Screening Women Out?

Pay Transparency in Job Search

Amen Jalal*†

August 31, 2025

Job Market Paper

Very frequently updated. Click **HERE** for latest version.

Abstract: Women sort into low-paying firms for otherwise similar jobs. This paper explains why and offers a policy solution. Large, high-paying firms are more likely to withhold pay information from job ads, and less likely to offer flexibility. Women value flexibility more than men, and apply at lower rates to these firms – but only when pay is not posted in job ads. When pay is posted, women's applications are in fact more responsive to pay than men's. To test whether mandating pay transparency can direct women toward large, high-paying firms, I field an experiment on Pakistan's leading job search platform. The study randomizes mandatory versus optional salary disclosure among ∼20,000 jobs from ∼8,900 firms. The intervention leaves large-firm salaries and amenities unchanged. Yet, women's applications to these firms rise by 95%, and men's by 59%, reversing the gender gap in sorting. I conclude that women do not knowingly "purchase" flexibility with pay. Rather, they turn to flexibility when its price is concealed.

*London School of Economics (LSE); u.jalal@lse.ac.uk.

†I am grateful to G^{2}LM|LIC (IZA) for funding. I thank my advisors Nava Ashraf, Oriana Bandiera, Gharad Bryan, and Robin Burgess for invaluable support. For very helpful feedback, I also thank Tim Besley, Emily Breza, Camille Landais, Alan Manning, and Paula Onuchik. I also benefited from comments by Zoë Cullen, Asim Khwaja, Gabriel Leite Mariante, Maitreesh Ghatak, Marta Guasch-Rusinol, Guy Michaels, Canishk Naik, Steve Pischke, Nina Roussille, Pol Simpson, Kartik Srivastava, and seminar participants at LSE, UC Berkeley, IZA, Stanford, Warwick, HEC Paris, Royal Holloway, Exeter, LUMS, EAYE, Discrimination & Diversity Workshop, Canadian Labor Economics Forum, and Scotland and North England Conference in Applied Microeconomics. Isra Imtiaz and Azka Shahid provided excellent research assistance. The project is registered in the AEA RCT Registry (0013788) and got ethics approval from LSE.

1 Introduction

Female education and employment have been rising globally. Yet, women remain concentrated in small, low-paying, and low-productivity firms. In Pakistan, where this study is based, women's college attainment and employment rates have both doubled since 1990, but gender gaps in salaried and large-firm employment have also widened in parallel. This illustrates a broader pattern: globally, women cluster in low-paying firms despite rising income and education, accounting for 15-50% of the gender pay gap (Card et al., 2016; Sorkin, 2017; Morchio and Moser, 2021; Vattuone, 2023; Bassier and Gautham, 2025). Whether women actively prefer low-paying firms, or face barriers to accessing high-paying ones is key to designing policy, but remains largely unknown.

This paper shows that missing wage information about high-paying firms drives women to low-paying employers. Contrary to the conventional wisdom that women trade wages for non-wage amenities, I find that women are more responsive to pay than men when salaries are disclosed in job posts. But when pay is undisclosed, neither gender can infer it from other sources. While both equally lack information, their responses under uncertainty diverge: men search randomly over the pay distribution, while women avoid higher pay at large firms.

This divergence stems from three features of large firms: they pay more, they are more likely to conceal pay from job-seekers, and they offer fewer family-friendly amenities like flexibility. Thus, pay is hidden precisely when it is high enough to compensate for inflexibility. Since women value such amenities more than men, this combination especially screens them out. However, pay transparency closes gender gaps in sorting without requiring amenities to change. This shows that rather than women preferring flexibility to pay, they do not know the price of flexibility.

To establish this explanation, I start by documenting baseline patterns in sorting, preferences, and pay disclosure using data from over 29 million job applications to $\sim$60,000 firms on Pakistan's largest online job search platform. This is supplemented by firm and worker surveys and a discrete choice experiment to capture job-seeker preferences. These data inform a simple job-search model that formalizes the paper's argument, and delivers testable predictions. I evaluate these predictions through a large-scale field experiment on the platform, randomizing mandatory versus optional salary range disclosure in over 20,000 job posts. These jobs span 8,900 firms and 310,000 applicants.

Conventional data sources provide limited insight into why gendered sorting occurs. Matched employer-employee records only capture final hires – missing the search environment that shapes sorting, and conflating workers' preferences with firms' hiring decisions. Public job posts help reveal search environments, but if pay is not disclosed in these posts, it is inaccessible to workers *and* researchers. This prevents tests of whether information frictions bind in two respects. First, without access to hidden salaries, we cannot test whether disclosure is driven by high- or low-salary jobs and firms. If non-disclosure is random, job seekers may infer hidden wages from posted ones, but if it reflects selection, pay transparency can inform beliefs and shift behaviors. Second, we cannot test whether men and women sort differently on pay depending on its disclosure status.

My data overcome both gaps: I observe workers' consideration sets and applications, and capture salaries even when they are not posted. This is possible because the platform requires firms to report salaries internally even when hidden publicly.[1]

The data reveal a sharp disconnect between real and perceived selection into pay disclosure. Large firms are 45% more likely to offer pay above the market average but 34% less likely to disclose it in job ads – making salaries more likely to be hidden when higher, for otherwise similar jobs.[2] Yet 89% of surveyed job-seekers believe hidden salaries are equal to or lower than posted ones, revealing a clear window for policy action.[3] Large firms also offer fewer family-friendly amenities. They are 55% less likely to provide flexible hours and 57% less likely to offer remote work. Meanwhile, the discrete choice experiment shows that women are 7 percentage points (pp) more likely than men to prefer remote jobs and 6 pp less likely to prefer jobs at large firms.

Men and women respond similarly to salary information, but differently to its absence. When salaries are visible, a log-point increase in pay is associated with a 21% gain in male applications and 26% in female applications. But when salaries are hidden in otherwise similar jobs, this pattern reverses: male applications become unresponsive to pay – reflecting the lack of information – while female applications turn negative. Women's negative coefficient is driven entirely by jobs at large-firms, where a log-point increase in hidden salaries is associated with a 6% decline in applications.

These patterns motivate a model where large firms pay more, offer less flexibility, and conceal salaries, while small firms do the opposite. Men and women value wages equally, but women place more weight on flexibility. When pay is hidden, workers form expectations. Uncertainty reduces the perceived value of undisclosed wages, discounting the appeal of large firms that hide pay. Two testable predictions follow. First, pay disclosure should raise applications to large firms for both genders. Second, the increase should be larger for women, as their stronger preference for flexibility places more of them on the margin where revealed wages just offset the amenity loss. Thus, while the friction facing men and women is the same, the penalty it imposes is gendered.

The field experiment implements the design implied by the model. It removes pay uncertainty to test whether applications shift toward large firms, and whether this shift is larger for women.[4] If revealing wage information has no effect on search, observed gender gaps reflect gendered preferences over non-wage amenities or firm type. If, instead, pay transparency reallocates search, both must be true: women learn new wage information, and are willing to choose pay over flexibility, but only when the trade-off is made transparent. In that case, information can reduce sorting gaps even in the presence of gendered amenity preferences.

[1] Firms have incentives to report truthful ranges. Many purchase a screening service where the platform shortlists relevant applications on behalf of the firm. Additionally, once applications are received, firms can filter applications by whether their expected salary falls within the range it provided for the job.

[2] This is conditional on the job's occupation, city and industry fixed effects, skills, education and experience requirements, schedule, and career-level.

[3] These beliefs do not differ meaningfully by gender.

[4] Both tests were pre-registered.

I find that pay transparency increases applications to treated jobs for both genders by 49%. Effects are strongest where salaries were hardest to learn but are now revealed to be higher – at large firms. In control, these firms are 30% more likely to hide salaries, and their minimum and maximum pay are 32% and 21% higher, respectively, while the range is narrower. Treatment does not affect large-firm salaries, so they remain both higher and more compressed. This drives a differential search response: applications to large-firm jobs surge by 66% (versus 23% at small firms).

The increase in applications is disproportionate for women. Treatment increases male applications by 59% and female by 95%.[5] This reverses the gender gap in sorting: in control, women are 20.4 pp less likely than men to apply to large versus small firms. In treatment, this gap *flips* in favor of women, to 43 pp, highlighting the role of incomplete information in driving gendered sorting.[6]

As women shift toward large firms, they also substitute away from flexibility. To track where in the amenity distribution applications to treated jobs originate, I compare the relative appeal of key amenities within, and then across treatment arms. By developing a natural language algorithm, I classify ads offering remote work, flexible hours, transport support, safe conditions, and equal-opportunity language – features known to matter more for women. Among these, remote work stands out. In control, remote jobs without posted salaries receive 134% more female applications and only 59% more male applications. Yet large firms are 62% less likely to offer remote options.

Salary disclosure does not change the distribution of amenities, but it sharply reduces their appeal. Female responsiveness to remote jobs falls from 134% to 74%. Other amenities also fade in importance as the appeal of large firms rises from zero to 61%. In contrast, male responsiveness to amenities remains stable regardless of the information regime. These results show that information can reallocate women's search toward high-paying firms even if the amenity mix is unchanged.

A potential concern with this analysis is that amenities are not randomly assigned. Thus, unobserved factors may influence both the supply of amenities and application behavior in ways that interact with treatment but are undetected by average treatment effects. To address this, I cross-randomize amenities with salary disclosure in the lab. Each job-seeker evaluates four pairs of stylized jobs in a discrete choice experiment that varies salary visibility, a large-firm label, remote work, transport subsidy, and equal-opportunity language. I then test whether women's preference for flexibility weakens when large-firm salaries are made transparent. When pay is concealed, women are 13 pp less likely to choose on-site jobs at large firms versus similar remote jobs. Once large-firm pay is disclosed, remote work ceases to matter, and choice of large-firm jobs rises monotonically with pay.[7] Thus, rather than women preferring flexibility to pay, flexibility helps offset uncertainty about

[5] Pooling jobs across firm size, disclosure increases male applications by 46% and female applications by 60% (p=0.06 for the gender difference).

[6] Specifically, in control, large firms receive 82% more male and 62% more female applications than small firms. By contrast, in treatment, large firms get 131% more applications than small firms, and 174% more female applications.

[7] Salaries, when visible are randomly drawn between 110-120% of workers' expected wage

pay. In fact, salary transparency itself emerges as the most valued job attribute in this exercise for both genders, increasing women's application odds by a factor of three relative to remote work.

As transparency reshapes search, it may generate firm-side externalities. In baseline surveys, 63% of non-disclosing firms correctly expected applications to rise under transparency, and 72% anticipated average quality to weakly decline (40% expected a strict decline). These beliefs imply two concerns. First, that screening is costly: if identifying strong candidates were easy, firms would maximize applications. Second, that adverse selection may occur, with higher posted salaries attracting less suitable applicants. Moreover, 71% report that transparency may constrain wage-setting, with potential hires anchoring on posted salaries even if they merit less. This shows that high-paying firms hide salaries strategically, to induce self-screening and preserve bargaining flexibility.

Firms' concerns are partly validated: average applicant quality in treated jobs does decline, though modestly. At large firms, the share of applicants meeting skills criteria falls by 6% (no change at small firms), and the education match drops by 1% overall. Hiring also takes 13% longer for treated large-firm jobs, reflecting the higher volume of applications. Similar modest negative effects on average quality are documented in Slovakia's transparency reform (Skoda, 2022).

Despite small declines in average quality, transparency improves the quality of top applicants, especially women. I define high-ability workers as the top decile of a market-wide index of CV characteristics. In control, 90% of jobs already attract at least one high-ability man, but only 57% attract a high-ability woman. Transparency raises this by 11% for women, and only 1% for men. However, lower-ability applications (the bottom decile) also rise – by 2% for men and 7% for women – though these increases are significantly smaller than the gains among women at the top. Similarly, the top three applicants in treated jobs have CVs that are 0.05-0.13 SD stronger on various dimensions, compared to the top three in control. These top-three are also 4% more likely to include a woman.

Firms' post-experiment choices reveal that many over-estimated the costs of pay transparency relative to gains in applicant quality. In the six weeks after the experiment, large never-treated firms continue to hide salaries in 80% of job postings, while fully treated firms are 25 percentage points (30%) less likely to do so. These choices provide a revealed-preference measure of how firms trade off higher screening costs against improved applicant quality.

A central policy question is how these supply- and demand-side changes interact to shape women's hiring outcomes. While pay transparency draws more women to large firms, it also heightens competition, potentially reducing women's hiring prospects. Yet, treated jobs are 7% more likely to download a female CV after reviewing applications on the portal, and they report similar female hiring rates post-experiment. This enables a back-of-envelope extrapolation: a representative woman is 18% more likely to be hired by a large firm in treatment than control – not because treated large firms become more likely to hire women, but because women become more likely to apply.

