to full-time jobs requiring at least a three-year college degree, yielding 515,434 postings. We further exclude postings with missing information on key variables (e.g., occupation, industry, or job description), removing an additional 5,009 observations. The final job postings sample consists of 510,425 observations.

Although China does not publish a unified measure of total job vacancies, two official indicators provide useful benchmarks: 12.06 million new urban jobs created in 2022 and 10.76 million college graduates in the same year, as reported by the National Bureau of Statistics of China (State Council of the People's Republic of China, 2023; Ministry of Education of the People's Republic of China, 2022b). Taken together, these figures

suggest that our two-month snapshot of job postings likely captures a meaningful share of vacancies requiring college-level education nationwide.

Within the job postings dataset, we identify age-related requirements by extracting explicit mentions of age ranges or upper limits from job descriptions. Specifically, we search for phrases such as "under X years old," "below X years old," and "aged X to X" to capture upper age constraints. We classify entry-level jobs as those requiring fewer than three years of experience and managerial jobs as those requiring three or more years.[7]

The second data source is an applicant dataset containing detailed records for 15,000 randomly selected job seekers aged 21–45 with bachelor's degrees who applied for at least one full-time position in August 2022. We focus on Guangdong Province as the study setting, as it has the largest employed population among all provinces in China.[8] The large volume of job vacancies in Guangdong generates substantial job postings and provides sufficient variation to study employers' stated hiring requirements. For context, Guangdong recorded 1.32 million newly created urban jobs in 2022 (Guangdong Provincial Bureau of Statistics, 2023), or approximately 110,000 per month on average. Our applicant sample therefore captures a substantial share of the skilled labor force in the province.

We track these sampled applicants from July to September 2022, a period during which most are actively searching on the platform. Of the initial 15,000 individuals, we exclude 3,363 applicants with missing key variables or who used the batch-apply function, as well as 30 individuals who did not apply to any identifiable positions, yielding a final sample of 11,607 applicants. These individuals submitted a total of 333,705 job applications over the three-month window. The dataset records whether firms provided positive feedback to each application, defined as the firm selecting options such as "Let's chat," "Make a call," or "Schedule an interview," or expressing equivalent intent in the conversation text.[9]

The job-posting dataset provides a broad overview of age discrimination in China's high-skilled labor market. Panel A of Table 1 reports summary statistics for job postings. Fifteen percent of postings specify an explicit upper age limit, with an average cutoff of 36.3 years. The most common threshold is 35 years, accounting for 36 percent of postings with age restrictions. This pattern reflects the well-documented "35-year-old phenomenon" in China, whereby employers systematically dismiss or avoid hiring workers once they reach their mid-30s. 6 percent indicate an explicit age minimum and 69 percent post a

[7] Our main results are robust to alternative definitions—entry-level positions requiring less than one year of experience and managerial positions requiring more than five years. Appendix 2 further validates this classification using a large language model that incorporates job title information; 90.1 percent of postings are classified identically by both methods.

[8] For example, of the 733.51 million employed workers nationwide in 2022, Guangdong accounted for nearly one-tenth (National Bureau of Statistics of China, 2023).

[9] When job seekers submit resumes, the application appears on the firm's conversation page. The system evaluates both explicit response options and conversation text to infer the firm's intent to proceed.

minimum experience requirement. All told, 70 percent of postings have either an explicit or an implicit minimum age operating through experience. Figures E9 and E10 in Appendix 5 display the full distributions of posted upper and lower age limits.

Managerial positions are less likely to impose explicit age limits; however, when such limits are present, the average cutoff is higher, consistent with the expectation that more experienced workers tend to be older. Notably, explicit age restrictions remain prevalent even among managerial postings, aligning with theoretical predictions that age discrimination can operate across all levels of the job hierarchy. Managerial positions also more frequently require positive signals from prior employment—such as management experience at listed companies, state-owned enterprises, or in team leadership roles—and offer substantially higher posted wages (approximately 2,158 USD per month, compared to 1,332 USD for entry-level positions).

The applicant dataset provides richer information on application behavior and employer responses, enabling a more granular analysis of the hiring process. Panel B of Table 1 summarizes the applicant sample. On average, individuals submit 29 applications, with 68 percent targeting entry-level positions, and 30 percent of applications receiving positive feedback. Mean prior monthly wages—defined as earnings from the most recent job reported by the applicant—are approximately 1,429 USD. Nearly half of applicants are male, 17 percent graduated from top-tier universities,[10] and more than 70 percent majored in STEM fields (including science, engineering, economics, and management).

4 Empirical Analysis

4.1 Age Discrimination in China

Posting an age maximum does not necessarily imply that firms strictly enforce it. To address this concern, we examine whether explicit age limits affect workers' application decisions and firms' hiring responses using a regression discontinuity design focused on postings that specify an upper age bound. We implement this analysis using the applicant dataset, where the running variable is applicant age relative to the posted maximum, measured in months. We assume that, in the absence of an age limit, applications would decline smoothly with age relative to the cutoff. This assumption would fail if applicants either strategically misreport their age to fall just below a posted cutoff, or selectively enter the applicant pool based on their age relative to common cutoffs—for example, if workers

[10] Top-tier universities refer to the 147 institutions designated under China's "Double First-Class" initiative, based on the official list released by the Ministry of Education of the People's Republic of China; see Ministry of Education of the People's Republic of China (2022a). Graduation from these institutions is widely regarded as a strong positive labor market signal.

close to a round-number threshold are disproportionately discouraged from registering on the platform. Figure E11 in Appendix 5 shows the age distribution of applicants in our sample. The distribution is smooth with no bunching at the round-number ages most commonly associated with posted maximums (e.g., 30, 35, 40), providing reassurance that neither form of manipulation is present in a way that would threaten identification.

Our regression model is:

$$\begin{aligned} Applications_{aj} = &\alpha + \theta \cdot \mathbb{1}\left(a > Limit_j\right) + \beta_1 (a - Limit_j) + \beta_2 (a - Limit_j) \cdot \mathbb{1}\left(a > Limit_j\right) \\ &+ \delta Z_j + \varepsilon_{aj}, \end{aligned} \tag{4}$$

where $Applications_{aj}$ denotes the number of applications from individuals of age a to job posting j, and $Limit_j$ is the posted maximum age. Z_j controls for firm size, ownership type, education requirements, and posted wages. The coefficient θ captures the discontinuous change in applications at the cutoff. The bandwidth is selected following Calonico et al. (2014), and standard errors are clustered at the firm level taking into account that firm unobservables are likely to be correlated across the ages of applicants.

Figure 1a shows the relationship between age relative to the cutoff and application counts. Applications decline sharply as age approaches the limit, reflecting applicants' awareness of reduced hiring prospects. The regression discontinuity estimate indicates that there are 0.00161 fewer applications per job posting just above the age threshold than just below it. This corresponds to a 17 percent decline at the cutoff, suggesting that applicants take posted age limits seriously when deciding where to apply.

We next estimate an analogous specification using the number of positive firm responses (callbacks) as the dependent variable. The results, shown in Figure 1b, closely mirror those for applications: callbacks decline smoothly as applicants approach the cutoff and drop discretely by nearly 20 percent at the threshold, confirming that stated age limits are binding in practice. The combination of applicant self-selection and employer decision-making generates a measurable discontinuity in opportunities for older applicants at the cutoff.

Focusing on callback rates rather than counts provides a complementary perspective but introduces an additional challenge: applicants who stop applying above the age cutoff alter the composition of the applicant pool on either side of the threshold. Consequently, differences in callback rates may partly reflect selection rather than firms' treatment of otherwise comparable applicants. Interpreting these differences as evidence of discrimination therefore requires the assumption that unobserved characteristics are balanced across the cutoff, conditional on observable covariates.[11]

[11] Even if this assumption does not hold exactly, the preceding analysis indicates that, collectively, both

Noting these caveats, we estimate the following individual-level specification:

$$PositiveResponse_{ij} = \alpha + \theta \cdot \mathbb{1}\left(a_i > Limit_j\right) + \beta_1(a_i - Limit_j) + \beta_2(a_i - Limit_j) \cdot \mathbb{1}\left(a_i > Limit_j\right) + \gamma X_i + \delta Z_j + Firm_j + \varepsilon_{ij}, \tag{5}$$

where $PositiveResponse_{ij}$ is an indicator equal to one if applicant i of age a receives a positive response from job j, $Limit_j$ is the posted age maximum, and $\mathbb{1}(a_i > Limit_j)$ indicates whether the applicant is above the stated age cutoff. The vector X_i contains observable applicant characteristics (e.g., gender, education, and field of study), and Z_j includes job-level controls such as firm size, ownership, education requirement, and posted wage. Standard errors are two-way clustered at the firm and individual levels, taking into account that an applicant's attractiveness is correlated across firms and that some firms are systematically more selective than others. Figure 1c plots individual-level regression residuals against age relative to the cutoff based on the applicant data. Callback probabilities decline smoothly approaching the limit, with applicants just above the cutoff experiencing an approximately 5 percentage point reduction in callback rates—a 13.7 percent decline relative to the pre-cutoff mean.[12]

Only a minority of firms explicitly state age-based restrictions, yet firms without posted age limits may nonetheless engage in similar practices. Based on the applicant data, Figure 2a plots positive response rates for entry-level jobs by applicant age, separately for postings with and without stated age maximums. Although this evidence is descriptive, the decline in response rates with applicant age is consistent with discriminatory behavior by firms. The patterns are nearly identical across the two groups. While firms with explicit age restrictions exhibit higher overall response rates, both groups peak at age 25 and decline at a similar rate thereafter. This suggests that explicit restrictions are not the sole indicator of discriminatory behavior; many firms appear to act on age-based preferences even when they do not state them in their job postings.

Appendix 3 reports robustness checks confirming that these discontinuities are stable across bandwidth choices, survive exclusion of observations near the cutoff, and are unchanged when age distance is measured in quarters rather than months.

4.2 Job Ladders

We have shown that many firms discriminate against older workers in entry-level positions. A central feature of our model, however, is that discrimination operates in both directions:

applicants and firms behave in a manner consistent with binding age discrimination.

[12] Implementing a donut RD excluding observations within one year of the cutoff yields an even larger discontinuity of 17 percent decline (see Appendix 3).

certain jobs are effectively reserved for older workers, reflected in explicit age minimums (5.9 percent of postings) and experience requirements (68.8 percent). While discussions of age discrimination often focus on disadvantages faced by older workers, it is important to note that the jobs restricted to older applicants offer substantially higher pay, with managerial positions averaging 2158.24 USD per month compared to 1332.2 USD for entry-level roles. This is consistent with the model's prediction that wages rise with ladder position.

To examine this formally, we regress log posted wages on job ladder indicators in the job-posting data:

$$\ln(PostedWage_j) = \alpha + \gamma_1 JobLadderIndex_j + \gamma_2 PositiveSignal_j + \delta Z_j + Occupation_j + \varepsilon_j \tag{6}$$

where $JobLadderIndex_j$ is proxied by either a managerial job indicator or the minimum age implied by the posting,[13] and $PositiveSignal_j$ indicates whether the posting requires management experience at foreign-invested or state-owned enterprises, or project leadership. We cluster correct the standard errors at the firm level to take into account systematic differences compensation levels and job requirements across firms. Table 2 reports the results. Managerial jobs offer substantially higher posted wages than entry-level jobs (column 1), positions requiring positive signals of prior performance also pay a premium (column 2), and each additional year in the minimum age requirement is associated with a 8.7 percent wage increase (column 3). Column (4) examines the combined impact of managerial jobs and positive signals on posted wages, finding that both factors contribute significantly. The final column includes both the independent variables for positive signal and minimum age requirement. These results confirm that jobs are stratified by ladder position in the way our framework assumes, and foreshadow why older workers applying to entry-level positions are viewed skeptically.

Another implication of our model is that workers sort toward jobs aligned with their career stage. Using the applicant data, Figure 2b shows that nearly all applications from workers in their early twenties target entry-level positions, a share that falls steadily to around 20 percent by the forties. Because the share of applications to managerial jobs is simply the complement, this pattern implies that older workers disproportionately target higher-level positions.

In our model, match success is most likely when workers apply to jobs on the appro-

[13] Defined as the stated minimum age, or minimum required experience plus 21. In China, age 21 often appears as the starting working age for many graduates because it aligns closely with the structure of the country's education system and typical school progression.

priate rung of the ladder. The evidence of binding age maximums and minimums already suggested such sorting, but Figure 2c provides direct confirmation. This figure plots the mean positive response rates by applicant age, with dots representing age-specific averages and a local linear (with 95% confidence intervals) superimposed to highlight the overall pattern. For entry-level positions, the positive response rate declines more steeply with age (blue line) compared to managerial jobs (red line). After age 30, applicants applying to entry-level positions experience a sharp decline in positive response rates.

We formalize this analysis with the following specification in the applicant data:

$$PositiveResponse_{ij} = \alpha + \eta_1 Age_i + \eta_2 EntryJobs_{ij} + \eta_3 EntryJobs_{ij} \times Age_i + \gamma X_i + \delta Z_j + \varepsilon_{ij} \tag{7}$$

Age_i denotes the age of applicant i, and $EntryJobs_{ij}$ indicates whether the job posting j that individual i applies to is an entry-level job. The coefficient η_3 estimates how the response rate varies with age when applicants apply for entry-level jobs. This analysis restricts the sample to workers over 25 years old, as most applicants under 25 primarily apply for entry-level jobs. The coefficient of interest, η_3, is -0.006, as shown in Figure 2c. This indicates that with each one-unit increase in age, the positive response rate for applicants applying to entry-level jobs declines by 0.6 percentage points compared to those applying to managerial jobs. This pattern is consistent with the job ladder framework, in which success depends on applying to positions aligned with one's career stage.

4.3 Adverse Selection

A central implication of our theory is that adverse selection of older workers at entry-level positions requires that those who advance to managerial roles are positively selected. One observable mechanism of this selection is firms' demand for productivity signals: only 5.7 percent of entry-level jobs require a positive signal, compared to 24 percent of managerial jobs. When stratified by experience requirement, the fraction requiring such signals rises monotonically, reaching 45 percent for jobs requiring more than 10 years of experience. Moving up the ladder demands increasingly strong evidence of high productivity.

The key implication of our model is that older workers applying to entry-level positions are adversely selected. We test this by using prior wages as a proxy for labor market productivity, which is defined as the monthly earnings from the most recent job reported by the job seeker. From the applicant data, Figure 3a shows that, conditional on demographics and education, the managerial-entry prior wage gap is positive at every age, suggesting that higher-productivity workers are more likely to target higher-level jobs.

The extent of adverse selection must, however, be understood relative to the average

worker of the same age. Formally, the difference in average quality between an age-a entry-level applicant and the overall pool of age-a workers is

$$\Delta_a = q_{ea} - \left[\omega_a q_{ea} + (1-\omega_a) q_{ma}\right], \tag{8}$$

where q_{ea} and q_{ma} are prior wages of entry-level and managerial applicants at age a, and ω_a is the share applying to entry-level jobs. Because nearly all young workers apply to entry-level positions, Δ_a is small at younger ages. While young applicants are negatively selected relative to the small number who apply to managerial jobs, they remain quite similar to the overall pool of workers their age. In contrast, most older workers target managerial positions. An older worker who nonetheless applies to an entry-level job is therefore a residual from a pool that has been repeatedly skimmed — at every career stage, the higher-productivity workers sorted away, leaving behind those who failed to accumulate qualifying signals. Consequently, older applicants to entry-level jobs are strongly negatively selected—not only relative to their peers applying for managerial roles, but also relative to the broader pool of older workers.

To demonstrate this formally, we estimate the following equation using the applicant data:

$$\ln(PriorWage_{ij}) = \alpha_a + \zeta_a \cdot EntryJob_{ij} + \mathbf{Z}_j'\gamma_a + Occupation_j + \varepsilon_{ij} \tag{9}$$

where ζ_a captures the difference in average applicant quality between entry-level and managerial applicants at age a, with occupation fixed effects and standard errors clustered at the applicant level. Figure 3b shows that entry-level applicants earn prior wages approximately 20 percent below their managerial-applying peers, with the gap slightly widening with age.[14]

Figure 3c uses the observed fractions of applicants to entry-level and managerial jobs to calculate selection relative to the overall pool at each age.[15] The gap is comparatively small for younger workers, who largely apply to entry-level positions. For older workers, however, the gap is substantial; those applying to entry-level jobs had prior wages approximately 16 percent lower than the overall regression-adjusted average for their age group.

[14] In Appendix 4 we ask whether this gap reflects pre-market credentials or productivity primarily revealed over the career, decomposing prior wages into a component explained by college selectivity, top-tier university status, gender, field of study, and graduation cohort, and a residual component. Managerial applicants are positively selected on both components, consistent with statistical discrimination on both observable and unobservable dimensions. The positive residual gap indicates that the sorting we document is not driven solely by pre-market credentials but also reflects productivity revealed over the career (Figure D8). These results are consistent with entry-level hiring managers being skeptical of both the observable and unobservable characteristics of older applicants.

[15] $\Delta_a = q_{ea} - \left[\omega_a q_{ea} + (1-\omega_a) q_{ma}\right]$ can be rewritten as $(1-\omega_a)(q_{ea} - q_{ma})$. Here, $q_{ea} - q_{ma}$ represents the average difference between applicants at age a who apply for entry-level jobs versus managerial jobs, i.e., ζ_a. Meanwhile, $(1-\omega_a)$ is the proportion of age a applicants who apply for managerial jobs.

Appendix 4 provides supplementary evidence on adverse selection using two alternative proxies for worker quality: an individual-level classification based on each applicant's portfolio of applications, and current employment status as reported on the platform. Both yield qualitatively identical conclusions.

Taken together, the empirical evidence indicates that age discrimination in China operates through both explicit and implicit channels, arising from formal age limits as well as firm behavior in the absence of posted restrictions. Applicants respond to these constraints by adjusting their application behavior, while firms reinforce them through differential callback rates. The job ladder framework provides a natural explanation for this dual pattern: older workers face disadvantages when applying for entry-level positions because their applications may signal negative selection, whereas younger workers are excluded from managerial roles until they accumulate sufficient signals of productivity. In this sense, age discrimination can be understood as a form of statistical discrimination, generating systematic disadvantages for workers whose age does not align with employers' expectations for a given job.

5 Conclusion

This paper provides new evidence on age discrimination in the Chinese labor market using unique data that link job postings, applicant behavior, and firm responses. We show that both explicit age limits and implicit employer behavior create binding barriers for workers whose age does not align with expectations for the job. Entry-level positions disproportionately exclude older applicants, while managerial positions are effectively reserved for those with strong prior signals of productivity. These patterns are consistent with a simple job ladder framework in which statistical discrimination arises naturally from the interaction of age, experience, and revealed productivity. More broadly, our findings highlight how age functions as a key sorting mechanism in labor markets, shaping both opportunities and constraints across the life cycle.

Our analysis does not rule out other explanations for discrimination. Ageist stereotypes and taste-based preferences likely account for an important share of what we observe. Yet the fact that the labor market both reveals and rewards productivity over time also means that some older workers will inevitably be left behind. Moreover, to the extent that labor market signals are noisy or imprecise, some productive older workers may be unfairly disadvantaged, reinforcing the salience of age discrimination in practice.

While age discrimination is legally permitted in our Chinese setting and age restrictions are often stated in job postings, the theoretical insights on the origins of age discrimination generalize more broadly. In any labor market where the most desirable jobs are rationed

to those who demonstrate high productivity, similar dynamics are likely to emerge. This suggests that even in settings where explicit discrimination is prohibited, the same underlying mechanisms could produce comparable patterns of disadvantage for older workers. Testing these implications in such contexts is an important direction for future research.

References

Altonji, J. G. and Pierret, C. R. (2001). Employer learning and statistical discrimination. *The Quarterly Journal of Economics*, 116(1):313–350.

Arcidiacono, P., Bayer, P., and Hizmo, A. (2010). Beyond signaling and human capital: Education and the revelation of ability. *American Economic Journal: Applied Economics*, 2(4):76–104.

Arrow, K., Ashenfelter, O., and Rees, A. (1973). Discrimination in labor markets. *The Theory of Discrimination*, pages 3–33.