A remaining puzzle is why workers fail to infer that hidden salaries at large firms are higher, or that high-amenity jobs have lower pay. First, available signals are misleading. When large firms

disclose pay, it is usually for lower-paying positions. Similarly, amenities are often positively correlated with wages due to confounding by unobserved job characteristics, even though theoretically they compensate for lower pay (see Lavetti (2023) for a review of the issue). Second, averages often mask noise and heterogeneity. Song et al. (2019) show that wage dispersion is greatest at the largest firms, where top talent coexists with automated tasks, and steep tenure returns depress entry wages. Third, even if signals point in the right direction, inference is cognitively demanding. Learning requires observing many posts, aggregating information, correcting for selection, and updating beliefs without feedback from the jobs workers did not get. Unsurprisingly, studies consistently find that workers misperceive pay even within their own firms and misjudge outside options even when they switch jobs frequently (Card et al., 2012; Cullen and Perez-Truglia, 2022; Jäger et al., 2024).

Could other mechanisms explain gendered sorting instead? I find that none appear sufficient to explain the observed patterns. There is no evidence that women are less informed about large-firm wages, less able to identify firm size, or concerned about discrimination at large firms. Gender gaps in network size or composition, and risk or ambiguity aversion also do not explain results. However, women are 29% more likely to prefer fixed salary offers over bargaining, consistent with Leibbrandt and List (2015). These bargaining averse women are also 17% more likely to apply to treated jobs. Nonetheless, bargaining aversion alone cannot explain why non-disclosure affects women differently by firm size: when pay is hidden, women randomize applications over the pay distribution for small-firm jobs but avoid high-wage jobs in large firms. A simple model extension accommodates bargaining aversion, which simply amplifies the cost of pay non-disclosure in inflexible jobs.

Contributions. A large literature following Card et al. (2016) has documented that women sort into low-paying firms, widening the gender pay gap. Yet the mechanisms remain poorly understood. This paper offers an explanation: high-wage firms are more likely to withhold salary information and less likely to offer family-friendly amenities. Women, who place greater weight on flexibility than men, tend to avoid these firms.

The paper brings together various strands of research usually studied in isolation, and shows how their interaction generates gendered sorting, while also advancing each strand individually. First, hypothetical choice experiments find that women value flexibility more than men (Mas and Pallais, 2017; Wiswall and Zafar, 2018), but Mas and Pallais (2020) note that these survey valuations diverge from observed sorting, with women often forgoing flexibility in practice, likely because it is bundled with less desirable attributes such as lower wages. I test this explanation: I replicate the gendered preference for flexibility in a discrete choice experiment, and then show that behavior diverges in the field when marginal preferences interact with constrained trade-offs and incomplete information. Specifically, when high-wage, high-amenity jobs are unavailable and high wages become newly visible for inflexible jobs, workers with a stated preference for flexibility shift toward those jobs. This implies that if lab studies vary wages and amenities jointly, reflecting their covariance in observational data, they may better approximate field behavior.

Second, research shows that high-paying firms often offer amenities that may be unappealing to women such as low flexibility, long hours, and competitive work environments (Blau and Kahn, 2017; Goldin, 2014; Cullen and Pakzad-Hurson, 2023; Niederle and Vesterlund, 2007; Niederle et al., 2014). This paper extends that literature to a low-income setting, showing how the supply of one amenity – remote work – shapes sorting in a context where women's mobility is highly constrained. Prior work in South Asia has responded to this same constraint, pursuing supply-side solutions: Garlick et al. (2025) brings women to work by offering transport in Pakistan; Ho et al. (2024) delivers work to women by creating new remote jobs for them in India. These approaches imply that for women to participate in the labor market, the nature of jobs – or associated amenities – may need to change. I show an alternative pathway: improving transparency around pay can reallocate women's search toward high-paying firms without changing the amenity infrastructure.

Third, a well-established literature shows that large firms offer wage premiums (Lucas, 1978; Moore, 1991; Troske, 1999). This paper shows that they are also more likely to withhold that information from job-seekers. Although existing work links wage disclosure to job or occupation characteristics (Banfi and Villena-Roldán, 2019; Batra et al., 2023), I document that variation in disclosure is primarily driven by firms. These findings reveal a new channel through which firms' recruitment and compensation practices can—often inadvertently—widen gender gaps, joining other documented practices such as flexible pay structures (Biasi and Sarsons, 2022), eliciting salary expectations and histories (Roussille, 2023; Hansen and McNichols, 2020), gender-targeted job ads (Kuhn and Shen, 2013), and subjective "potential"-based promotion criteria (Benson et al., 2021).

Finally, the paper contributes to the emerging literature on salary transparency mandates. Prior studies evaluate such policies observationally in Slovakia (Skoda, 2022), Austria (Frimmel et al., 2022), and the U.S. (Arnold et al., 2022). I advance this work in four ways. First, I present the first randomized controlled trial of pay transparency mandates, using a saturation design that identifies both direct and spillover effects. Second, I replicate and extend existing evidence in a new low-income setting: like existing work, posted salaries in Pakistan rise by about 3% on average while ranges narrow, and applicant quality declines modestly. Third, I overcome a major challenge in prior work, noncompliance with the mandate, that leaves undisclosed salaries unobserved by researchers (e.g., in Colorado, 30% of postings ignored disclosure requirements). Using internally reported salaries, I document novel patterns of firm selection into disclosure. Last, I extend the focus to gender, showing that transparency nearly doubles women's applications to high-paying firms.

The paper is organized as follows: Section 2 describes the setting and Section 3 the data, Section 4 documents facts motivating this research, Section 5 provides a theoretical framework, Section 6 outlines the design of the field experiment, Section 7 lists the empirical approach, Section 8 presents the results, Section 9 replicates the results in a lab setting, Section 10 discusses alternative supply-side mechanisms, Section 11 summarizes demand side impacts, Section 12 provides back-of-envelope estimates of women's hiring likelihood, and finally, Section 13 concludes.

2 Context

2.1 Pakistan

This research is set in Pakistan, which ranks last among 148 countries in the Global Gender Gap Index, reflecting the world's widest gender disparities in economic, educational, health, and political outcomes (World Economic Forum, 2025).[8] Female labor force participation is just 24.4% – far below countries at similar income levels – while, male labor force participation is at 81.0% (ILO, 2025). Even among college-educated women, employment rates barely exceed the national average (Bandiera et al., 2025). Inequality is evident not only at labor force entry but also within employment. Pakistan records one of the largest raw gender pay gaps globally at 34%, second only to India (34.5%), and the highest adjusted pay gap at 36.3% (ILO, 2018).[9]

Yet, Pakistan has made substantial progress in women's education and employment over the past three decades. As Appendix Figure A.1 Panel (a) shows, the female share of employment increased from about 14% to 23%, while the share of college-educated women rose from just 3% to 13%. Panel (b) shows that these gains have nearly halved gender gaps: male-to-female ratios fell from 6.0 to 3.4 in employment and from 2.5 to 1.3 in college education.

However, these gains in education and employment have not translated into access to better jobs. In low-income countries, casual wage and self-employment often reflect labor rationing, and provide lower returns and fewer formal protections (Breza et al., 2021). Panel (c) of Appendix Figure A.1 shows that despite gains in women's overall employment, the male-to-female ratio in large-firm employment more than doubled from 0.7 to 1.5, and the gap in salaried versus casual wage or self-employment widened from 1.4 to 2.0. In short, there has been substantial progress in whether women work, and in the human capital they bring to work, but less so in where their work takes place.

2.2 Job platform

I partner with the largest job search platform in Pakistan – bigger than its next five competitors combined. Over the course of my baseline data period (2019-2024), it facilitated 29.5 million job applications.

Firms. There are 90,640 firms registered on the platform, of which 69% posted at least one job in the duration covered by my data. The average firm posted 3.6 job ads. The platform hosts a diverse range of firm sizes: 22% have fewer than 10 employees, while 17% have more than 300.

[8] The Global Gender Gap Index measures gender parity across four dimensions: labor force participation and opportunity, educational attainment, health and survival, and political empowerment.

[9] The adjusted pay gap accounts for differences in education, age, employment type (full- vs. part-time), and sector (public vs. private).

Jobs. The platform receives over 100 new job postings daily, amounting to 326,145 job ads from January 2019 to April 2024. The platform features a broad range of jobs, spanning 63 industries and 63 occupations. Appendix Table B.2 highlights the top 10 industries and occupations, while Appendix Figure A.3 illustrates the diversity of job titles available. While the job mix skews toward white-collar occupations, there is substantial variation coverage of low-paying jobs, with about 27% of listings offering salaries at or below the monthly minimum wage.

Job-seekers. About 3.1 million job-seekers are registered on the platform in this period. The user base is younger and more educated than the national average, with a median age of 27 and 79.2% holding a Bachelor's degree or higher. Women represent 26.9% of all job-seekers, aligning with labor force participation patterns nationally. Although the set of workers and jobs the platform hosts is not nationally representative, it offers an unparalleled window into recruitment in Pakistan. Importantly, it features large variation in firm size and salary levels, a sizable population of female job-seekers, and a pool of applicants most relevant to large, high-paying firms. These characteristics make it an ideal setting to examine why women disproportionately avoid such firms.

2.3 Salary disclosure in job ads

Global context. Salary is a fundamental attribute of jobs, yet it is often missing from job ads globally. As shown in Figure 1, salary non-disclosure is widespread across countries, regardless of regional or economic context. One notable exception is Slovakia, where 83.2% of ads include salary information, following a legal mandate requiring disclosure of the minimum salary (Skoda, 2022).

Salary transparency laws. In recent years, there has been a large global shift toward pay transparency laws mandating employers to state salary information in job ads. In the United States, as of 2024, 14 states have enacted such laws. In March 2023, the Salary Transparency Act was also introduced in Congress, which if passed, would require US employers across the country to disclose salary ranges in job ads. In Europe, several countries including Austria, Denmark, Iceland, Finland, Germany, Norway, and Sweden have implemented similar measures, while the United Kingdom is piloting the reform with a group of volunteering firms. These reforms are part of a global shift toward pay transparency: over the past two decades, 71% of OECD countries have adopted laws addressing some aspect of salary disclosure (Cullen, 2024). Measures range from protecting workers who share pay information with colleagues from being penalized, to requiring firms to disclose internal pay data, including reporting of gender pay gaps. As such policies proliferate, assessing their effectiveness and broader economic consequences has become increasingly important.

3 Data and Descriptive Statistics

Administrative data. The job platform generates data on registered job-seekers, job postings, firms and applications. Job characteristics are described in Appendix Table B.1. The table shows that visible salaries range from a minimum of PKR 42,517 (SD: 43,895) to a maximum of PKR 68,183 (SD: 67,561), while hidden salaries are generally higher, with a minimum of PKR 56,152 (SD: 61,529) and a maximum of PKR 98,143 (SD: 107,424).[10] The average job requires less than one year of prior experience (Mean: 0.90 years, SD: 1.89). Job postings attract an average of 75.5 applications (SD: 270.02) and the average job requires around three skills (Mean: 3.25, SD: 2.44). A majority of positions are full-time (94%) with daytime schedules (79%). Only 5% of jobs mandate screening questions as part of the application process. Most jobs target mid-career professionals (65%), followed by entry-level positions (29%), while a small fraction are available for interns or students (4%) and department heads (1%). Bachelor's degrees are the most common educational requirement, listed for 65% of the jobs. Firms that post jobs vary substantially in size, with small businesses (1-50 employees) accounting for about half the postings (54%).

Job-seeker characteristics are described in Appendix Table B.3. The average age is 28.1 years (SD: 6.93), and candidates typically possess 3.27 years of work experience (SD: 5.01). Academic performance, measured by GPA, averages 2.85 (SD: 0.80). Job-seekers submit around 8.84 applications (SD: 46.14). The workforce skews male, with women comprising 27% of job-seekers. Job-seekers' average current salary is PKR 45,998 (SD: 52,795), while expected salaries are slightly higher, averaging PKR 49,441 (SD: 54,151). Mid-career professionals make up the largest group (41%), followed by entry-level candidates (29%) and interns or students (26%). Senior-level workers, such as department heads (3%) and executives (1%), represent a smaller fraction of job-seekers. Over half of job-seekers hold a Bachelor's degree (53%), with an additional 26% possessing a Master's degree. High school or less accounts for a small share (2%), while 7% hold a certificate or diploma, and 1% have earned a PhD.

Appendix Figure A.2 highlights large gender gaps even in this highly educated sample. Women are 10.9 pp (47.6%) more likely than men to hold a Master's degree, yet 13.9 pp more likely to be confined to entry-level positions. Conditional on occupation, they earn 25% less than men and expect to earn 19% lower salaries. Women also apply to fewer jobs (7.4 fewer on average, or 49.1% less) than men, and the maximum posted salary of jobs they target is 17% lower than for men. These differences underscore how application behavior may drive downstream gender gaps in employment and pay.

Baseline firm surveys. I conduct baseline surveys with key decision-makers in over 300 firms to understand how firms make salary disclosure decisions. These firms are randomly sampled from the

[10] For reference, the national minimum wage is PKR 37,000 as of July 2024.

platform's registered list, conditional on the firm having posted at least one job in the past year and a minimum of five in the past three years. The sample is stratified such that two-thirds of the firms hide salaries while one-thirds do not. Respondents are well-positioned to speak to hiring practices: 81% of respondents have the authority to unilaterally modify their firms' job ads, 90% are involved in setting salaries, and the median respondent has 6 years of hiring experience. A third of respondents are from human resources, 41% are from general administration (e.g., CEOs), and the remainder are department heads. Although the response rate is low (15%), it is not systematically different across firm size, industry, salary non-disclosure frequency, market share of jobs and applications, mean salary, female share of applicants, and other characteristics (see Appendix Figure A.5). The survey covers information about why firms choose to disclose or hide salary ranges, and their beliefs about how application volumes and average quality will change if they reverse their disclosure decisions.

Endline firm surveys. Following the conclusion of the field experiment, I conduct endline surveys with firms exposed to the intervention. The survey covers 4,146 firms – 46.6% of those included in the experiment. As described in the baseline survey section above, respondents are firm representatives registered on the platform who tend to be knowledgeable about, or directly involved in hiring and wage-setting decisions. To limit respondent fatigue and minimize differential attrition, firms that posted more than five jobs were asked about a random subset of five, resulting in data on 5,685 jobs (28.3% of all experiment jobs). The response rate does not differ by treatment status. The survey was intentionally concise to maximize the response rate, asking only about the number of candidates hired, the number of women hired, and the time taken to fill a vacancy.

Endline job-seeker surveys. I field an online survey with 579 job-seekers from the platform to investigate the mechanisms underlying their application behavior observed during the experiment. A total of 15,492 job-seekers (5% of all applicants) were invited via email and given one week to take the survey. To incentivize participation, respondents were offered a reward of PKR 5,000 (USD 17.69). The invitation list was randomly drawn from the approximately 310,000 job-seekers who were active during the experiment, defined as having submitted at least one application. Among respondents, two-thirds completed the full survey, with an average completion rate of 78.1%. The survey captures job-seekers' beliefs, preferences, and behaviors related to salary disclosure and application decisions. It includes modules that assess the accuracy of respondents' beliefs about hidden salaries, application volumes (i.e., perceived competition), and firm size – benchmarked against platform administrative data. It also measures preferences over risk, bargaining, and ambiguity; the latter is elicited through a standard incentivized ambiguity aversion game following Ashraf et al. (2009). To examine preferences over job attributes, the survey includes a discrete choice experiment completed by 438 respondents (56% women). In this survey experiment, respondent evaluates

four job pairs (eight vignettes total), where salary visibility, large-firm labels, remote work, and equal-opportunity language are independently randomized.

4 Motivating facts

Firms play a central role in shaping wage inequality: similar workers earn markedly different salaries depending on their employer (Goux and Maurin, 1999; Gruetter and Lalive, 2009; Card et al., 2013, 2018). Among firm-level characteristics, size stands out as a central force in shaping the organization and compensation of labor. A long-standing literature documents a robust "firm-size wage premium" – the empirical regularity that larger firms pay more to observationally similar workers, across countries and over time (Lucas, 1978; Moore, 1991; Oi and Idson, 1999; Colonnelli et al., 2018). In the U.S., for example, workers in firms with 500+ employees earn wages more than 30% higher than those in firms with fewer than 25 workers – a disparity as large as the gender wage gap. This premium is only partly explained by differences in sorting and firm characteristics (Troske, 1999), suggesting that large firms not only attract talent but also help produce it, through greater investments in training and firm-specific human capital. These patterns make large-firm jobs especially valuable. Thus, understanding access to them is key to explaining labor market inequality. Building on this background, the following facts show how large firms' recruitment practices interact with workers' job search behavior, shaping gender gaps in sorting across firms.

Large firms pay more for otherwise similar jobs. I show that the firm-size premium extends beyond workers to jobs: higher-paying positions are disproportionately concentrated in large firms. Nearly half (45%) of jobs with above-market-average salaries are posted by firms with at least 50 employees. I adopt this 50-employee threshold to define "large firms" because it corresponds to the median firm size in the job-level data and frequently marks the point at which labor regulations begin to apply, though the results are robust to alternative definitions. Importantly, the concentration of high-paying jobs in large firms emerges prior to any sorting of workers across firms or the creation of surplus through firm–worker complementarities, underscoring that the premium is rooted in jobs themselves. Jobs posted by large firms offer 23.1% higher minimum and 17.5% higher maximum salaries than those posted by smaller firms. To isolate size-specific pay differences from other job characteristics, Figure 3 Panel (a) residualizes salaries on a rich set of job-level controls, including fixed effects for occupation, industry, city, required education and experience, career level, and work schedule. The resulting measure of excess pay remains strongly correlated with firm size: relative to small firms (1-10 employees), firms with over 300 employees pay about 20% more for otherwise similar jobs.

Large firms are less likely to disclose salary information to job-seekers. Figure 3 Panel (b) shows that the probability a job ad omits salary information rises sharply with firm size. While only 39% of job postings from the smallest firms (1-10 employees) hide salary, the rate climbs to 75% among firms with 301-600 employees and remains high at 72.6% for firms with over 600 employees. This pattern suggests that salary non-disclosure is the most common among the largest and highest-paying firms.

Salary disclosure is a firm, not job attribute. The firm-level patterns above suggest a notable asymmetry: the jobs that offer the highest pay are also most likely to conceal it. That this reflects firm behavior – larger firms both pay more and disclose less – is one interpretation. An alternative explanation is that jobs with hidden salaries differ systematically from those with transparent pay, and are simply more common in large firms. Appendix Figure A.7 helps distinguish between these channels by leveraging a novel feature of the platform: firms must report salary ranges internally but can choose whether to display them to job-seekers.[11] While previous research has explored which types of jobs are more likely to hide salary information, this feature allows me to go further by examining whether salary levels rise with non-disclosure even within jobs of similar types.[12] Surprisingly, they do. For otherwise similar jobs, salaries are more likely to be hidden when they are higher: a one-log-point increase in the minimum, midpoint, or maximum salary is associated with a 2.6, 8.6, and 9.4 percent increase in the probability of non-disclosure. While job-level traits such as career level, education, and experience requirements do predict non-disclosure,[13] they explain only a small share of the overall variation. Figure 2 shows that a rich set of job characteristics accounts for no more than 12.7% of the variation in disclosure status. In contrast, firm fixed effects alone explain 53.4% of this variation. Appendix Figure A.4 further confirms this result using the bayesian information criterion. Together, these results point to the firm, not the job, as the dominant source of variation in salary disclosure.

Job-seekers misperceive selection into pay disclosure. Job-seekers systematically misperceive the relationship between pay and transparency. In the job-seekers survey, 89% of respondents believe that hidden salaries are equal to or lower than disclosed ones. This creates a clear opportunity for pay transparency to correct beliefs and redirect search. Large firms play a central role in this

[11] Firms have incentives to report accurate salary ranges to access benchmarking tools and filter applicants by expected or current salary.

[12] Job-level traits do predict salary disclosure status. In particular, as jobs become more advanced in career level and require higher education and experience, they are more likely to conceal salary information from job-seekers. Appendix Figure A.6 illustrates this pattern: as experience and education requirements rise, the likelihood of salary non-disclosure also increases. Similarly, jobs at higher career levels are less likely to disclose salary ranges. These trends are consistent with findings from the U.S. and Chile (Batra et al., 2023; Banfi and Villena-Roldán, 2019).

[13] Jobs requiring more education and experience, or situated at higher career levels, are more likely to conceal salary information. See Appendix Figure A.6. These patterns align with prior findings from the U.S. and Chile (Batra et al., 2023; Banfi and Villena-Roldán, 2019).

distortion. Not only are they significantly less likely to disclose salaries overall, but their tendency to conceal increases steeply with salary: in large firms, a one-log-point increase in the posted salary is associated with a 9.0 pp rise in the probability of non-disclosure, nearly double the 4.6 pp increase observed in small firms. This means that job-seekers exposed only to disclosed salaries are likely to form pessimistic beliefs about large-firm compensation. Thus, a salary transparency intervention may deliver its largest belief correction – and potentially its strongest behavioral effect – among high-paying, large-firm jobs.

Large firms also offer fewer 'family-friendly' amenities. As firms grow larger, employees may have less leverage to negotiate flexible work arrangements, and may face more stressful and demanding work environments (Blau and Kahn, 2017; Goldin, 2014; Cullen and Pakzad-Hurson, 2023; Niederle and Vesterlund, 2007; Niederle et al., 2014). Yet such flexibility can be particularly valuable to women (Mas and Pallais, 2017; Wiswall and Zafar, 2018), making it important to examine how access to these non-wage amenities varies with firm size. For a limited time prior to my study, my partner job platform gathered information about non-wage amenities in job ads. Appendix Table B.4 uses this baseline data to study the correlation between availability of such amenities and firm size, controlling for job characteristics such as industry and occupation fixed effects, education and experience requirements, stated gender preferences in job ads, and the career level of the job. It shows that large firms are 26 pp (58%) less likely to offer flexible work hours, 8 pp (57%) less likely to offer remote work, 7 pp (39%) less likely to offer transport support for commuting. Through text analysis of job descriptions in my experimental sample, I replicate these patterns in Section 8.4.

Women value flexibility more than men. To examine gendered preferences over job amenities, I embed a discrete choice experiment (DCE) in the job-seeker survey. Each respondent evaluates four pairs of stylized job postings that vary randomly along five attributes: salary visibility, large-firm label, remote work, transport subsidy, and inclusive language ("equal opportunity employer"). I estimate preferences using a conditional logit model, where the probability that worker w selects job alternative j from choice set i is given by:

$$\Pr(\text{choice}_{wij} = 1) = \frac{\exp(\boldsymbol{\alpha}' \mathbf{X}_{wij})}{\sum_{k \in C_i} \exp(\boldsymbol{\alpha}' \mathbf{X}_{wik})}$$

Here, $\mathbf{X}_{wij}$ is a vector of job attributes (e.g., salary visibility, large-firm label, remote work, transport subsidy, inclusive language), and $\boldsymbol{\alpha}$ is a vector of corresponding preference weights. The observations are at the worker $\times$ job alternative level, within a binary choice set C_i. Table 1 reports the resulting coefficients, separately by gender. Women place significantly greater value on flexible work arrangements: they are 6.51 pp more likely than men to choose remote jobs. Conversely, women are 6.24 pp less likely to prefer jobs at large firms. By contrast, gender gaps are small

and statistically insignificant for other amenities, including salary visibility, inclusive language, and transport subsidies.

Women are as responsive to salaries as men, but only when revealed. These patterns raise a key question: do men and women respond differently to the combination of high pay and low transparency, especially when it is concentrated in large, low-flexibility firms? To answer this, I turn to application behavior and show that while women are just as responsive to salary as men when it is disclosed, they react very differently when it is not. Figure 4 examines how job applications vary with salary by gender, depending on whether the wage is disclosed. Application counts are normalized by gender-specific means to account for the fact that women comprise only 27% of job-seekers in the sample and submit fewer applications overall. All estimates control for industry, occupation, and city fixed effects, as well as detailed job characteristics such as experience and education requirements, career level, number of vacancies, required skills, and job schedule. Panel (a) shows that men and women respond similarly to wage information: when salaries are visible, both apply more to higher-paying jobs, and women respond even more strongly than men (26% vs. 21% increase in applications per log-point increase in wage). However, when salaries are hidden, this pattern diverges. Women's application rate flips sign, declining with wage, while men's becomes flat. Thus, men appear to apply indiscriminately across the wage distribution when pay is unknown, consistent with random or undirected search. In contrast, women expressly avoid high-paying jobs when the wage is hidden: a log unit increase in hidden pay corresponds to a 4% decline in application likelihood. This slope is different from men's at $p<0.01$. Panel (b) shows that this negative relationship for women is entirely concentrated in large firms. Among large-firm jobs with undisclosed salaries, a log-point increase in pay is associated with a 6% drop in applications from women, while no such pattern is observed at small firms. Men again show no wage-responsiveness under opacity, regardless of firm size. The divergence in behavior under salary opacity highlights a potential source of allocative inefficiency, where the absence of wage information leads women to self-select out of high-paying roles they may have otherwise pursued.

5 Theoretical Framework

This section formalizes how lack of salary transparency interacts with gender-specific amenity preferences to shape sorting across large and small firms. The primitive forces underlying this model come directly from the empirical facts in Section 4.

5.1 Jobs

There are two job types, indexed by $j \in \{L, S\}$ where L is a job at a large firm and S is a job at a small firm. There are three key differences between jobs of type L and S: large firms pay higher wages, provide fewer family-friendly amenities, and conceal pay. These are summarized in the table below:

$w_L > w_S$	Large firms pay more
$a_L < a_S$	Large firms offer fewer amenities
$d_L = 0,\ d_S = 1$	Large firms hide wage; small firms disclose

Here $w_j > 0$ is the salary and a_j is a scalar amenity index, higher values indicating more flexibility. The binary indicator d_j equals 1 when the wage is disclosed (visible to job-seekers) and 0 when it is hidden. I adopt the assumption that L always conceals and S always reveals pay for two reasons. First, it restricts focus to the tension that empirically matters most: high wages coincide with low transparency. Second, it keeps the exposition simple by yielding easy to visualize cut-off rules. However, this simplification is without loss of generality. Allowing each firm type to disclose with probabilities $\Pr(d_L = 1) = \pi_L < \pi_S = \Pr(d_S = 1) \in (0, 1)$ leads to the same comparative-statics and predictions.