Becker, G. S. (1971). *The Economics of Discrimination*. University of Chicago Press, Chicago, 2nd edition.

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

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.

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

Cassidy, H., DeVaro, J., and Kauhanen, A. (2016). Promotion signaling, gender, and turnover: New theory and evidence. *Journal of Economic Behavior & Organization*, 126:140–166.

DiPrete, T. A. and Soule, W. T. (1988). Gender and promotion in segmented job ladder systems. *American Sociological Review*, pages 26–40.

Farber, H. S., Silverman, D., and Von Wachter, T. (2016). Determinants of callbacks to job applications: An audit study. *American Economic Review*, 106(5):314–318.

Gibbons, R. and Waldman, M. (1999). A theory of wage and promotion dynamics inside firms. *The Quarterly Journal of Economics*, 114(4):1321–1358.

Guangdong Provincial Bureau of Statistics (2023). Guangdong statistical document (pdf) [in chinese]. https://stats.gd.gov.cn/attachment/0/517/517175/4146083.pdf. Accessed April 17, 2026.

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

Johnson, R. W. and Neumark, D. (1997). Age discrimination, job separations, and employment status of older workers: Evidence from self-reports. *Journal of Human Resources*, 32(4):779–811.

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

Lazear, E. P. (1992). The job as a concept. In William J. Bruns, J., editor, *Performance Measurement and Incentive Compensation*, pages 183–215. Harvard Business School Press, Cambridge.

Lin, W. (2025). Age discrimination causes unemployment: Evidence from the "35-year-old phenomenon" in china. *China Economic Quarterly International*, 5(2):147–159.

Luo, Q., Zhang, L., and Zhang, X. (2025). Age discrimination in hiring: Evidence from online job ads. Technical report, RF Berlin-CReAM Discussion Paper Series.

Marinescu, I. and Wolthoff, R. (2020). Opening the black box of the matching function: The power of words. *Journal of Labor Economics*, 38(2):535–568.

Ministry of Education of the People's Republic of China (2022a). Notice on the announcement of the second round of double first-class universities and disciplines list. http://www.moe.gov.cn/srcsite/A22/s7065/202202/W020220214318455516037.pdf. Issued jointly with the Ministry of Finance and the National Development and Reform Commission; accessed 2026-04-10.

Ministry of Education of the People's Republic of China (2022b). Number of college graduates reaches 10.76 million in 2022. http://www.moe.gov.cn/. Official statistic reported by MOE; accessed 2026-04-10.

National Bureau of Statistics of China (2023). China statistical yearbook 2023. https://www.stats.gov.cn/sj/ndsj/2023/indexeh.htm. Accessed April 17, 2026.

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

Neumark, D., Burn, I., and Button, P. (2016). Experimental age discrimination evidence and the heckman critique. *American Economic Review*, 106(5):303–308.

Neumark, D., Burn, I., Button, P., and Chehras, N. (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.

Ng, T. W. H. and Feldman, D. C. (2012). Evaluating six common stereotypes about older workers with meta-analytical data. *Personnel Psychology*, 65:821–858.

Phelps, E. S. (1972). The statistical theory of racism and sexism. *The American Economic Review*, 62(4):659–661.

Pinkston, J. C. (2006). A test of screening discrimination with employer learning. *Industrial and Labor Relations Review*, 59(2):267–284.

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

Spence, M. (1978). Job market signaling. In *Uncertainty in Economics*, pages 281–306. Elsevier.

State Council of the People's Republic of China (2023). China's job market remains stable in 2022. http://english.www.gov.cn/archive/statistics/202301/17/content_WS63c61e87c6d0a757729e592e.html. Accessed: 2026-04-10.

Stypinska, J. and Turek, K. (2017). Hard and soft age discrimination: the dual nature of workplace discrimination. *European Journal of Ageing*, 14(1):49–61.

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

Waldman, M. (1984). Worker allocation, hierarchies and the wage distribution. *The Review of Economic Studies*, 51(1):95–109.

6 Tables and Figures

Table 1: Summary Statistics: Job Postings and Applicant Characteristics

Panel A: Job postings data of large sample

Variables	**All jobs**	**Entry-level jobs**	**Managerial jobs**
Observations	510,425	344,907	165,518
Education requirement			
Three-year college	59.8%	65.9%	47.1%
University	37.6%	31.5%	50.5%
Master degree and above	2.6%	2.6%	2.4%
Have experience requirement	68.8%	53.9%	100%
Age requirement			
Specifies minimum age	5.9%	5.5%	6.9%
Average minimum age	25.4	23.8	28.0
Explicit or implied minimum age	70.4%	56.1%	100%
Specifies maximum age	14.9%	16.2%	12.1%
Average maximum age	36.3	35.0	39.9
Maximum age ≤ 30	3.2%	4.5%	0.5%
Maximum age 31-40	9.2%	10.0%	7.7%
Maximum age 41+	2.4%	1.7%	3.9%
Employer Characteristics			
Management experience required	11.5%	5.7%	23.6%
Mean posting monthly wage (USD)	1600.06	1332.20	2158.24
Small firm (<100 employees)	35.8%	37.6%	32.2%
State Owned Enterprise	7.7%	7.6%	8.0%
Number of Applications	26.4	24.1	31.4

Panel B: Applicant's characteristics

Variables	**All workers**
Observations	11,607
Average number of applications	28.8
Apply for entry-level jobs	68.2%
Getting positive response	30.4%
Mean prior wage (USD)	1428.7
Male	48.2%
Graduated from top-tier university	17.4%
Majored in STEM	73.1%

Notes:

In Panel A, entry-level jobs are defined as positions requiring less than three years of experience. In contrast, managerial jobs typically require three or more years of work experience. The variable "management experience" identifies requirements for prior managerial or high-quality professional experience. It is constructed based on a keyword mapping of the job descriptions. Relevant keywords include "management experience," "experience at listed companies, state-owned enterprises, or foreign enterprises," and "experience leading teams."

In Panel B, "Prior wage" is defined as the monthly earnings from the most recent job reported by the job seeker. Top-tier university refers to the top 147 university rated by Ministry of Education in China in 2022. STEM includes science, engineering, economics, and management. Other majors include agriculture, medicine, humanities, law and social sciences, education, and art.

Table 2: Posted Wages and Minimum Age Requirements

Dependent variable:	log of posted wage				
	(1)	(2)	(3)	(4)	(5)
Managerial jobs	0.371***	-	-	0.348***	-
	(0.002)	-	-	(0.003)	-
Positive signal	-	0.258***	-	0.152***	0.130***
	-	(0.003)	-	(0.003)	(0.003)
Minimum age/experience in years	-	-	0.087***	-	0.082***
	-	-	(0.001)	-	(0.001)
Observations	510,425	510,425	510,425	510,425	510,425
R-squared	0.371	0.305	0.380	0.378	0.385
Job posting characteristics	Y	Y	Y	Y	Y
Occupation FE	Y	Y	Y	Y	Y

Notes: This table analyzes the job posting dataset with controls for firm size, ownership type, education requirements, and occupation fixed effects. Standard errors are clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

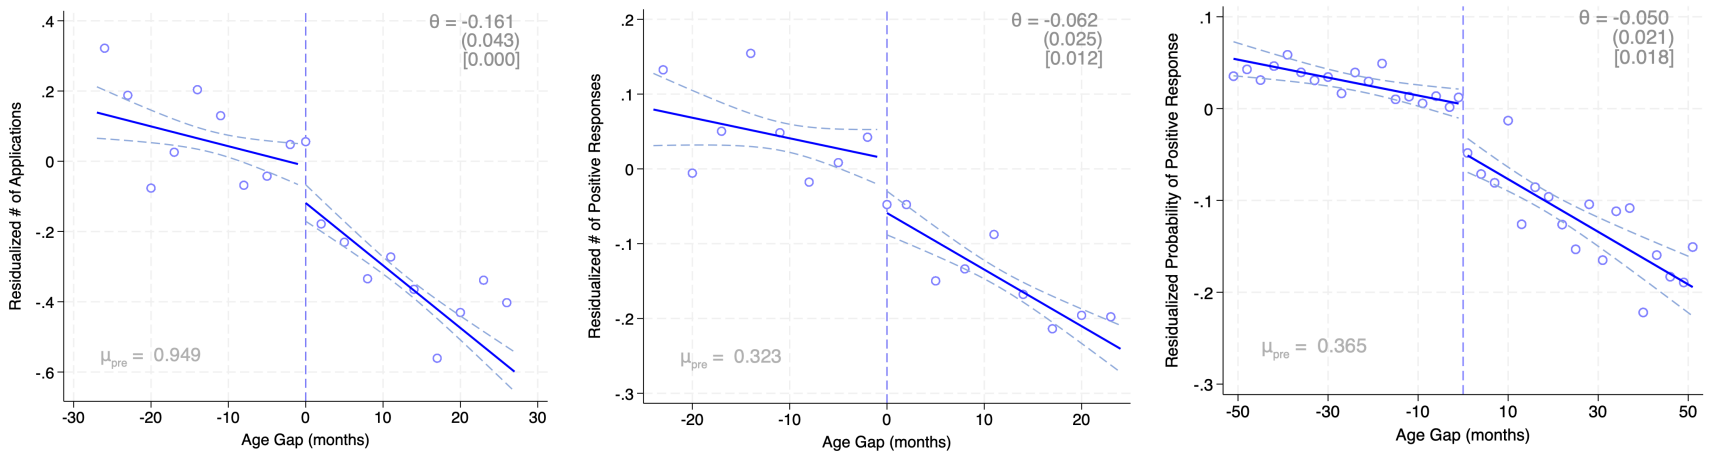

(a) Age relative to upper age limit and the number of applications at job-posting level

(b) Age relative to upper age limit and the number of positive responses at job-posting level

(c) Age relative to upper age limit and positive response rate at application level

Figure 1: Upper Age Limit and Job Application Outcomes

Notes: This figure is based on applicant data. Subfigures (a) and (b) present job-posting-level analyses controlling for firm size, ownership, education requirement, and wage, while subfigure (c) shows an application-level analysis that additional accounts for applicant characteristics such as gender, school tier, and major. Each panel applies local linear regression with a triangular kernel, using MSE-optimal bandwidths of 28, 24, and 53 months for (a), (b), and (c), respectively. Each dot represents the average of the outcome variable in three-month bins. The line are generated from linear regression with 95 percent confidence intervals. The estimated coefficients θ, standard error (shown in parentheses) and p-value (shown in brackets) are generated from equations (1) to (3), respectively. For readability, we multiply the outcome variable by 100 in the job-posting-level analyses shown in panels (a) and (b). μ_{pre} represents the pre-cutoff mean of the outcome variable. Figures without covariate controls display patterns similar to those with covariates.

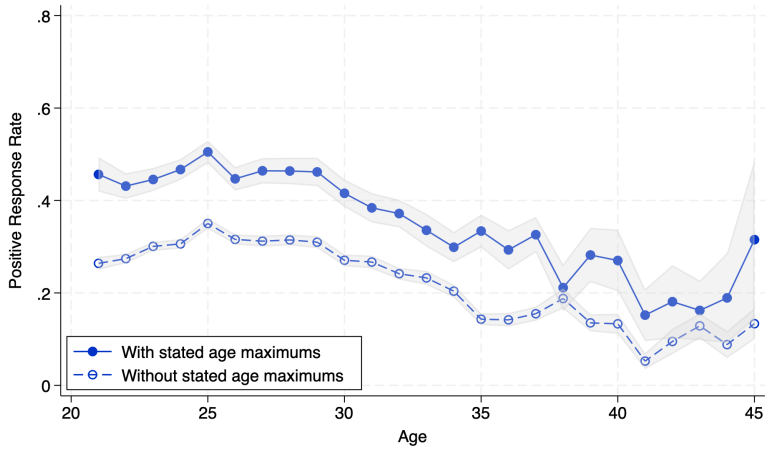

(a) Positive Response Rates for Entry-Level Jobs With and Without Maximum Age Limits

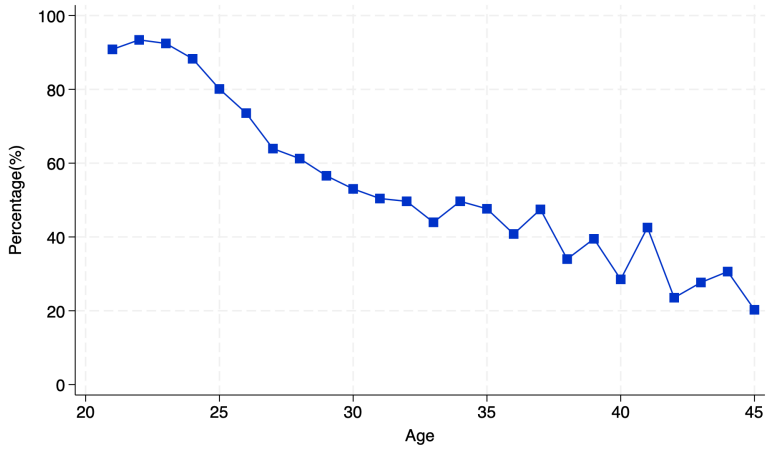

(b) Percentage of Applicants Applying for Entry-Level Jobs by Age



(c) Positive Response Rates by Age for Different Job Types

Figure 2: Application Patterns and Response Rates for Applicants by Age

Notes: This figure is based on applicant data. Panel (a) presents the positive response rates comparing jobs with and without age maximum requirements. Panel (b) displays the proportion of applicants at each age level applying for entry-level positions. Panel (c) illustrates the positive response rates across different age groups when applying for various job types. η_3 captures the slope difference in positive response rates between entry-level and managerial job applications across ages.

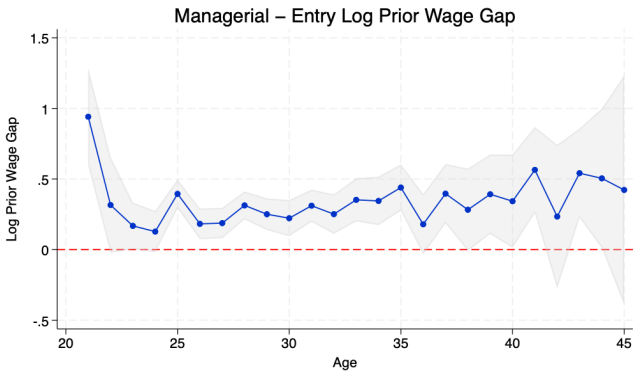

(a) Prior Wage Gap by Age for Workers Applying to Different Jobs

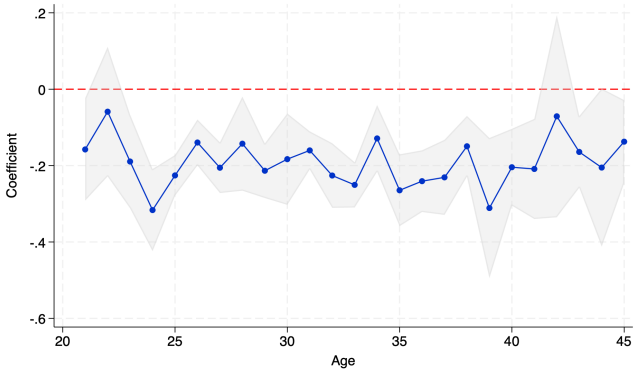

(b) Regression-adjusted Difference in Prior Wage between Workers Applying to Entry-level Jobs and Workers Applying to Managerial Positions

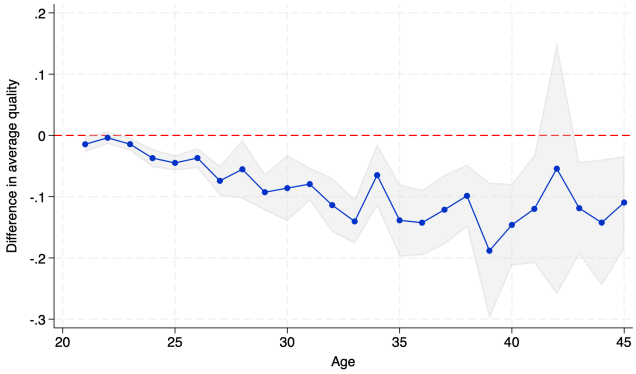

(c) Difference in Average Quality between Workers Applying to Entry-level Jobs and an Average Worker of the Same Age

Figure 3: Older Workers Applying to Entry-level Jobs Are Negatively Selected

Notes: This figure is based on applicant data. In Panel (a), the line shows the gap in prior wage levels between applicants to managerial and entry-level jobs by age group. We control for gender, top-tier university status, major fixed effects, and graduation-year fixed effects, with standard errors clustered at the applicant level. Panel (b) presents the coefficients for ζ_a for each age a in Eq. (9). Panel (c) displays the difference in average quality (Δ_a) between workers applying to entry-level jobs and the average worker of the same age, based on prior wage data. Based on Eq. (8), we calculate the Δ_a by $\zeta_a \times (1 - \omega_a)$. Shaded areas indicate 95% confidence intervals. For panels (b) and (c), the confidence intervals are derived from 500 bootstrap replications.

Appendix 1 Theoretical Appendix

This appendix provides proofs of Propositions 1 and 2 from Section 2, followed by figures illustrating the key model dynamics, and a discussion of model extensions.

Appendix 1.1 Proofs

Proof of Proposition 1. Since only relative proportions matter, normalize the initial population of each type of worker as $u_h(0) = 1$ and $u_\ell(0) = \xi$. Signals arrive at Poisson rate α for high types and β for low types, so the populations of workers without a signal evolve as $u_h(t) = e^{-\alpha t}$ and $u_\ell(t) = \xi e^{-\beta t}$. The probability that an age-t entry-level applicant is a high type is

$$h(t) = \frac{u_h(t)}{u_h(t) + u_\ell(t)} = \frac{1}{1 + \xi\, e^{(\alpha-\beta)t}}.$$

An employer hires into an entry position if and only if $h(t)p_h + (1-h(t))p_\ell \geq w_N$. Since $h(t)$ is strictly decreasing in t, setting this condition with equality yields the unique threshold

$$T_N = \frac{1}{\alpha - \beta} \ln\left(\frac{p_h - w_N}{\xi(w_N - p_\ell)}\right).$$

The argument of the logarithm is positive under the assumption $p_h > w_N > p_\ell$. For $t \leq T_N$ the expected productivity weakly exceeds w_N, so employers hire; for $t > T_N$ it falls strictly below w_N, so employers do not hire. □

Proof of Proposition 2. Let $s_h(t) = 1 - e^{-\alpha t}$ and $s_\ell(t) = \xi(1 - e^{-\beta t})$ denote the (relative) populations of workers with a genuine signal. Under imperfect screening with accuracy $\tau \in (1/2, 1)$, the mass of type-j workers accepted is $\tau s_j(t) + (1-\tau)u_j(t)$. The average productivity of accepted applicants is

$$\begin{aligned} P(t) &= \frac{[\tau s_h(t) + (1-\tau)u_h(t)]\, p_h + [\tau s_\ell(t) + (1-\tau)u_\ell(t)]\, p_\ell}{\tau s_h(t) + (1-\tau)u_h(t) + \tau s_\ell(t) + (1-\tau)u_\ell(t)} \\ &= p_\ell + (p_h - p_\ell)\, \frac{\tau + (1-2\tau)e^{-\alpha t}}{\tau(1+\xi) + (1-2\tau)\left(e^{-\alpha t} + \xi e^{-\beta t}\right)}. \end{aligned} \tag{10}$$

Boundary values. At $t = 0$: $P(0) = \frac{p_h + \xi p_\ell}{1+\xi}$, which equals the population average and is strictly less than w_M by assumption (1). As $t \to \infty$: $P(t) \to \frac{p_h + \xi p_\ell}{1+\xi} < w_M$ as well, since both $e^{-\alpha t}$ and $e^{-\beta t}$ vanish.

Initial slope. A direct calculation shows $P'(0) > 0$, so $P(t)$ is increasing near $t = 0$.