5.2 Workers

A continuum of workers chooses where to apply. Each worker is characterized by their gender $g \in \{f, m\}$ and an amenity preference parameter $\gamma \geq 0$. Conditional on gender, γ is distributed continuously according to the cumulative distribution function $P_g(\gamma)$ with probability density $p_g(\gamma)$. Women have greater taste for flexible amenities:

$$\gamma_f \ \text{ first-order stochastically dominates } \ \gamma_m.$$

Workers of both genders value wages equally, with constant relative risk aversion (CRRA) $\theta \in (0, 1)$. Gender neutrality in θ is supported by two facts. First, when wages are disclosed, the responsiveness of male and female applications is quantitatively similar (26% versus 21%; Figure 4). Second, measured gender differences in risk and ambiguity aversion are statistically insignificant (Columns 1-2 of Table B.19).

Utility from a job offering wage w and amenity level $a \in (0, 1)$ is:

$$U_g(w, a; d) = \begin{cases} w^\theta \, a^{\gamma_g} & \text{if } d = 1 \quad \text{(salary visible)}, \\ \mathbb{E}[w^\theta] \, a^{\gamma_g} & \text{if } d = 0 \quad \text{(salary hidden)}, \end{cases}$$

where the expectation is taken over the unconditional wage distribution $F(w)$:

$$\mathbb{E}[w^\theta] \;=\; \int w^\theta \, dF(w).$$

When the salary is posted, a job-seeker observes the exact wage w and derives its full utility from $w^\theta a^{\gamma_g}$: the standard CRRA transformation of income, and the value they attach to a job's amenities. When the salary is not posted, the worker can no longer condition on w as it is now unknown. They therefore evaluate the job at the certainty-equivalent wage, $\mathbb{E}[w^\theta]$. This expectation is taken over the distribution of market wages that workers plausibly know from observing vacancies that disclose pay, as well as their own and peers' earnings history.[14]

The function $x \mapsto x^\theta$ is strictly concave because $\theta \in (0, 1)$. Thus, the Jensen's inequality implies:

$$\mathbb{E}[w^\theta] \;<\; \big(\mathbb{E}[w]\big)^\theta.$$

i.e., the expected utility of the wage is lower than the utility of the expected wage.

The model makes an implicit assumption that workers have correlation neglect: they observe job attributes but do not infer across them, so U_g depends only on (w, a, d) and not directly on j; and hidden wages are evaluated from the overall distribution rather than from j or a. This is motivated by incentivized belief elicitation in survey data: workers of both genders correctly recognize in 73% of cases when a posting is from a large firm (Column 2, Table B.14), yet 89% believe that hidden salaries are the same as or lower than posted ones, despite the fact that large firms both pay more and disclose less. Moreover, when asked to guess undisclosed wages, workers are inaccurate regardless of firm size and gender (Table B.12).[15] Similarly, wage expectations do not depend on amenities: while theory often treats wages as compensating for disamenities, empirically wages and amenities are often positively correlated because of unobserved confounders (Lavetti, 2023). Hence, the direction in which amenities should inform wage beliefs is ambiguous.

[14] Using a common prior keeps the algebra simpler. The model can be extended – while preserving the core insights – to allow for biased beliefs where workers weight observed wages differently, or update from personal networks. These would amplify the benefit of transparency for relevant subgroups, without altering its sign.

[15] This is consistent with a broader literature showing that workers have limited information about pay: they lack knowledge of wages even within their own firms and are poorly informed about outside options even when they switch jobs frequently (Card et al., 2012; Cullen and Perez-Truglia, 2022; Jäger et al., 2024). One reason may be that while large firms offer higher average pay, they also exhibit considerable wage dispersion, particularly at the point of entry (Song et al., 2019).

5.3 Sorting under different information regimes

Imagine a worker can only apply to one job. They would then choose to apply to the job that yields higher utility. In this world of constrained application choices, we can lean on the model to predict how salary transparency may affect sorting of male and female workers across small and large firms. Let (w_L, a_L) and (w_S, a_S) denote, respectively, the wage and amenity bundles offered by a large-firm job L and a small-firm job S.

Control regime: wage hidden at L and revealed at S. Absent a wage disclosure mandate, large-firm jobs do not reveal salaries. When salaries are not disclosed, the worker must compare a known wage at the small firm to an uncertain wage at the large firm, subject to the Jensen penalty. This discounts the appeal of large firm salaries, as illustrated in the upper panel of Figure 5: the gray line L_C (large firm under control) lies below the full-information utility that would be attainable from large firms under transparency (the coral L_T line).

Workers prefer the large-firm job if the expected utility from its (hidden) wage and lower amenities exceeds that of the small-firm job:

$$\mathbb{E}[w^\theta] \cdot a_L^{\gamma_g} \;\geq\; w_S^\theta \cdot a_S^{\gamma_g}.$$

Rearranging and taking logs:

$$\log \mathbb{E}[w^\theta] - \theta \log w_S \;\geq\; \gamma_g \cdot (\log a_S - \log a_L).$$

Let $\Delta w \equiv \log(w_L/w_S)$ be the wage gap between large and small firms, $\Delta a \equiv \log(a_S/a_L)$ be the amenity gap between small and large firms , and $\kappa \equiv \theta \log w_L - \log \mathbb{E}[w^\theta]$ be the uncertainty penalty reflecting the difference in a transparent large firm salary and its expected utility under opacity, which is positive for $\theta \in (0, 1)$.

Using these, the inequality above can be rearranged into a threshold rule:

$$\gamma_g \;\leq\; \gamma^C \;\equiv\; \frac{\theta \Delta w - \kappa}{\Delta a}. \tag{1}$$

This defines a cut-off γ^C that partitions workers by their valuation of amenities. In the top panel of Figure 5, this threshold appears as the left dashed vertical line labeled γ_C^*. The intersection of this line with the utility functions shows where workers are indifferent between large and small firms under the control regime. Those with low amenity demand ($\gamma_g \leq \gamma^C$) choose the large-firm job despite its hidden wage, while those with higher amenity demand ($\gamma_g > \gamma^C$) prefer the more flexible small-firm option. This cut-off depends on three forces: it rises with the wage premium large firms

offer over small firms (Δw), and falls with the amenity gap between small and large firms (Δa), as well as the extent of the uncertainty penalty (κ).

When the wage gap between large and small firms becomes arbitrarily large ($\Delta w \to \infty$), the cut-off diverges to infinity, and every worker, regardless of how much they care about flexibility, selects the large-firm job. But as the wage gap narrows, γ^C moves leftward, and amenity preferences begin to have bite. Fewer workers are willing to forego flexibility for a now-smaller, and still uncertain, wage advantage. Eventually, as the wage gap shrinks to the point where $\Delta w = \kappa/\theta$, the cut-off collapses to zero. In this case, only the worker with $\gamma_g = 0$ is indifferent between the two jobs; all others strictly prefer the small firm. If the wage gap shrinks further ($\Delta w < \kappa/\theta$), the large-firm job is strictly dominated and never chosen under uncertainty.

Transparency regime: salary revealed at both L and S. Mandating disclosure removes the Jensen discount, allowing workers to directly compare $w_L^\theta a_L^{\gamma_g}$ versus $w_S^\theta a_S^{\gamma_g}$. In Figure 5, this is shown by the coral line L_T (large firm under transparency), which lies above the gray control line L_C because uncertainty is eliminated. The new cut-off is therefore:

$$\gamma_g \;\le\; \gamma^T \;=\; \frac{\theta\,\Delta w}{\Delta a} \;=\; \gamma^C + \frac{\kappa}{\Delta a} \;>\; \gamma^C. \tag{2}$$

In the figure, this appears as the right dashed vertical line γ_T^*. Transparency restores the full value of the higher wage, shifting the threshold right by exactly $\kappa/\Delta a$. The resulting switcher band consists of workers who previously preferred the small firm due to uncertainty about wages, but who are now willing to trade off amenities for a higher observed salary. The new cut-off inherits the same comparative statics as before, except that it no longer falls with κ because $\kappa = 0$ under transparency.

Switching interval. Pay transparency therefore expands the set of workers who prefer the large firm from $[0, \gamma^C]$ to $[0, \gamma^T]$. All workers with $\gamma \in \left(\gamma^C, \gamma^T\right]$ – the *switcher band* – move from S to L once wages are posted. This switchers' interval is highlighted by the gray shading in Figure 5.

Gender-specific switching mass The lower panel of Figure 5 reveals that transparency has differential effects by gender. The blue curve shows the cumulative distribution function $P_m(\gamma)$ for men, while the red curve shows $P_f(\gamma)$ for women. Because women have stronger preferences for amenities on average (first-order stochastic dominance), the female CDF lies to the right of the male CDF. Let $\Delta_g = P_g(\gamma^T) - P_g(\gamma^C)$ denote the fraction of gender g that lies in the switcher band. The switching masses are illustrated by the vertical arrows along the y-axis: Δ_m (in blue) shows the fraction of men in the switcher band, while Δ_f (in red) shows the corresponding fraction for women.

Because the female density $p_f(\gamma)$ exceeds the male density $p_m(\gamma)$ for all γ in that interval[16], we obtain:

$$\Delta_f > \Delta_m \tag{3}$$

i.e., wage disclosure redirects a larger share of female than male applications toward high-pay, low-amenity jobs at large firms. Visually, this is evident from the fact that the red switching arrow Δ_f is longer than the blue switching arrow Δ_m in the figure.

5.4 Testable predictions

The framework above delivers two sharp, testable implications:

P1. *The transparency effect.* Revealing wages ($d_L = 1$) increases the application probability to large firms for both genders: $\Pr(L \mid \text{Treatment}) > \Pr(L \mid \text{Control})$.

P2. *Gender-heterogeneous response.* The increase is larger for women because their higher amenity preference places more of them on the marginal frontier where the revealed wage premium is just enough to offset the amenity loss: $\big[\Pr_f(L \mid T) - \Pr_f(L \mid C)\big] > \big[\Pr_m(L \mid T) - \Pr_m(L \mid C)\big]$.

5.5 Model takeaways

The framework nests two explanations for gender sorting. First, *preferences*: women value flexibility more, shifting their demand toward small firms ($a_S > a_L$). Second, *information frictions*: pay uncertainty discounts the wage premium at large firms. When wages are posted, women's higher amenity taste still pulls them toward S, but the removal of the information friction brings a disproportionate share of them to L. In this sense, flexibility is not preferred to pay in general; rather, it compensates for uncertainty about pay. In this sense, while the friction facing men and women is the same - uncertainty about w_L – its outcomes are gendered due to differential amenity preferences. Finally, though it takes both preferences and frictions to generate these sorting patterns, information alone can undo it. The next section tests these implications in a large-scale field experiment.

The model also informs the design needed to disentangle frictions from preferences. Observational data overstate the role of preferences by inferring them from choices made under incomplete information. Specifically, they do not account for the impact of perceived risk when workers sort on uncertain wage expectations rather than true wages. By contrast, lab experiments recover true preferences by removing frictions, but miss how frictions distort behavior at the margin. Existing lab

[16]The CDF gap $\Delta P(\gamma) \equiv P_m(\gamma) - P_f(\gamma)$ is weakly decreasing throughout the switcher band $(\gamma^C, \gamma^T]$ which is a mild strengthening of first order stochastic dominance that holds, for example, under single-crossing/location-shift families. Empirically, this corresponds to switcher bands that capture moderate amenity tastes: high enough that many women are near indifferent at baseline, but not so high that essentially no one is willing to move to L.

designs also draw wages and amenities independently, though they often co-vary outside the lab. A field experiment that randomizes pay transparency can remove wage frictions (as in the lab), while preserving realistic trade-offs (as in observational data) to capture marginal responses.

6 Field experiment: Design

The experiment is designed to test two key predictions of the model: (i) that pay transparency increases job applications to large firms, and (ii) that this effect is stronger for women. Both were pre-registered. To implement the experiment, I partner with Pakistan's largest online job search platform. Together, we randomly assigned 20,088 job ads posted by 8,906 firms to either mandatory or optional salary disclosure.

Job-level randomization. Treatment is assigned at the time of job posting. When a firm representative initiates a new job post, a random number generated in the background determines whether the position is allocated to the treatment or control group. Appendix Figure A.11 confirms that these draws follow a uniform distribution. In the control condition, firms retain the ability to conceal salary information from job-seekers, as illustrated in Appendix Figure 6, Panel (a). In the treatment condition, this option is removed. Instead, the interface displays the message: "To find you the best match, the salary for this job will be displayed in the ad, as part of a reform. Click here to learn more." This prompt is shown in Panel (b) of Figure 6. Additional information is provided on demand. Clicking the "Learn More" link opens a pop-up window (Figure A.9, Panel (a)) that explicitly informs the firm that the platform is conducting an experiment to assess whether salary transparency improves job search, and that job ads are being randomly selected to disclose salaries. This ensures that assignment to treatment is understood as the result of randomization, not of any particular firm behavior or job characteristic. In doing so, the platform centrally manages firm beliefs about the new information regime in a consistent and transparent manner.