Unique interior maximum. One can verify that $P''(t) < 0$ whenever $P'(t) = 0$, so any critical point is a strict local maximum. Together with the boundary conditions and $P'(0) > 0$, this

implies $P(t)$ has exactly one interior maximum.
Conclusion. Either $P(t) < w_M$ for all t—in which case managerial hiring is never profitable—or there exists a unique interval $[T_m, T_M]$ on which $P(t) \geq w_M$. Employers screen and hire applicants in this age range and reject all others. □

Appendix 1.2 Model Extensions

Several aspects of the stylized model can be generalized without substantially changing the results.

Worker heterogeneity and multiple signals. Workers could have a richer set of productivity types, with signal arrival rates that increase with productivity. If workers can accumulate multiple signals, the analysis of any given signal quantity proceeds similarly; the exact age thresholds would vary with signal quantity, but the qualitative pattern is preserved. A small positive cost of application would not discourage embellished applications to managerial jobs, since the potential gains from passing screening are large.

Aging out of each rung. The two-rung model delivers a sharp prediction that is worth emphasizing: just as entry-level positions have an effective age maximum T_N, managerial positions have an effective age maximum T_M. Workers do not simply age into a rung and remain there indefinitely—they must continue to accumulate signals of excellence commensurate with their age throughout their careers. A single early signal is not sufficient to sustain access to the upper rung as low types gradually catch up in signal acquisition. In a richer k-rung model, this logic applies at every transition: each rung has both an entry floor and an effective ceiling, and employers at every level rationally update their beliefs about applicant quality based on the combination of age and observed signals. Workers who fail to maintain a trajectory consistent with their career stage find themselves squeezed from both directions—no longer competitive for lower rungs, yet increasingly disadvantaged on higher ones as well.

Directed applications. As presented, the model assumes that only workers without a signal apply to entry-level jobs, while all workers apply to managerial jobs. This is an approximation for another aspect of the job ladder: workers primarily apply upward. One would expect workers to find entry-level jobs (in a matter of months) much faster than they receive a signal (in a matter of years). Before workers reach age T_m , most will have found entry-level jobs. This means that most workers will be applying to managerial positions using on-the-job search, and such workers have little reason to apply for a lateral position.

Under this interpretation, entry-level applications come from the unemployed (e.g. the few that have yet to find a job or had a job separation, most of whom have no signal). Managerial applications come from the many workers in entry-level jobs. Of course, there is a small flow of unemployed workers with signals, who might consider entry jobs as a stepping stone to managerial positions. However, if screening quality and the wage differential are sufficiently large (specifically, $\tau w_m > w_n$), they prefer the managerial job, generating the same outcome as our base model between ages T_m and T_M. The expanded model would need an outside option (e.g. leaving the workforce, relying on social insurance, or dropping to some lower tier of job) that workers turn to beyond age T_M.

Appendix 1.3 Model Figures

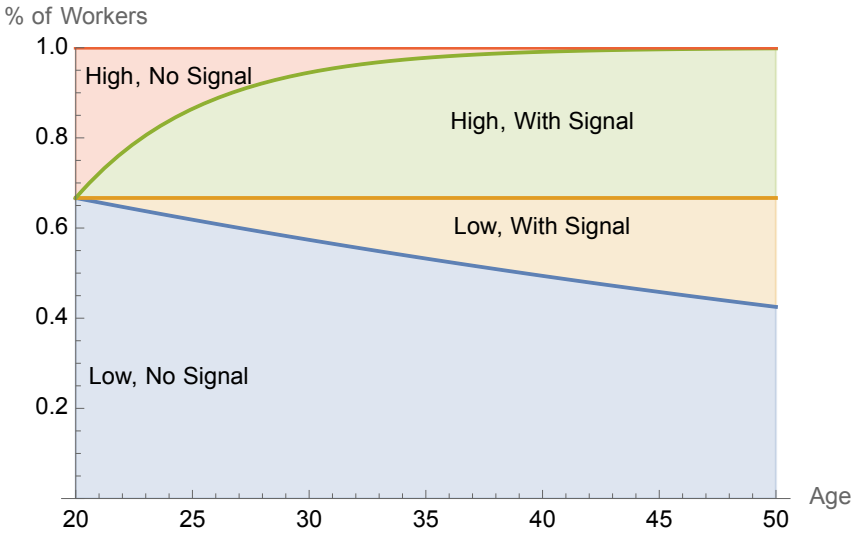

Figure A1: **Distribution of worker types with age.** The figure plots the share of workers in each type-by-signal category as a function of age, under parameters $\xi = 2$, $\alpha = 0.18$, and $\beta = 0.015$. High-type workers (top two regions) accumulate signals at rate α and transition from the "High, No Signal" group to the "High, With Signal" group faster than low-type workers (bottom two regions) transition at rate β. As a result, the pool of workers without a signal becomes increasingly dominated by low types as workers age, driving the adverse selection among entry-level applicants established in Proposition 1.

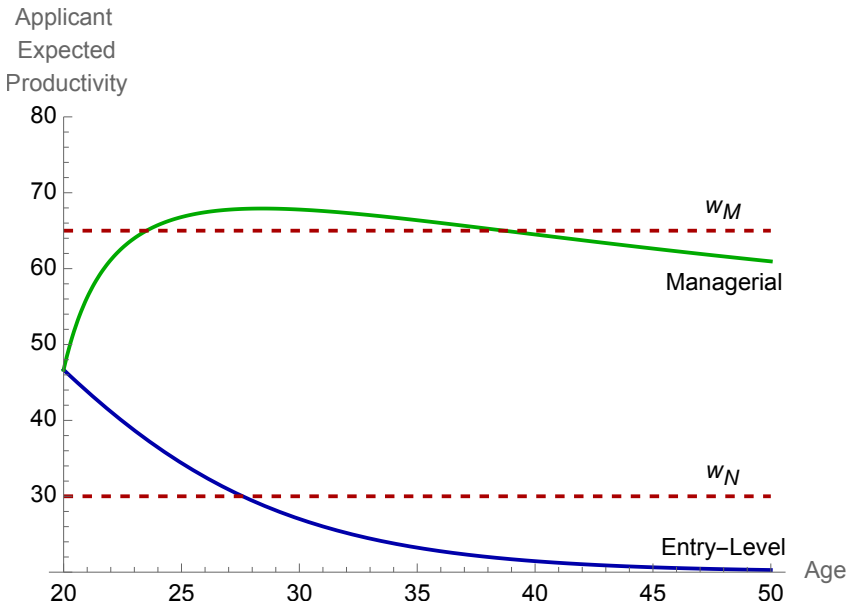

Figure A2: **Expected productivity of applicants by age and job type.** The lower solid line plots the expected productivity of entry-level applicants (workers without a signal) as a function of age; the upper solid line plots the expected productivity of managerial applicants who pass signal screening, $P(t)$ from equation (3). Dashed lines indicate the minimum wage floors w_N and w_M for entry-level and managerial positions respectively. Parameters are $p_h = 100$, $p_\ell = 20$, $\tau = 0.85$, with ξ, α, and β as in Figure A1. Entry-level employers hire only applicants whose expected productivity meets w_N, yielding the age cutoff T_N of Proposition 1. Managerial employers hire only when screened productivity meets w_M, yielding the age window $[T_m, T_M]$ of Proposition 2.

Appendix 2 Job Type Classification

In this paper, the initial classification for job type is based primarily on the explicit experience requirements stated in the job postings. However, this experience-only criterion probably carries risks in handling the ambiguous middle range of experience (e.g., 3-5 years). To assess whether our classification might introduce bias, we conduct a robustness check by employing a Large Language Models (LLM) to incorporate job title information into the classification criteria. Job titles convey important signals about a position's hierarchy and the level of experience required (Marinescu and Wolthoff, 2020), offering exactly the type of information needed to enhance the accuracy of our classification method.

LLMs are proficient at understanding and inferring textual content. Here, we utilize the Qwen-max Model, a generative pre-trained transformer, to classify the positions. We input the job title and experience requirements of a position to the model, then prompt it to generate appropriate classification results. We define entry-level jobs are junior roles requiring only little to no prior professional experience. They are designed as career entry points, often with structured training and opportunities for advancement. Managerial jobs involve overseeing teams or projects, with responsibility for decision-making and resource coordination. Relevant work experience is required to qualify for this position.

Here is the prompt that we use in the model:

> *You are a senior HR expert, specializing in job evaluation and classification.*
>
> *Please classify the following job position into one of these categories:*
>
> *-Entry-level jobs: junior roles that require little to no prior professional experience. They are designed as career entry points, often include structured training, and offer opportunities for advancement.*
>
> *-Managerial position: responsible for organizing, coordinating, and directing the work of others, requires certain work experience*
>
> *Please output in the following JSON format*
> *Note that it can only be entry-level jobs or managerial jobs, do not output your thinking and analysis process, give me the closest classification result.*
>
> *Job title: {job title}*
> *Required years of experience: {exp requirement}*

Table B1 displays that the distribution of classification results using LLM does not diverge significantly from those based on experience requirements. In the overall sample, 90.1% of job postings are classified consistently by both methods, with 214,247 (64.2%) labeled as entry-level positions by both, and 86,424 (25.9%) labeled as managerial positions

by both. As a robustness check, we also replicated our analysis using an LLM-based job classification. Given its strong alignment with the experience-based measure, the results are nearly identical to those in the main text. For the sake of conciseness, we omit the detailed output here.

Table B1: Cross-Tabulation of Two Job Classification Methods

Experience-Based Classification	**LLMs Classification**		**Total**
	Entry-Level	**Managerial**	
Entry-Level	214,247	10,377	224,624
	(64.2%)	(3.1%)	(67.3%)
Managerial	22,657	86,424	109,081
	(6.8%)	(25.9%)	(32.7%)
Total	236,904	96,801	333,705
	(71.0%)	(29.0%)	(100%)

Notes: This table presents the distribution of job postings that applicants apply for.

Appendix 3 Robustness analysis

In this section, we conduct multiple robustness checks to verify the consistency of our results across different estimation approaches. First, we perform a donut-hole regression to exclude observations near the treatment threshold. As illustrated in Figure C3, the estimates remain statistically significant and the coefficients are stable, indicating that the documented discontinuity is not driven by observations near the cutoff.

We next examine the robustness of our results to the choice of bandwidth. Our primary specification relies on the MSE-optimal bandwidth selected via the data-driven procedure described by Calonico et al. (2014). To test sensitivity, we re-estimate the model using bandwidths ranging from 50% to 150% of this baseline. Figure C4 plots the treatment effect estimates across varying bandwidths. The solid point represents the baseline estimate using the optimal bandwidth (100%), while the hollow points correspond to estimates using bandwidths ranging from 50% to 150% of the baseline, in 5% increments. The estimates remain quantitatively similar to the original results, demonstrating that our findings are not an artifact of a particular bandwidth choice. However, narrower bandwidths may yield less precise results, which consistent with the bias-variance trade-off.

Third, we alter the granularity of our key independent variable to test the sensitivity of our results to the measurement unit of age distance. Our baseline specification uses a month-based measure; here, we re-calculate the age distance using quarters (three-month intervals). As shown in Figure C5, the estimated discontinuity remains statistically significant when using this coarser unit of measurement.

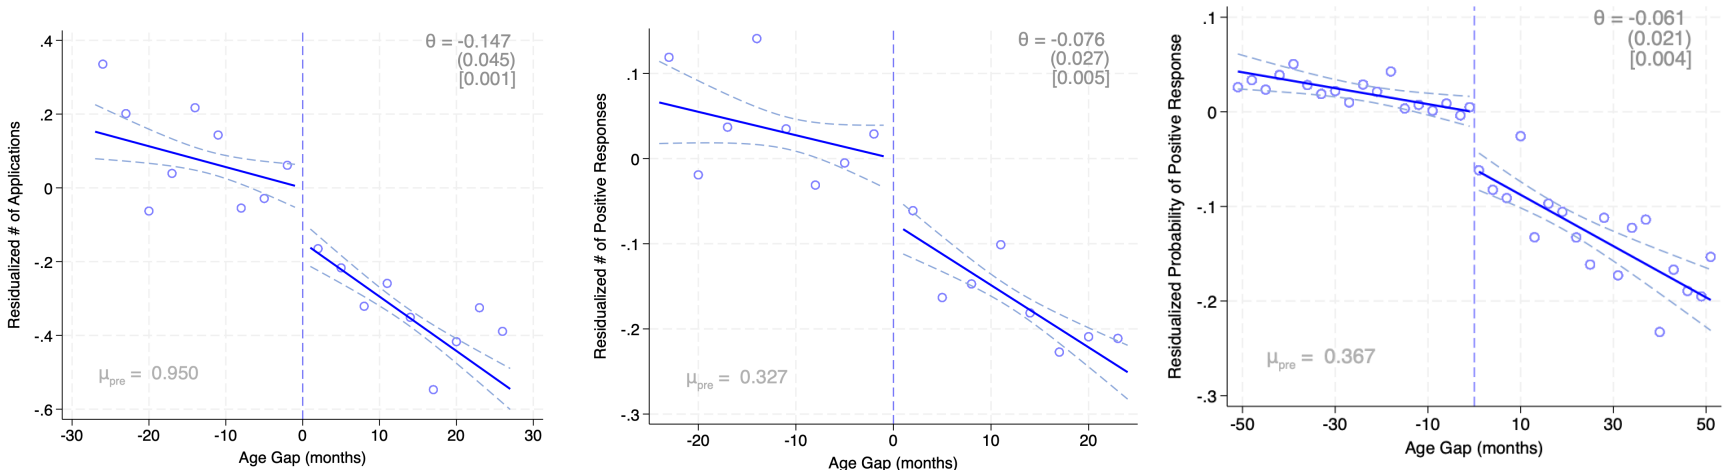

(a) Age relative to cutoff and the number of applications at job-posting level (Donut)

(b) Age relative to cutoff and the number of positive responses at job-posting level (Donut)

(c) Age relative to cutoff and positive response rate at application level (Donut)

Figure C3: Robustness Check: Age Cutoff and Job Application Outcomes (Donut Regression)

Notes: This figure is based on applicant data and presents robustness checks using a donut regression that excludes observations within a $\pm$1-month window around the age cutoff. Subfigures (a) and (b) present job-posting-level analyses controlling for firm size, ownership, education requirement, and wage, while subfigure (c) shows an application-level analysis that additional accounts for applicant characteristics such as gender, school tier, and major. The estimated coefficients θ, standard error (shown in parentheses) and p-value (shown in brackets) are generated from equations (1) to (3), respectively. For readability, we multiply the outcome variable by 100 in the job-posting-level analyses shown in panels (a) and (b). μ_{pre} represents the pre-cutoff mean of the outcome variable. Figures without covariate controls display patterns similar to those with covariates.

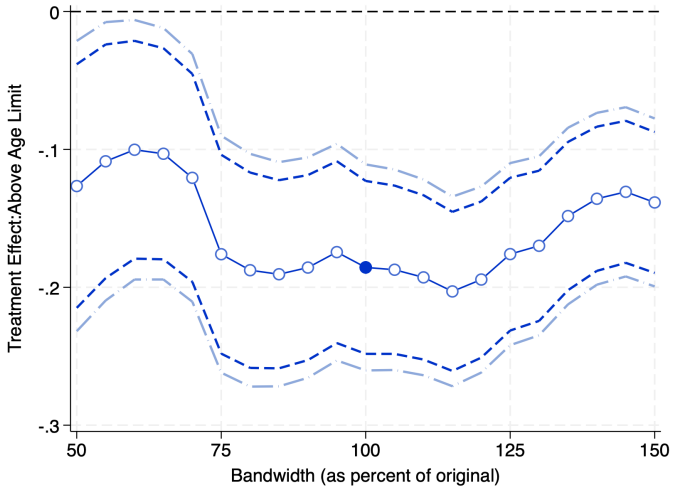

(a) Number of applications at job-posting level

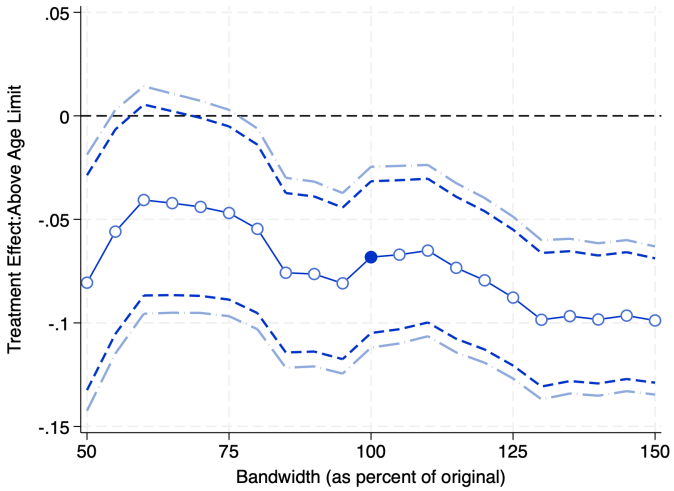

(b) Number of positive responses at job-posting level

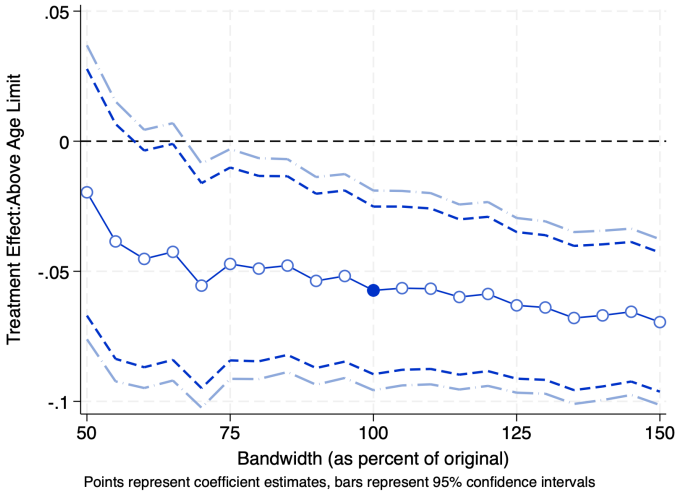

(c) Positive response rate at application level

Figure C4: Robustness Analysis: Bandwidth Sensitivity

Notes: This figure is based on applicant data and presents sensitivity tests across bandwidths ranging from 50% to 150% of the original MSE-optimal bandwidth. The solid blue point is the primary estimate. Darker blue dashed lines represent 95% confidence intervals, while lighter blue dashed lines show 90% confidence intervals.

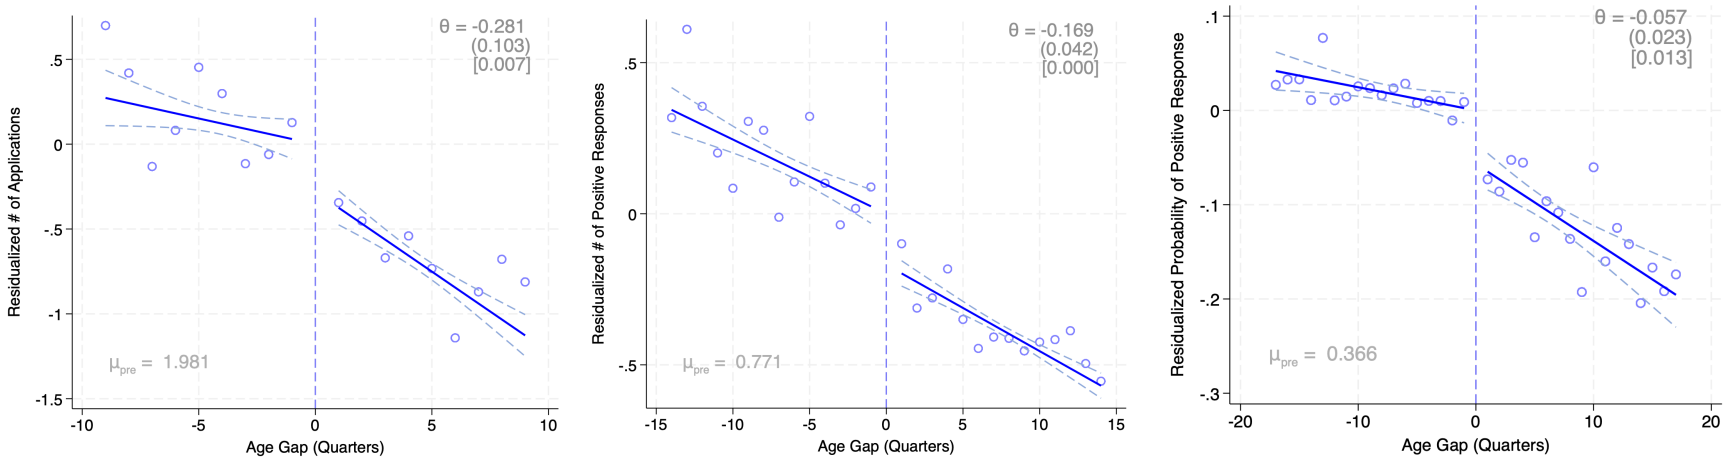

(a) Age relative to cutoff and the number of applications at job-posting level

(b) Age relative to cutoff and the number of positive responses at job-posting level