Opt out procedure. All treated firms are initially defaulted into the experiment, but those that actively request to opt out are permitted to do so at a moderate administrative cost. As shown in Figure A.9, Panel (a), firms with concerns about the study are invited to contact the platform. These inquiries are typically handled by phone. During these calls, firm representatives are read a pre-scripted, detailed explanation outlining the study's objectives and the randomization procedure. Firms that still wish to opt out must explicitly confirm their preference to withdraw, despite the stated rationale, and await manual processing to receive an override enabling salary non-disclosure. Whereas standard job posts are published within one business day, opting out typically introduces a delay of two to three days. The experiment is explicitly designed to approve all such opt-out requests in order to prevent any attrition of firms or job ads from the platform. This is particularly useful in

this setting because, even when firms opt out of public disclosure, they are still required to privately report salary ranges to the platform. This feature enables analysis of selection into non-compliance. This is a key advantage of this study's environment, as existing studies of wage transparency laws frequently document substantial non-compliance but are unable to examine its correlates, given that undisclosed salaries remain inaccessible to both job-seekers and researchers. For example, despite legal mandates, 30% of job ads in Colorado omit salary information from job ads (Arnold et al., 2022).

Managing inattention. To account for potential inattention, the platform displays an additional pop-up in treatment immediately before the job post is finalized. This message reminds the treatment group that the salary information they have entered will be publicly visible. This confirmation screen is presented in Figure A.9, Panel (b).

Managing strategic manipulation. Because randomization is implemented at the job ad level, firms cannot circumvent treatment assignment by refreshing the webpage or reposting the same job under an identical title. The platform ensures consistency by linking treatment status to the initial attempt to create a given job post – even if the post is left incomplete and resumed at a later time. To further safeguard against evasion, the platform employs a dedicated sales team that reviews all ads before they go live. If a firm attempts to bypass treatment by making superficial modifications to the job title while posting a substantively identical position, the sales team contacts the firm, standardizes the title to match the earlier post, and reassigns the original treatment status to the updated ad. This enforcement mechanism ensures fidelity to the experimental design and limits scope for strategic non-compliance. To address any residual inconsistencies due to platform oversight, the analysis collapses all job titles posted by the same firm into a single observation. Each such observation is assigned the treatment status corresponding to the job's first appearance on the platform during the experiment.

Saturation design. The experiment assumes that a given job's treatment assignment does not affect outcomes for its competitors (the stable unit treatment value assumption). To assess the validity of this assumption and detect potential spillover effects across jobs competing for the same applicants, I embed a two-stage cluster saturation design into the randomization process. In the first stage, I construct clusters representing distinct labor market segments, defined by occupation–industry pairs. To ensure that these clusters are relatively self-contained, I use historical application data to identify overlap in applicant pools across cells and merge any pairs with more than 10% mutual overlap. This process reduces 1,118 occupation–industry combinations to 474 disjoint clusters, which I illustrate in Figure A.10. I then randomly assign each cluster to one of two treatment saturation conditions: in high-saturation clusters, 75% of job ads are assigned to treatment; in low-saturation

clusters, only 25% are treated. This variation in local treatment density generates exogenous differences in job-seekers' exposure to disclosed salaries within each market. In the second stage, treatment is randomized at the job level within each cluster. I stratify assignment of clusters by their historical number of job postings to maintain balance in the overall share of treated and control jobs across the experiment. This design supports identification not only of the direct effects of salary disclosure on treated jobs, but also of spillover effects by comparing outcomes across markets with differing intensities of competitor treatment – thus enabling empirical tests of interference.

7 Field experiment: Empirical Approach

I estimate intent-to-treat effects of salary disclosure on job outcomes using two complementary specifications. The primary outcomes of interest include posted salaries, application volumes by gender and measures of applicant quality.

To estimate treatment effects on log posted salaries and applicant quality, I estimate the following ordinary least squares (OLS) specification:

$$Y_j = \alpha + \beta T_j + \epsilon_j \tag{4}$$

where Y_j denotes the outcome for job post j, and T_j is a binary indicator equal to 1 if the job is assigned to the treatment group and 0 otherwise, irrespective of whether the firm complies with the assigned disclosure status. Since some treated posts may ultimately choose not to disclose salaries and some control posts may voluntarily disclose them, the estimated coefficient β captures a lower bound.

To examine impacts on application volume, I estimate the following Poisson model and report incidence rate ratios:

$$Y_j = \exp(\alpha + \beta T_j) \tag{5}$$

The Poisson specification is better-suited for count data and, importantly, allows treatment effects to be interpreted as percentage changes relative to the control group. This is especially useful given the baseline gender imbalance on the platform: men outnumber women more than three to one, leading to mechanically higher male application counts for the average job. Percent-based estimates offer a scale-adjusted comparison across groups and avoid misleading interpretations based on absolute levels. For context, I also report OLS estimates to assess whether these percentage changes translate into economically meaningful differences in application volumes.

8 Field experiment: Results

8.1 First stage and compliance

Table B.5 confirms high compliance with treatment. It reports first-stage results, where the dependent variable is an indicator for whether the job ad includes a visible salary. Column 1 shows that 56.2% of control group job posts disclose salary information. Assignment to treatment increases this likelihood by 40.1 pp. This implies a non-compliance rate of just 3.7% among treated jobs. Most of this non-compliance is driven by large firms. Column 2 disaggregates results by firm size. Salary disclosure is more common at baseline among small firms (64.0%) than large firms (44.8%). Among small firms, treatment increases disclosure by 35.8 percentage points, yielding near-perfect compliance (non-compliance rate: 0.3%). Among large firms, treatment increases disclosure by 47.0 pp – resulting in a somewhat higher, but still moderate, non-compliance rate of 8.3%.

Additionally, Figure A.13 shows that the number of job postings and active firms remained stable following the treatment, relative to baseline levels, suggesting that the mandate for salary disclosure did not lead to a reduction in platform participation.

Negligible non-compliance and firm exit rates reflect the platform's substantial market power. Most firms treat platform requirements such as salary disclosure as necessary conditions for accessing the country's largest pool of job-seekers. In the case of this intervention, they likely weigh the potential costs of disclosure against the opportunity cost of losing access to this market. Large firms have more leverage, as their exit could reduce the platform's appeal to job-seekers. Thus, they feel better positioned to request various exceptions to platform rules – even though in this case such exceptions are granted to all who ask. Since such firms are few, the platform is able to maintain high compliance regardless.

8.2 Advertised salaries

Salary disclosure requirements can raise or lower posted wages. Salaries may increase as firms compete for talent – especially since firms that previously withheld pay were offering higher wages and are now induced to reveal them. Conversely, posted salaries may fall if disclosure enables implicit or explicit collusion among firms, or if firms preempt demands from incumbent workers by offering lower starting salaries. The intervention may also affect the width of posted salary ranges: firms seeking to obscure exact pay may broaden the range to reduce precision, a potential form of strategic non-compliance with the reform's intent.

Table 2 considers treatment effects on the log minimum and maximum advertised salaries, as well as the size of the salary range, computed as the maximum minus minimum salary, divided by the midpoint of the range. Column 1 shows that the effect of the treatment on the log of the maximum advertised salary is null. Column 2 examines the treatment's impact on the log of the minimum

advertised salary. The results indicate that treatment leads to a 3% increase in the minimum salary (significant at the 5% level). The treatment has a negative and statistically significant effect of -0.03 (p<0.01), indicating that firms tend to post narrower salary ranges in response to the disclosure requirement. This is contrary to the expectation that firms would widen their ranges to avoid providing meaningful information to job-seekers. This finding aligns with evidence from the firm survey, in which firms report that applicants anchor their wage expectations to the top half of the salary range. Widening of salary ranges may therefore lead to tougher salary demands during negotiations. These results align closely with recent evaluations of salary transparency mandates. Arnold et al. (2022) find that a reform in Colorado increased both posted and contracted salaries by 3.6%, while Skoda (2022) reports a 3% increase in Slovakia following a similar policy. Although I lack data on contracted wages, the fact that these studies observe similar effects on posted and realized pay suggests that the impacts documented here likely extend to workers' final compensation as well.

8.3 Applications

Salary range disclosure in job ads should have no impact on workers' search behavior if there are no information frictions that salary transparency resolves, or if workers do not adjust their job search based on salary information. To test whether salary information matters, I consider treatment impacts on job views and applications. Views capture whether workers click on job summaries to access its details. These summaries, examples of which are shown in Appendix Figure A.17, include salary range information when available, and thus, an applicant choosing to view further details of the job may be responding to the salary information. Table 3 shows that views on job ads with salary range information increase by 45% (Panel A, Column 1) while job applications increase by 49% (Panel A, Column 4). The treatment effect on applications is 14 pp larger for women (Panel A, Column 6) than for men (Panel A, Column 6), with the difference significant at $p = 0.06$, as reported in Panel A, Column 6. This suggests that salary information plays an important role in determining where job-seekers direct search, particularly for women.

Large firms. If salary information matters for directing search, the largest effects should occur where salaries were previously harder to access, but are ex-post revealed to be high. As pre-registered, jobs at large firms match these characteristics exactly.[17] Table 4 Panel A, Columns 2-3 show that while views on treated small-firm jobs increase by 22%, they increase by 59% for jobs at large firms, with the difference significant at p-value=0.02. Similarly, Columns 5-6 show that applications to jobs at small firms increase by 23% but they increase by 66% for large firms (p-value=0.02).

[17]I classify firms with more than 50 employees as large. This cutoff approximately splits the sample in half. See Appendix Table B.1 for the distribution of the firm size.

Why do these effects concentrate among large firms? First, these firms are significantly more likely to withhold salary information. Table B.5 shows that in the control group, large firms are 19.2 pp (30%) more likely to hide pay ranges. Second, large firms offer higher salaries. Table 5 shows that in the control group, large firms' maximum posted salaries are 21% higher than those of small firms, and minimum salaries are 32% higher. Treatment does not affect salary levels at large firms. Although the average increase in posted salaries is driven primarily by small firms, their wages remain lower on average than those at large firms. Finally, salary ranges at large firms are narrower than at small firms in both treatment and control groups (Column 3), meaning disclosure provides a more precise wage signal for these jobs. Taken together, these results suggest that workers respond strongly when new information reveals higher pay.

Women. On the supply side, the model predicts that salary transparency affects both men and women, but the effect is larger for women. Figure 7 shows treatment effects by gender. In the full sample (left), treatment increases male applications by 46.1% from the control mean of 65.7 applications, and female applications by 60.2% (control mean = 15.9). This amounts to a gender gap of 14.1 pp (p-value<0.1). In the large-firm subsample (right), where treatment resolves wage uncertainty and reveals a wage premium, the effects and the resulting gender difference are even larger. Male applications increase by 58.8% (from a mean of 89.9), while female applications nearly doubled (95.2% from a control mean of 20.6), yielding a gender gap of 36.4 pp. As shown in Figure 8, these treatment effects not only close, but reverse the gender gap in directed search to large firms. The left panel shows the control condition, where large firms receive 82.0% more male applications and 61.5% more female applications than small firms – a gender gap of 20.4 pp in favor of men. The right panel shows the treatment condition, where salary information is disclosed. In this regime, large firms receive 130.9% more male applications and 173.5% more female applications relative to small firms, reversing the gender gap: the relative difference in women's applications to large versus small firms is 42.6 pp larger than men's. These findings are consistent with the model's prediction that women respond more strongly when high wages are made transparent at less flexible firms. A tabular version of results on both job ad views and applications is presented in Appendix Table B.6, while a corresponding graph overlaying OLS and Poisson treatment effects can be found in Appendix Figure A.14.

Randomization inference. Causal interpretation of the large increase in job applications rests on random assignment. To formally assess whether observed treatment effects could plausibly have arisen by chance, I conduct a randomization inference procedure with 1,000 placebo alternative assignments. The RI procedure proceeds in three steps. First, I compute the actual treatment effect for each outcome—defined as the coefficient on salary disclosure in a regression of application counts, estimated separately for male and female applicants and for large and small firms. Sec-

ond, I construct a distribution of placebo treatment effects by re-randomizing treatment assignment 1,000 times, ensuring that each placebo draw matches the actual share of treated jobs within the relevant firm-size subgroup. Third, I compute the randomization inference *p*-value as the proportion of placebo estimates with absolute values greater than or equal to the observed treatment effect.