(c) Age relative to cutoff and positive response rate at application level

Figure C5: Robustness Check: Age Cutoff and Job Application Outcomes (Quarter-Based Measure, Donut=1)

Notes: This figure presents robustness checks based on applicant data, with the age gap between applicants' ages and job posting age requirements measured in quarters. Subfigures (a) and (b) present job-posting-level analyses controlling for firm size, ownership, education requirement, and wage, while subfigure (c) shows an application-level analysis that additional accounts for applicant characteristics such as gender, school tier, and major. Each panel applies local linear regression with a triangular kernel, using MSE-optimal bandwidths of 9, 1, and 17 quarters for (a), (b), and (c), respectively. The estimated coefficients θ, standard error (shown in parentheses) and p-value (shown in brackets) are generated from equations (1) to (3), respectively. For readability, we multiply the outcome variable by 100 in the job-posting-level analyses shown in panels (a) and (b). μ_{pre} represents the pre-cutoff mean of the outcome variable. Figures without covariate controls display patterns similar to those with covariates.

Appendix 4 Supplementary Evidence on Adverse Selection

In Figure 3a, we compare applicants' prior wages based on the types of jobs they apply for. A potential issue with this comparison is that individual applicants may apply for multiple job types, creating repeated observations across categories. This could introduce bias if an applicant's inherent characteristics are correlated with their application choices. To verify our results, we classify each applicant as either a junior or senior worker based on the proportion of their applications targeting each job category. This classification allows us to better investigate the relationship between job type applied and prior wage.

Specifically, we define applicants as junior workers if more than 50% of their applications are for entry-level jobs. We define applicants as senior workers if at least 50% of their applications are for managerial jobs. Figure D6 shows a positive prior wage gap between managerial-dominant and entry-dominant applicants. This provides evidence that applicants who apply more frequently to managerial jobs tend to be positively selected on prior wages.

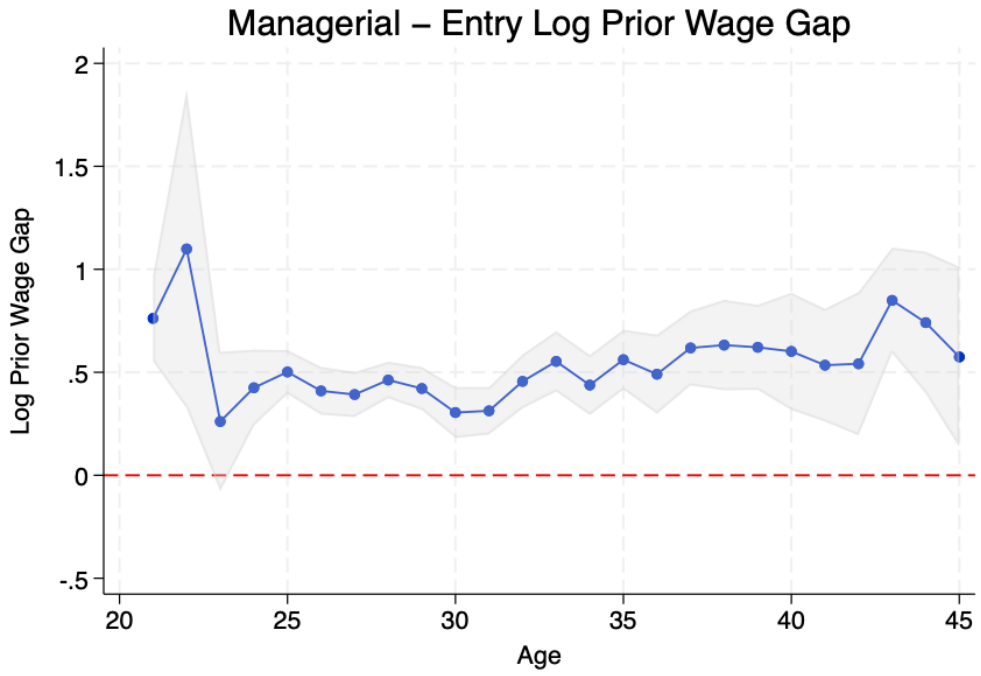

Figure D6: Prior Wages by Age for Workers Applying to Different Jobs: Individual Level

Notes: This figure is based on applicant data. In this figure, we categorized applicants by the percentage of job types they apply for. We define applicants as junior workers if they apply for more than 50% entry-level jobs. We define workers as senior workers if they apply for 50% or more managerial jobs. Shaded areas show 95% confidence intervals.

Prior wages are just one proxy for quality, which may reflect firm-specific pay-setting or reporting error as well as ability. As a complementary check, we use current employment status—an independent proxy for applicant quality—rather than prior wage information. This is based on self-reported information from the website indicating whether an applicant is "employed but seeking new opportunities." From Figure D7, we find a generally positive employment probability gap between applicants to managerial and entry-level jobs across most age groups, indicating that applicants to managerial positions are more likely to be currently employed. Even when we use employment status as a proxy for applicant quality, we still find evidence that applicants to managerial jobs are of higher quality than those applying for entry-level jobs.

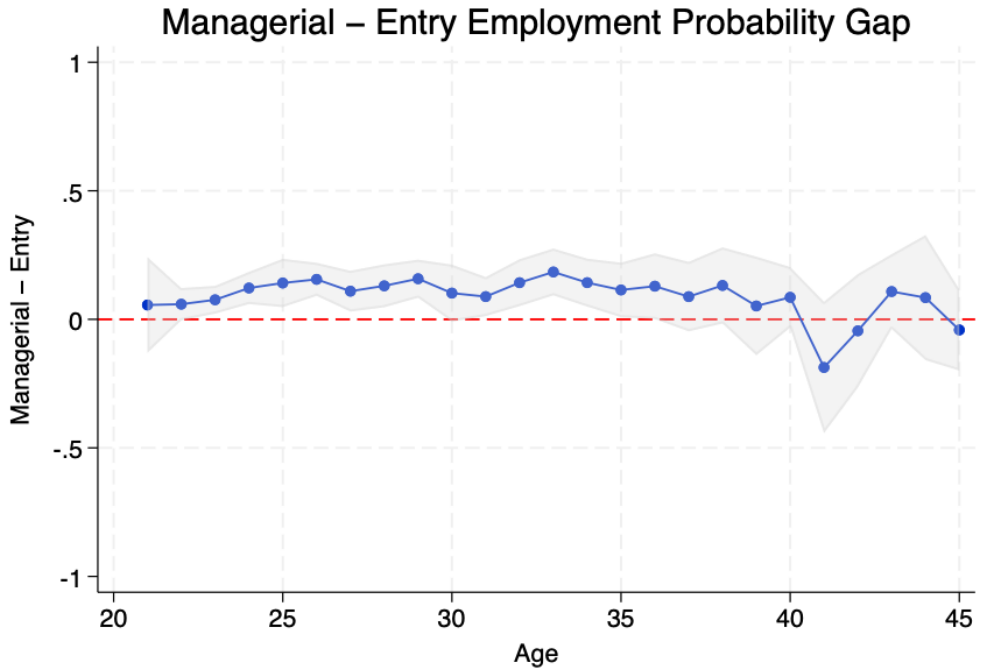

Figure D7: Employment Probability by Age for Workers Applying to Different Jobs

Notes: The figure shows the employment probability gap between applicants to managerial and entry-level jobs by age. The gap is estimated from application-level regressions controlling for age, with standard errors clustered at the applicant level. Shaded areas show 95% confidence intervals.

Our use of prior wages as a proxy for productivity in Section 4.3 raises a natural question: does the adverse selection we document operate through fixed, pre-market characteristics—such as the quality of a worker's education—or through characteristics that evolve and are revealed over the course of a career? To shed light on the question empirically, we decompose prior wages into a component explained by fixed educational credentials and a residual component, and ask how each contributes to the entry–managerial

gap.

We construct the matched sample by linking resume records to college admission cutoff data. We first standardize college names in the resume data and then merge them with a panel of Guangdong college admission cutoffs by standardized college name and inferred admission year. Because the resume data do not contain applicants' province of origin, we assume that all jobseekers took the college entrance examination in Guangdong Province. We therefore use the science-track admission cutoff in the corresponding admission year as the measure of college selectivity. We then construct our primary measure of school quality as the five-year average of within-year standardized admission cutoff scores. [16] The inferred admission year is calculated as the graduation year minus four. Under this procedure, we successfully match 60.65 percent (7,040 of 11,607) of jobseekers to admission cutoff information.

We estimate

$$\ln(PriorWage_i) = \mathbf{X}_i'\theta + u_i, \tag{11}$$

where $\mathbf{X}_i$ collects cross-sectional, pre-market measures of educational attainment: the school quality index (the five-year average of within-year standardized admission cutoffs; see footnote 16 for the percentile alternative), an indicator for attending a top-tier university, gender, field-of-study fixed effects, and graduation-cohort fixed effects. The fitted value $\mathbf{X}_i'\hat{\theta}$ is the *observable component*—the part of prior wages attributable to credentials fixed at labor-market entry—and the residual $\hat{u}_i$ is the *residual component*, which our framework interprets as productivity revealed over the career. We calculate the difference between managerial and entry-level applicants of the same age, and plot each resulting gap by age.