Appendix Figure A.12 visualizes the resulting placebo distributions and overlays the actual treatment effect in red. Each panel corresponds to a separate subgroup (men/women $\times$ large/small firms). Vertical black dashed lines indicate the average placebo effect ($\bar{\beta}_{RI}$), while red dashed lines show the true treatment effect (β_{TE}). In all cases, the observed treatment effect lies far to the right of the placebo distribution. For example, among large firms, the observed increase in female applications ($\beta_{TE} = 19.63$) is extreme relative to the placebo distribution averaging 0.02, with a randomization inference p-value of 0.007. Male applications to large firms are also starkly different: the true treatment effect of 52.79 is well outside the simulated null, yielding an RI p-value of 0.001.

Among small firms as well, observed increases in applications remain statistically distinguishable from zero: for female applications, the randomization inference p-value is 0.042; for male applications, it is effectively zero. These results provide robust confirmation that treatment effects on applications are not artifacts of random chance, but instead reflect systematic shifts in job-seeker behavior in response to a specific randomization regime.

Spillovers. I find no evidence that the treatment effects on job applications are driven by spillovers across jobs or firms. First, control jobs receive similar volumes of male and female applications regardless of the treatment intensity of neighboring jobs. In Table B.10, Columns 1–4, the coefficients on *High saturation cluster* are small and statistically insignificant across all models. For example, the estimated effect on male applications is 4.06 (s.e. = 3.04) in Column 3, and –0.38 (s.e. = 2.90) in Column 4. This rules out meaningful crowding in or out of job-seekers from untreated jobs located near treated ones.

Second, control firms do not become more likely to disclose salaries when their competitors are more heavily treated. As shown in Table B.10, Column 5, the coefficient on "High saturation cluster" is precisely estimated at zero (–0.00, s.e. = 0.02), indicating no change in salary visibility for control jobs located in treated clusters.

Third, treatment effects do not appear to vary with the saturation of neighboring jobs. In Table B.11, I condition on whether the job is located in a high-saturation cluster. Treatment effects on applications remain large and statistically significant, particularly among large firms. For example, the interaction between "Treated Job" and "Large firm" for female applications in Column 1 is 17.67 ($p < 0.01$), and in Column 5 (male applications) is 40.96 ($p < 0.01$), showing that treatment effects persist even when adjusting for cluster-wide exposure.

8.4 Family-friendly amenities

The average treatment effects establish that disclosure draws more women to large firms. What these averages leave unanswered is where in the amenity distribution those additional applications come from. Do women merely stack large-firm jobs to their pile of flexible job applications or do they re-optimize? To unpack this, I look inside each experimental arm and track how the relative appeal of key non-wage amenities shifts once wages are made visible. Within treatment and control groups, I then relate application counts to these amenities and firm size in a poisson specification, conditioning on the full set of other amenities and job characteristics. The resulting coefficients tell us how attractive a given attribute is after netting out other attributes, as well as the overall surge in applications generated by the intervention.

Classifying amenities. I begin by classifying job amenities by developing a context-specific natural language processing algorithm that labels every job posting according to whether it advertises remote work, flexible hours, transport support, a safe-work environment, or an equal-opportunity statement. These amenities are chosen based on extensive prior work suggesting women may value them differentially.

Supply of amenities. Table 6 examines how the prevalence of these job amenities varies by firm size, separately for control and treatment jobs. Consistent with the baseline facts reported in Section 4, across both groups, large firms are significantly less likely to offer remote work and flexible hours. In the control group (Panel A), large firms are 5 pp (50%) less likely to offer remote work and 2 pp (40%) less likely to offer flexible arrangements, relative to small firms (Columns 1 and 3). These gaps persist under treatment (Panel B), where large firms are 6 pp (60%) less likely to offer remote work and 3 pp (60%) less likely to offer flexible hours – both significant at the 1% level. In contrast, there are no systematic differences in the likelihood of offering transport assistance (Column 2). Meanwhile, large firms are more likely to mention a safe environment for women (1 pp) and to present themselves as equal opportunity employers (1–2 pp; Columns 4–5)

Demand for amenities: Measurement To see how these amenities shape job search under different information environments, I estimate Poisson regressions of application counts by gender on firm size and amenities, controlling for job observables. This is represented by the equation below:

$$Y_j = \exp\left(\alpha + \gamma X_j + \Lambda(T_j \cdot X_j)\right)$$

where Y_j are gender-specific application counts directed to job j, T_j is a treatment indicator for job j, and X_j includes a vector of job characteristics such as education and experience requirements, career level, and industry and occupation fixed effects. Table 7 reports the results from this exercise,

estimating how specific job-level amenities predict the number of applications submitted by women and men, separately for control jobs with hidden salaries (Columns 1–2) and treatment jobs with disclosed salaries (Columns 3–4). Each coefficient reflects the relative change in application volume associated with a particular amenity, controlling for other job characteristics.

Demand for amenities: Results. In the control group, remote work is the strongest predictor of applications for both genders: jobs offering remote work receive 134% more applications from women and 59% more from men (both $p < 0.01$). Safe environment and equal opportunity language also significantly increase applications, especially for women: the presence of either amenity raises female applications by 42–70% and male applications by 40%. Flexible work hours and transport support, however, do not significantly influence application volumes in the absence of salary disclosure. Figure A.15 formally tests for gender differences in responsiveness to each amenity by comparing coefficients across Columns 1 and 2. Remote work emerges as the most salient gender gap, increasing female applications 75.3 pp more than male applications ($p < 0.01$). The next largest gap is for equal opportunity language (30.9 pp, $p < 0.10$). However, since large firms are more likely to include equal opportunity language in their job ads, this amenity is unlikely to explain the observed gender gap in applications to small versus large firms.

Table 7 Column 3 and 4 show that under treatment, the predictive power of amenities for women's applications fades considerably. Remote work remains significant but its effect declines to 74% for women though it stays relatively stable at 51% for men (both $p < 0.01$). In contrast, flexible work arrangements become negatively associated with application volume: women submit 35% fewer applications and men 27% fewer to jobs offering flexibility (both statistically significant). Transport support also becomes negatively associated with applications, reducing female applications by 44% and male applications by 26%. Meanwhile, the effect of large firm status becomes significantly stronger once salaries are disclosed. In the control group, large firms receive 13% more female applications (not statistically significant) and 28% more male applications ($p < 0.01$). Under treatment, however, large firms receive 61% more female applications and 42% more male applications (both $p < 0.01$).

Taken together, these patterns suggest that in the absence of pay transparency, remote work plays a large role in attracting applicants, particularly women. Salary disclosure reduces the appeal of remote work and other family-friendly amenities. Instead, it redirects search toward large firms, even when those firms are less likely to offer the amenities women value. This compositional evidence deepens the headline results in two ways. First, it decomposes the average treatment effect, showing that the surge at large firms is fueled by a decline in the relative appeal of flexibility rather than by indiscriminate stacking of new applications on an unchanged portfolio. Second, it suggests that what looked like a strong, fixed taste for flexibility under hidden wages largely dissipates once the wage component of the bundle becomes visible.

9 Lab experiment: Cross-randomizing amenities and salaries

A potential concern with this analysis is that amenities are not randomly assigned. While I control for industry, occupation, education, experience, and career level, unobserved differences may still influence both the supply of an amenity and its appeal to job-seekers. This limits the causal interpretation of within-arm estimates. Across arms, salary disclosure is randomly assigned and does not alter the distribution of advertised amenities. Comparisons across arms therefore reflect how job-seekers respond to a stable set of amenities under different information regimes. Still, without random assignment of amenities, average effects may mask unobserved heterogeneity in how the kinds of jobs offering a given amenity interact with salary disclosure to shape applications.

Design. To address these concerns, I implement a discrete choice experiment with 438 job-seekers, in which salary visibility is cross-randomized with family-friendly amenities in a controlled setting. Each respondent evaluates four pairs of stylized job postings (eight job profiles in total), where the key amenities – salary visibility, a large firm label, remote work, transport support, and equal opportunity language – are independently randomized. Salaries, when visible are randomly drawn between 110-120% of workers' expected wage. To anchor beliefs about the job beyond these traits, each vignette also displays a fixed job title, location, education requirement, and schedule. The first three attributes are personalized using each respondent's current or preferred job title, location, and education level. Additionally, each job is attributed to a fictitious firm name designed to be neutral and uninformative about the firm's industry or other characteristics. Figure A.16 presents three illustrative examples of vignettes generated through this experimental design.

Estimation. I estimate whether women pivot from flexibility to higher pay when large firms reveal salaries – and when the revealed salaries rise. To do so, I estimate two models. The first is a linear probability model (LPM), which leverages both within- and across-choice-set variation:

$$\Pr(Y_{ij} = 1) = \alpha + \gamma \cdot Remote_j + \sum_s \delta_s \cdot Salary_{sj} + \sum_s \lambda_s \cdot (Remote_j \times Salary_{sj}) + \mathbf{X}_j'\beta + \varepsilon_{ij}$$

where $Y_{ij} = 1$ if individual i selects job j; $Remote_j$ indicates whether the job offers remote work; $Salary_{sj}$ are indicators for salary levels $s \in \{low, medium, high\}$; and $\mathbf{X}_j$ includes other randomized job-level attributes: diversity language and transport support. The interaction terms λ_s capture how the appeal of remote work varies across salary levels.

The second specification is a conditional logit model, which includes fixed effects for each choice pair and relies solely on within-set variation:

$$\Pr(Y_{ij} = 1 \mid \text{ChoiceSet}_i) = \frac{\exp\left(\gamma \cdot \text{Remote}_j + \sum_s \lambda_s \cdot (\text{Remote}_j \times \text{Salary}_{sj}) + \mathbf{X}_j'\beta\right)}{\sum_{k \in \text{ChoiceSet}_i} \exp\left(\gamma \cdot \text{Remote}_k + \sum_s \lambda_s \cdot (\text{Remote}_k \times \text{Salary}_{sk}) + \mathbf{X}_k'\beta\right)}$$

Results. Table 8 reports estimates from the linear probability model, while Table B.7 presents the conditional logit results. Both tables disaggregate results by gender and salary level. Panel A of each table shows effects for women; Panel B for men. Despite differences in estimation strategy and identifying variation, the results are highly consistent across the two approaches. Table 8 shows that when salaries are hidden, women are 13.2 pp more likely to choose remote over on-site jobs. However, this preference declines once salaries are disclosed and no longer significantly influences job choice at any wage level. At the highest salary tier, the remote premium falls to 6.9 pp and is statistically insignificant. Conditional logit estimates in Table B.7 confirm this pattern: under pay opacity, remote jobs are 14.9 pp more likely to be chosen, but this effect disappears entirely once salaries are made visible, falling to 3.6 pp (statistically insignificant) at the highest salary level. Men's preference for remote work does not vary with salary disclosure. At no wage level – and in neither specification – is the likelihood of choosing a remote job over an on-site job significantly different from zero. If anything, at the highest salary tier, men show a slight preference for on-site jobs over remote ones.

Takeaways. These results suggest that although women value flexibility more than men, flexibility does not drive their job search behavior when salary information is available. Nor do women appear to trade off wages for flexibility – there is no evidence that they accept lower pay in exchange for remote work when pay information is available. Instead, flexibility becomes relevant only in the absence of salary information. When large firms disclose pay, women reallocate their search toward higher-paying jobs, with application likelihood rising monotonically with salary levels. In fact, Table 1 shows that salary disclosure is the single most influential job attribute in women's application decisions. Among all randomized features in the discrete choice experiment, it has the strongest effect on the likelihood of job selection – raising the odds by a factor of 2.65. This effect is substantially larger than that of remote work (OR = 1.38), transport subsidies (OR = 1.69), or equal opportunity language (OR = 1.53). Taken together, these results underscore that wage transparency – not flexibility – is the dominant driver of women's job search.

10 Alternative mechanisms

A large literature has documented a range of gender differences that could plausibly affect sorting and labor market outcomes. In this section, I use data from the job-seeker surveys (described in Section 3) to test each of these alternative explanations.

10.1 Gender gaps in information or inference

Gender differences in labor market experience or attachment may create disparities not only in the information men and women hold about jobs, but also in how they interpret non-wage signals in job ads and form beliefs about wages, competition, and firm characteristics, especially when firms withhold such information from public view. To measure these gaps directly, I present job-seekers with two job ads each from the platform, and ask them to guess characteristics such as wages, competition, and firm size. Their responses are incentivized with financial rewards for accuracy, and benchmarked against firm-reported truths or administrative data. Importantly, the ads are drawn from each respondent's own occupation, ensuring that any gender differences in accuracy reflect information relevant to their own job search rather than unequal knowledge of the broader labor market. I find that such gender gaps in information or inference cannot account for the patterns of sorting I document.

Hidden salaries. To test whether men and women differ in their ability to infer pay, I present job seekers with two ads from the same platform that do not disclose salaries, one from a small firm and one from a large firm. Respondents are asked to guess the salary range and are offered a financial reward if their guess falls within 10% of the true salary reported by the firm to the platform. Table B.12 shows that while both men and women have little information about hidden wages, the gender differences in guessing accuracy are not statistically significant. Columns (1)–(2) report the probability respondents guess the minimum of the hidden salary range within 10% of the true value. Female coefficients are small and insignificant: 0.018 for small firms and 0.03 for large firms, compared to male means of 0.069 and 0.09, respectively. Columns (3)–(4) repeat the exercise for the maximum of the salary range. Again, there is no evidence of a female disadvantage: women's relative accuracy rates (0.04 at small firms and –0.04 at large firms) are statistically indistinguishable from men's (men's means: 0.07 and 0.14, respectively). Importantly, for both men and women, accuracy does not differ meaningfully by firm size. Overall, these results indicate that women are not systematically less informed about hidden salaries. Gender differences in sorting are therefore unlikely to be driven by informational gaps about undisclosed wages.

Competition. A second possibility is that men and women differ in their knowledge of competition for jobs. If women expect greater competition for positions with hidden salaries, particularly at

large firms, they may be less inclined to apply to such jobs. To test this, I asked job-seekers to guess the number of male and female applications received for the same two ads used in the salary exercise, one from a small firm and one from a large firm. Respondents received a financial reward if their guess fell within 10% of the true number of applications observed on the platform. Table B.13 reports the results. Columns (1)–(2) show the probability of guessing the number of female applications within 10% of the true value. The female coefficients are –0.01 for small firms and –0.01 for large firms, compared to male means of 0.05 and 0.04, respectively. Neither are statistically significant. Columns (3)–(4) repeat the exercise for male applications. Again, there is no evidence of differences in beliefs about competition by gender: women's accuracy rates (0.003 for small firms and –0.01 for large firms) are statistically similar to men's means of 0.04. As with salaries, both men and women appear to have little information about the level of competition for jobs, and any differences across gender are small and statistically insignificant. This suggests that gender gaps in sorting are not driven by asymmetric beliefs about the number of applicants competing for jobs with hidden salaries at large firms.

Firm size. A third possibility is that men and women differ in their knowledge of firm size. Women may value the opportunity to work in large firms, but if they are less able to infer firm size from job ads, they may systematically under-apply to such jobs. To test this, I asked job-seekers to guess the number of workers employed at the firms in the same two job ads. I then classified responses as correct if they placed the firm on the right side of the large-firm threshold (>50 workers). Table B.14 reports the results. Columns (1)–(2) show the probability of correctly classifying a firm as small or large. For small firms (Column 1), the female coefficient of 0.042 is statistically indistinguishable from a male mean of 0.36. For large firms (Column 2), the female coefficient is –0.02, compared to a male mean of 0.728. In both cases, the female coefficients are small and statistically insignificant. While there is little evidence of a gender gap in beliefs about firm size, both men and women are far more accurate when the ad comes from a large firm. This suggests that women knowingly under-apply to large-firm jobs.

10.2 Gender gaps in network size and quality

Another potential explanation is that men and women differ in the size or quality of their social networks, which could affect the flow of information they receive about jobs prior to application. If men are embedded in networks that are larger or higher quality, they may have better access to information about wages, competition, or firm characteristics, and thus be better able to target applications. Table B.15 reports gender differences in self-reported network characteristics. Column (1) measures network size as the number of close friends or relatives with whom the respondent interacts at least once every two weeks. Column (2) reports the share of these contacts that provide job search advice. Column (3) reports the share of the network that is employed, and Column

(4) the share of employed contacts that earn more than the respondent. The results show broadly similar networks across genders, with one exception: women report a smaller share of employed contacts (–0.09 pp), relative to a male mean of 0.79. However, Column (5) shows that the share of a respondent's network that is employed does not predict a higher share of applications to treated jobs for either men or women. Thus, although men report 11.2% more employed contacts, this difference does not translate into more targeted applications. Consistent with this, across all incentivized belief-elicitation tasks described above (guessing salaries, the number of applications, or firm size) there are no systematic gender differences in accuracy. Figure 4 further reinforces this point: as discussed previously, when salaries are hidden, men's application behavior is not responsive to actual wages. If men's networks conveyed more accurate information about hidden salaries, we would expect their application slope to be upward-sloping with respect to pay. Instead, the flat slope suggests that men's different network composition does not result in more accurate beliefs or in different search strategies.

10.3 Perceptions of discrimination at large firms

Another possibility is that women avoid large firms because they expect to be paid less than equally qualified men. Once salaries are disclosed, they may update these beliefs or expect that discrimination becomes harder to sustain given a public commitment to wages. Table B.16 provides evidence against the prevalence of such perceptions. Columns (1)–(2) examine respondents' expectations about their own salary offers. After guessing the minimum and maximum salary for a large-firm job, respondents reported what monthly salary they thought they themselves would be offered. The outcomes capture whether this expected offer fell within 10% or 20% of their previously stated maximum. Female coefficients are small and statistically insignificant (–0.038 in Column 1 and –0.054 in Column 2), relative to male means of 0.318 and 0.576. This indicates that women do not perceive themselves as substantially disadvantaged in their own offers relative to men. However, measures about own salary offers conflate beliefs about personal fit or ability with general expectations of gender discrimination. Column (3) addresses this limitation by eliciting second-order beliefs. Respondents were shown two otherwise identical candidate profiles, one male and one female, and asked how much the firm would pay to each of them. Column (3) reports whether respondents' answers implied a positive pay gap between the male and female candidate despite otherwise similar résumé characteristics. Women are slightly more likely than men to report that the male candidate would receive a higher offer (0.05 pp), but the difference is statistically insignificant. Overall, only a minority of both male and female respondents perceive discrimination of this kind (22% and 27% respectively). Together, these results suggest that women neither expect systematically worse own offers at large firms nor perceive discrimination against female candidates as widespread. Perceptions of discrimination therefore appear unlikely to explain women's under-representation in large-firm jobs.

10.4 Intra-household bargaining

Women may face household constraints when deciding whether to apply for jobs. In particular, they may need to consult or seek permission from household members before applying. When salaries are undisclosed, especially in jobs that also lack flexibility, returns are more ambiguous, making it harder to justify such applications alongside household responsibilities. Table B.17 documents that women are 19.3 pp more likely than men to report seeking someone's permission before applying. For men, permission is most often sought from fathers, and the likelihood declines with age. For women, younger respondents are also most likely to seek permission from fathers, but as they get older they also report seeking permission from spouses. As a result, women remain just as likely to require permission at older ages, whereas men become less likely to do so. However, while women are more likely to seek permission from household members, this does not appear to drive their responsiveness to wage disclosure. Table B.18 examines whether this household constraint translates into different responses to treatment. Column (1) replicates the gender gap in permission-seeking. Column (2), however, shows that permission-seeking status does not predict a higher share of applications to treated jobs for either genders. This null result should not be taken to imply that domestic constraints are unimportant. On the contrary, it is highly plausible that the gendered division of labor within households is the key source of women's stronger preference for flexibility. What my data show is simply that this channel does not manifest in the form of reported permission-seeking or household-bargaining. Rather, household constraints are likely embedded more deeply in the trade-offs women face between job characteristics and domestic responsibilities.

10.5 Gender differences in risk, ambiguity and bargaining aversion

I next examine whether gender gaps in preferences over risk, ambiguity, or bargaining could account for differences in application behavior under different information regimes.

Measuring risk aversion. I measure ambiguity aversion using a standard incentivized game following Ashraf et al. (2009). Respondents choose between two urns for a cash prize. In the "known risk" urn, the composition of red and yellow balls is fully disclosed, so the probability of winning is transparent. In the "ambiguous" urn, the composition is unknown, but respondents can choose which color counts as a win. Choosing the risky urn over the ambiguous urn indicates higher ambiguity aversion.

Measuring ambiguity aversion. To measure risk attitudes, I rely on a subjective self-assessment question: "On a scale of 1 to 10, with 1 being not willing at all and 10 being very willing, how willing are you to take risks, in general?" This contrasts with the more traditional lottery-based

elicitations of risk preferences. As reviewed in Bertrand (2011), such subjective measures have been shown to be as reliable as lottery-based methods.

Measuring bargaining aversion. Finally, I elicit bargaining preferences using a hypothetical job-offer scenario. Respondents are asked to imagine receiving two otherwise identical job offers: one from a firm that sends a final offer letter with a fixed salary, and another from a firm that invites them to negotiate the salary before making a final offer. Respondents can only respond to one firm. Choosing the non-negotiable offer indicates a preference to avoid bargaining, while choosing the negotiable offer indicates willingness to engage in bargaining.

Results. Table B.19 reports the results. Columns (1)–(3) show gender gaps in preferences. Women are no more likely than men to be ambiguity averse (Column 1). They report somewhat lower willingness to take risks (–0.63 on the 1–10 scale, Column 2), though the effect is modest relative to the male mean of 5.94. By contrast, women are substantially more likely to be bargaining averse (Column 3): the female coefficient of 0.18 pp (30%) is large and significant. Columns (4)–(5) link these preferences to job application behavior. Risk preferences are unrelated to the share of applications to treated jobs. In contrast, bargaining aversion predicts a higher share of applications to treated jobs, and the effect is concentrated among women: the female $\times$ bargaining-averse interaction is positive and statistically significant (0.12, $p < 0.1$). Together, the results point to modest gender differences in risk, no gender differences in ambiguity, but sizeable gender gaps in bargaining preferences. Moreover, bargaining aversion, but not risk or ambiguity aversion, help explain why women respond more strongly to wage disclosure in their application choices.

Implications. The evidence shows that women are substantially more bargaining averse than men, and that bargaining aversion predicts stronger responsiveness to salary disclosure. However, bargaining aversion on its own cannot explain the firm-size heterogeneity in women's behavior: when pay is hidden, women randomize across small firms in the same way as men, but avoid high-wage jobs in large firms, as shown in Figure 4. Thus, I view bargaining aversion as an additional force that amplifies women's sensitivity to wage opacity, rather than exclusively drive sorting differences. In the baseline model, the utility from a job without wage disclosure is

$$U = \mathbb{E}[w^{\theta}]\, a^{\gamma_g},$$

where θ governs gender neutral preferences over wages and γ_g captures the valuation of non-wage amenities a. Bargaining aversion can be incorporated by introducing a parameter $\beta_f < 1$ that effectively increases concavity in the wage dimension under non-disclosure for women alone:

$$U = \mathbb{E}[w^{\beta_f \theta}]\, a^{\gamma_g}.$$

Lower values of β_f magnify the uncertainty penalty attached to hidden wages, making undisclosed pay less attractive relative to women. This mechanism raises the gains from transparency for women. In short, bargaining aversion strengthens the effect of wage disclosure on women's application behavior, but it cannot by itself account for why wage opacity deters women specifically from high paying jobs at large firms. That pattern arises from the interaction of opacity with limited flexibility in large firms, while bargaining aversion acts as an amplifier.

11 The demand side

The experiment shows that salary transparency meaningfully reshapes women's search, directing them toward large firms with higher wages but lower flexibility. Given that women globally sort into lower-paying firms, this shift alone makes a strong case for policy action. But for policy to be effective, we must understand its broader equilibrium effects. That many firms currently hide salaries suggests that disclosure imposes costs on employers even if it delivers benefits to job-seekers. A full evaluation of transparency must therefore account for both sides of the market: how workers respond to new information, and how firms adapt when disclosure is no longer optional. This section examines firms' incentives to conceal salaries, the impact of mandated disclosure on the composition of their applicant pools, and how firms adjust their behavior in the aftermath of the intervention.

Why firms hide attractive salaries. Advertising high salaries plausibly helps attract top talent. Yet, large firms frequently withhold this information, suggesting that disclosure carries potential costs. To examine these motivations, I conducted a baseline survey with approximately 300 employers. Among firms that do not disclose salaries, 63% expect transparency to increase the number of applicants, while 72% anticipate that it will weakly reduce average applicant quality (40% expect a strict decline). A simple model presented in Appendix Section B rationalizes these beliefs – showing that advertising high salaries can lead to a decline in average quality. Separately, 71% of firms also report that salary disclosure may constrain wage-setting, as posted salaries anchor candidates' wage expectations regardless of their individual qualifications. These responses highlight three underlying concerns: first, that larger applicant pools raise the cost of screening; second, that higher posted salaries may invite applications from less-suitable candidates; and third, that disclosure reduces firms' flexibility in tailoring offers.

These concerns are not unfounded. In my setting, the average job receives 82 applications, and 70–80% of candidates meet basic education and experience requirements on paper. These numbers reflect a broader challenge: as educational attainment rises globally, a growing share of candidates meet minimum qualifications, but the quality of their education remains highly uneven, particularly in low-income countries such as Pakistan (Muralidharan and Singh, 2023). As a result, standard résumé signals are weak predictors of ability (Bertrand et al., 2024; Bandiera et al., 2024), often

requiring firms to rely on costlier processes, such as interviews, to identify suitable candidates. These facts lend empirical support to the theoretical prediction by Michelacci and Suarez (2006) that salary non-disclosure is more common when applicant ability is highly dispersed and the risk of adverse selection is high. It also complements prior evidence that firms often recruit through familiar networks to reduce uncertainty about match quality (Jayachandran, 2021; Chandrasekhar et al., 2020).