Figure D8 reports the result. The observable component (left panel) is small but consistently positive. The residual component (right panel) is likewise positive across nearly the entire age range, with managerial applicants persistently outscoring entry-level applicants. The reduced sample size due to imperfect matching to educational credentials accounts for the noisiness of the figure relative to the overall gaps. Averaging across ages, however, yields an average gap of 7.4 log points in the observable component and 22.1 log points in the residual component (See Table D2), from pooled regressions with age fixed effects and standard errors clustered at the applicant level. Managerial applicants are thus positively selected on both components, consistent with statistical discrimination on both observable and unobservable dimensions. This patterns shows the sorting we document is not driven solely by fixed pre-market credentials but also reflects productivity revealed over the career.

Because the admission-score series has incomplete coverage for earlier enrollment

[16] As an alternative measure of school quality, we also use the five-year average percentile rank of admission scores. The results are similar to those using the primary measure.

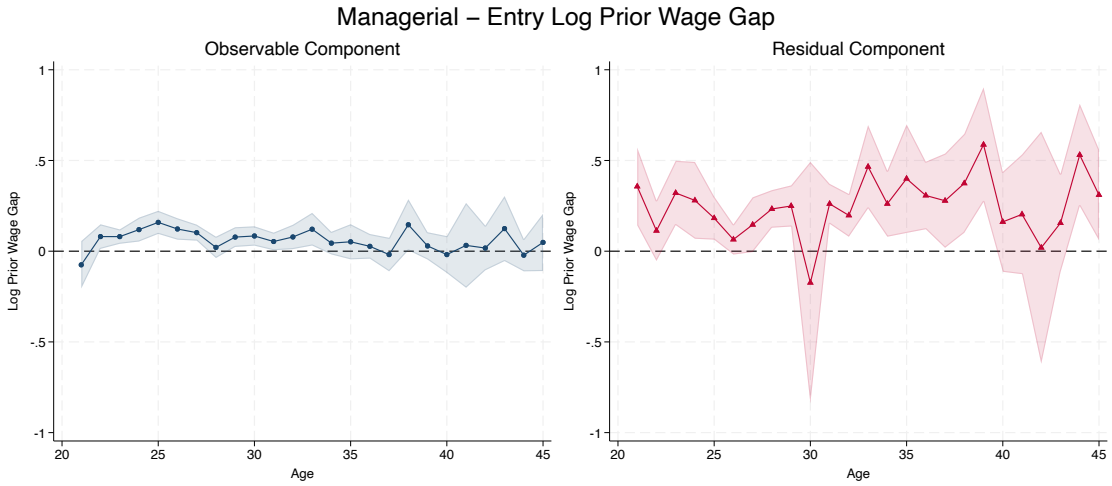

Figure D8: Decomposition of the Managerial–Entry Wage Gap into Observable and Residual Components, by Age

Notes: This figure is based on applicant data. Prior wages are decomposed via Eq. (11): the observable component is the predicted value from a regression of log prior wage on the school quality index (the five-year average of within-year standardized admission cutoffs), gender, a top-tier university indicator, field-of-study fixed effects, and graduation-cohort fixed effects; the residual component is the regression residual. Both components are differenced between managerial and entry-level applicants at each age. Shaded areas show 95% confidence intervals.

years, and older applicants necessarily enrolled earlier, the decomposition drops a disproportionate share of older applicants. Match rates are highest—above 55 percent, and reaching 69 percent and 81 percent for the 2022 and 2023 cohorts—for the graduation cohorts corresponding to the youngest applicants in our sample, and fall toward 40–50 percent for cohorts graduating before the mid-2000s, which correspond to applicants in their late thirties and forties. Estimates at older ages therefore rest on smaller, less representative samples and are correspondingly noisier, as reflected in the wider confidence intervals in Figure D8 beyond age 35.

Table D2: Managerial–Entry Wage Gaps in Prior Wage Components

	Mean Score		Percentiles	
	Observable Component	Residual Component	Observable Component	Residual Component
	(1)	(2)	(3)	(4)
Managerial jobs	0.074***	0.221***	0.075***	0.221***
	(0.008)	(0.034)	(0.008)	(0.033)
First-stage R^2	0.271		0.271	
Age FE	Yes	Yes	Yes	Yes
Observations	152,291	152,291	152,291	152,291
R^2	0.642	0.054	0.642	0.054

Notes: The dependent variables are the observable and residual components of log prior wage, respectively. The results are reported using two alternative measures of school quality: (i) the five-year average of within-year standardized admission cutoff scores and (ii) the five-year average percentile rank of admission scores, based on Guangdong's science-track admission cutoffs from 2015–2019. Standard errors are clustered at the applicant level and reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix 5 Age Distributions

Figures E9 and E10 display the distribution of explicitly posted age upper and lower limits across job postings in the large sample, and Figure E11 displays the age distribution of applicants in the applicant sample. The distribution of upper limits is heavily concentrated at round numbers, with age 35 accounting for approximately 36 percent of all postings that specify a maximum—a pattern consistent with the well-documented "35-year-old phenomenon" in China. Secondary mass points appear at ages 30, 40, and 45. The distribution of lower limits is concentrated at ages 22–25, reflecting the typical graduation age for Chinese college students, with secondary mass at ages 28–30, consistent with firms using experience requirements to screen for workers who have accumulated at least several years of demonstrated performance. Taken together, these distributions illustrate that posted age restrictions are not arbitrarily placed but cluster at career-stage-relevant thresholds that align closely with the theoretical predictions of T_N, T_m, and T_M.

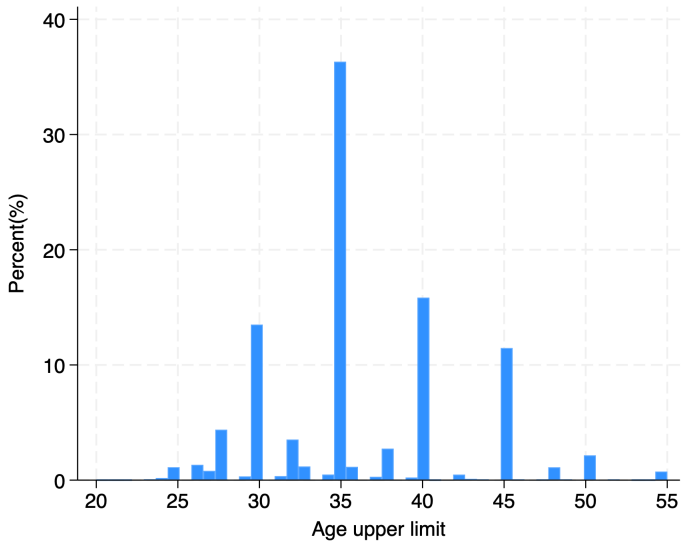

Figure E9: Distribution of Posted Age Upper Limits

Notes: This figure displays the distribution of explicitly posted age upper limits across all job postings in the large sample that specify such a restriction. Age is measured in years. The most common threshold is age 35, which accounts for approximately 36 percent of postings with an upper age limit.

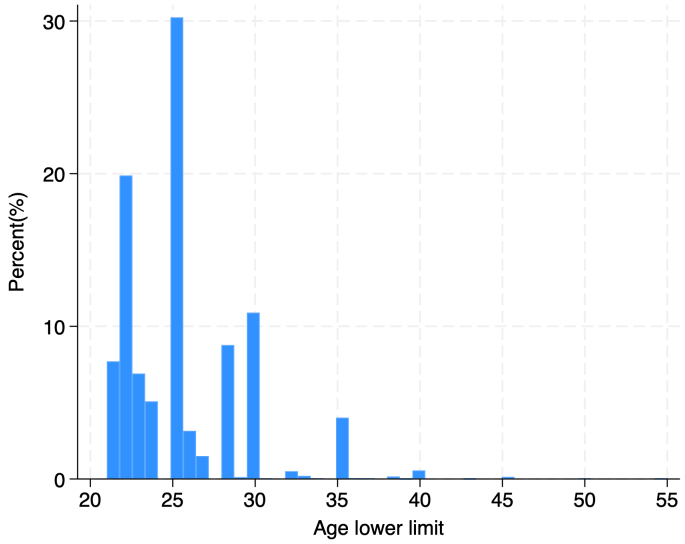

Figure E10: Distribution of Posted Age Lower Limits

Notes: This figure displays the distribution of explicitly posted age lower limits across all job postings in the large sample that specify such a restriction. Age is measured in years. Mass points at ages 22–25 reflect typical college graduation ages in China, while secondary concentration at ages 28–30 is consistent with firms requiring several years of demonstrated labor market performance.

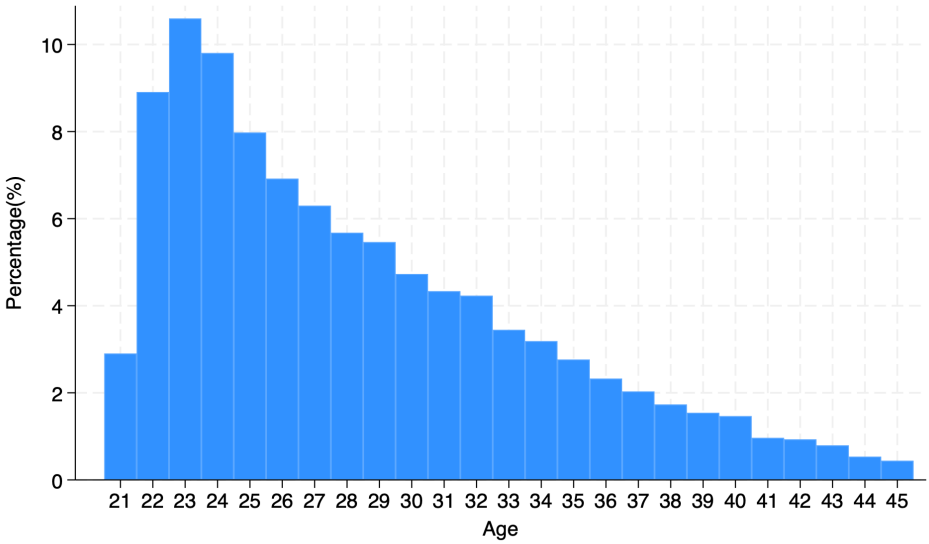

Figure E11: Age Distribution of Applicants

Notes: This figure displays the age distribution of applicants in the applicant sample. Age is measured in years. The distribution is smooth with no bunching at round-number ages commonly associated with posted age maximums (e.g., 30, 35, 40), supporting the validity of the regression discontinuity design.