Impacts on average applicant quality. While firms recognize that salary disclosure may increase the number of applicants, they also express concern that average applicant quality could decline to the point that the screening burden outweighs the benefits of a larger pool. Table 9 examines whether these concerns are substantiated in observable résumé characteristics. Panel A shows effects on the share of applicants who meet the experience requirements. The results are null for large firms, but for small firms, the share declines by 2 pp among female applicants and 1 pp among male applicants. Given control group means of 64–74%, this corresponds to a relative decline of 1.4–3.1% in the experience match of applicants. Panel B reports effects on education match. Treatment leads to a 1 pp decline for both male and female applicants at small and large firms. With control means between 88–92%, this represents a relative decline of roughly 1.1% in education match quality. Panel C shows the most variation in skills match. Among applicants to large firms, treatment reduces the share meeting all skill requirements by 3 pp for men and 2 pp for women. Given control means of 0.54 for men and 0.49 for women, this implies a 5.6% decline for men and a 4.1% decline for women. Taken together, these effects suggest modest declines in observed applicant quality, concentrated in skill-based matches at large firms. However, the consistently high control means for education and experience reinforce the descriptive finding that most applicants appear formally qualified, and that sorting on deeper, unobservable traits may be more difficult. Thus, while observable declines are small, they may signal broader reductions in unobserved applicant quality – consistent with firm concerns, and prior literature (Skoda, 2022).

Selection into the applicant pool. While average applicant quality matters for understanding screening costs, it is less informative about match quality: firms ultimately hire the best – not the average – candidate from their applicant pool. To assess how transparency affects selection into the applicant pool, Table 10 examines whether salary range disclosure changes the likelihood that high- or low-ability job-seekers apply to jobs at large firms. Appendix Table B.8 reports corresponding results for small firms. This analysis classifies workers based on their position in the market-wide distribution of résumé characteristics. Panels A and B report results for high-ability applicants, defined as those in the top decile (or top quartile for GPA) of each metric. Panel A shows effects for men; Panel B for women. Column 1 summarizes four résumé traits – education level, GPA, years of experience, and management experience – using a composite index. Columns 2–5 report treatment

effects separately for each component. Panels C and D repeat the analysis for low-ability applicants, defined as those in the bottom decile (or bottom quartile for GPA) of each distribution.

A key takeaway from Panels A and B is that high-ability men are already highly represented in applications to large firms even without salary disclosure. In the control group, 85–93% of applicants meet top-end thresholds across the four traits, and 92% fall in the top decile of the composite index. As such, the effects of transparency are modest: Panel A shows a 2.2 pp increase in the share of male applicants in the top decile of the index. Treatment effects on individual résumé traits range from 1.0 to 2.2 pp.

High-ability women, by contrast, are far less likely to apply under salary non-disclosure. In the control group, only 54–71% of applicants fall into the high-ability category for each trait, and just 63% are in the top decile of the composite index. Salary transparency significantly improves this: Panel B shows 7–8 pp increases across all measures, including a 7 pp increase in the share of top-decile women by the composite index, and 8.0–8.3 pp gains for GPA and experience. Gender differences in responsiveness are statistically significant in every column. These results indicate that transparency disproportionately draws in high-ability women.

Panels C and D show that transparency also increases the share of low-ability applicants, but to a lesser extent. Most large-firm jobs already receive applications from low-ability men, with control group shares between 92% and 98%. Treatment increases these shares by just 1.0–2.2 pp. For low-ability women, control group shares are also higher than for high-ability women: ranging from 75% to 85% for education, years of experience and management experience. This suggests that low ability women also apply broadly, like men. Treatment effects are also smaller than for high ability women, ranging from 2.0 to 5.0 pp. The composite index, which aggregates these traits, rises by 3.0 pp (Column 1).

While transparency does draw more low-ability women into the applicant pool, these increases are 3.2 pp smaller on the composite index than the corresponding gains among high-ability women. Thus, the evidence points to a favorable compositional shift: salary disclosure disproportionately attracts high-ability women, while also generating modest increases in lower-ability applications.

Quality of potential hires. If enough higher-quality applicants from the labor market are drawn to treated jobs, the quality of top candidates in treated applicant pools should improve. To test for this, I now consider whether the top candidates applying to treated jobs are higher quality than the top candidates applying to control jobs. To construct this measure, I identify the top three values in each applicant pool for five résumé traits: expected salary, education level, years of experience, GPA, and team size managed. An applicant is flagged if they appear in the top three on any of these traits, and a composite score is calculated based on how many traits they lead in. The three applicants with the highest scores are designated as the top candidates (ties are allowed). Table 11 presents the results of this analysis, with expected salary, education level, years of experience, GPA, and the team size

managed normalized using the mean and standard deviation of the control group to make effect sizes comparable across units. Column 1 shows that the average expected salary of these top candidates is no different in the treatment group compared to the control group. However, across Columns 2-5, we see that the education level, years of experience, GPA, and team size managed by these top candidates are 0.07-.0.14 SDs higher in treatment than control. Moreover, Column 6 shows that the likelihood that the top candidates include at least one woman is also 3 pp (5.2%) higher in treatment than in control. These effects are stable across small and large firms (Panels B and C respectively), but especially pronounced for large firms, where the likelihood a woman is among the top candidates is 7.1% higher in treatment than control.

Weighing the trade-offs: Firms' revealed preference. How do firms weigh the potential trade-off between higher screening costs and improved applicant quality under salary transparency? A natural test lies in their behavior after the experiment ends. Table 12 presents results from a firm-level analysis of post-experiment disclosure behavior in the six weeks following the experiment, as a function of the firm's degree of treatment exposure during the experiment.

Column 1 shows that, across all firms, a 10 pp increase in the share of the firm's job postings that were randomly assigned to salary disclosure leads to an 1.11 pp reduction in the probability that the firm hides salaries in subsequent postings. Column 2 shows no significant effect for small firms. But Column 3 reveals a strong and significant response among large firms: a 10 pp increase in treatment exposure reduces the likelihood of hiding salaries by 2.5 pp. Thus, firms who had all their job posts treated at 25 pp less likely to hide salaries compared to never treated firms (a 31% reduction).

This pattern of post-experiment behavior offers a revealed preference measure of how firms evaluate the net consequences of disclosure. The fact that firms with greater exposure to transparency are less likely to return to hiding salaries suggests that many initially overestimated the costs of attracting a broader applicant pool or underestimated the benefits of attracting high-ability candidates. In particular, the sharp response among large firms – those most likely to conceal pay in the first place – indicates that transparency shifted firms' beliefs about its value.

12 Women's hiring outcomes: Back-of-the-envelope calculation

While this paper primarily examines how pay transparency reshapes job search and sorting behavior, a natural downstream question is whether these changes also affect women's likelihood of being hired into better-paid jobs. To explore this, I collect job-level hiring data from firms, reporting the number of male and female applicants hired for each position. These data when combined with observed application behavior, allow me to estimate implied changes in a representative woman's likelihood of being hired.

This is a back-of-the-envelope extrapolation for two main reasons. First, the intervention is randomized at the job level, whereas the outcome of interest in this section is at the worker level: whether the average woman is more likely to be hired under salary disclosure. The calculation below therefore maps job-level quantities into worker-level hiring probabilities. Second, I field surveys at the firm level, not the worker level, which limits the granularity of the data. Surveying workers directly would have required interviewing most, if not all, of the 310,000 applicants to recover job-level hiring outcomes, an approach that was financially infeasible. From each firm, I instead collect information on up to five randomly selected jobs they posted during the experiment. Because firm representatives are typically busy and often unwilling to disclose personally identifiable information about hires to researchers or the platform, the survey asks only for aggregate counts: the total number of hires and the number of women hired per job. As a result, I cannot link individual outcomes to treatment status, and instead proceed under a random-matching assumption within applicant pools.

More broadly, estimating worker-level treatment effects is not the purpose for which the study was designed. The primary goal of the experiment is to evaluate whether transparency helps correct a misallocation of women's search. Hiring decisions, by contrast, reflect a wider set of factors beyond application behavior (e.g., discrimination, bargaining flexibility) that fall outside the scope of the research question in this paper. The extrapolation that follows is therefore presented as a suggestive exercise, rather than a definitive estimate.

Calculations. For a representative worker of gender $g \in \{f, m\}$, the probability of being hired is determined by two components: the probability of applying to a given job, and the probability of being hired conditional on applying:

$$\Pr(\text{hired}_g) = \Pr\big(\text{applies}_g\big) \times \Pr(\text{hired} \mid \text{applied}).$$

Let N_g denote the number of job-seekers of gender g in the labor market, and A_g the average number of applicants of gender g per job. Then:

$$\Pr\big(\text{applies}_g\big) = \frac{A_g}{N_g}.$$

Assuming firms hire randomly from the applicant pool, let H be the average number of hires per job and T the total number of applicants. Then:

$$\Pr(\text{hired} \mid \text{applied}) = \frac{H}{T}.$$

Combining both terms, the unconditional probability that a representative man or woman is hired becomes:

$$\Pr(\text{hired}_g) \approx \frac{A_g}{N_g} \cdot \frac{H}{T}.$$

To estimate the treatment effect of salary disclosure on this probability, I compare treatment and control groups:

$$R_g = \frac{\Pr(\text{hired}_g \mid \text{treatment})}{\Pr(\text{hired}_g \mid \text{control})} = \frac{(A'_g/T') \cdot H'}{(A_g/T) \cdot H}.$$

Appendix Table B.20 shows that jobs hire a similar number of male and female candidates across treatment and control. Thus, ($H' \approx H$), which simplifies the expression above to:

$$R_g = \frac{A'_g/T'}{A_g/T},$$

which captures the proportional change in unconditional hiring probability for gender g due to salary disclosure.This extrapolation is feasible because it relies only on observable quantities: the average number of applicants (A_g, A'_g) and hires per job (H, H'), while assuming that the size of the labor pool (N_g) remains constant across treatment arms. Using job-level data, I compute this ratio for both women and men. For women, the average number of female applicants per job is $A_f = 20.61$ in the control group and $A'_f = 40.24$ in the treatment group. The corresponding total number of applicants is $T = 110.45$ and $T' = 182.91$, respectively. This yields:

$$\frac{A_f}{T} = 0.1867, \quad \frac{A'_f}{T'} = 0.2200, \quad \Rightarrow \quad R_f = \frac{0.2200}{0.1867} = 1.178.$$

That is, a representative woman is 17.8% more likely to be hired by a large firm under salary disclosure than under concealment. This increase is not driven by changes in firm-side hiring behavior, but rather by the rise in female application rates relative to overall competition.

For men, the average number of male applicants per job is $A_m = 89.79$ in the control group and $A'_m = 142.58$ in the treatment group, yielding:

$$\frac{A_m}{T} = 0.8131, \quad \frac{A'_m}{T'} = 0.7795, \quad \Rightarrow \quad R_m = \frac{0.7795}{0.8131} = 0.959.$$

This implies that a representative man is 4.1% less likely to be hired under salary disclosure. These results highlight that, despite increased competition for treated jobs, salary transparency raises the probability that a representative woman is hired by a large firm. This occurs not because treated firms are more likely to hire women, but because women are equally likely to be hired once they apply – and considerably more likely to apply when salaries are disclosed.

13 Conclusion

This paper has shown that gendered sorting across firms arises from the interaction of three facts: large firms pay more but are less transparent about wages; they offer fewer flexible work arrangements; and women place somewhat greater weight than men on flexibility. When salaries are withheld, uncertainty discounts the value of high-wage, low-flexibility jobs, disproportionately reducing women's applications. Pay transparency removes this informational wedge, redirecting search toward higher-paying firms even though the underlying amenity landscape remains unchanged, thereby closing gender gaps in sorting. Résumé match quality declines slightly on average, but the top of the applicant pool improves, particularly with high-ability women. Treated large firms also become more likely to disclose salaries after the experiment, suggesting they overestimated the costs of screening and underestimated the benefits of attracting stronger applicants. The paper brings together administrative data from Pakistan's largest job portal, surveys of firms and job-seekers, a discrete choice experiment, and a large-scale field experiment that randomized mandated pay disclosure across over 20,000 job postings. A key takeaway is conceptual: gender gaps in sorting emerge from the joint presence of preferences and informational frictions, but information alone is sufficient to close them.

The results also help resolve a longstanding puzzle in the literature. Women consistently report stronger preferences for flexibility in hypothetical choice experiments, yet this does not map neatly onto observed application behavior. This paper shows why: outside the lab, wages and amenities covary, and the appeal of one attribute depends on trade-offs with the other. When wages are uncertain, flexibility becomes disproportionately valuable. Once wages are revealed, the perceived trade-off shifts, and women are more willing to substitute toward inflexible, high-wage jobs. Information thus reshapes not only beliefs about pay but also the willingness to trade flexibility for higher earnings, underscoring that women's flexibility preferences are neither innate nor immutable.

The findings carry clear policy implications. One approach to narrowing gender gaps in sorting is to alter the amenity mix offered by high-paying firms – for example, expanding flexible or remote work options. A cheaper and more immediately effective lever, however, is to improve pay information, which can redirect women's search and complement more costly structural reforms.

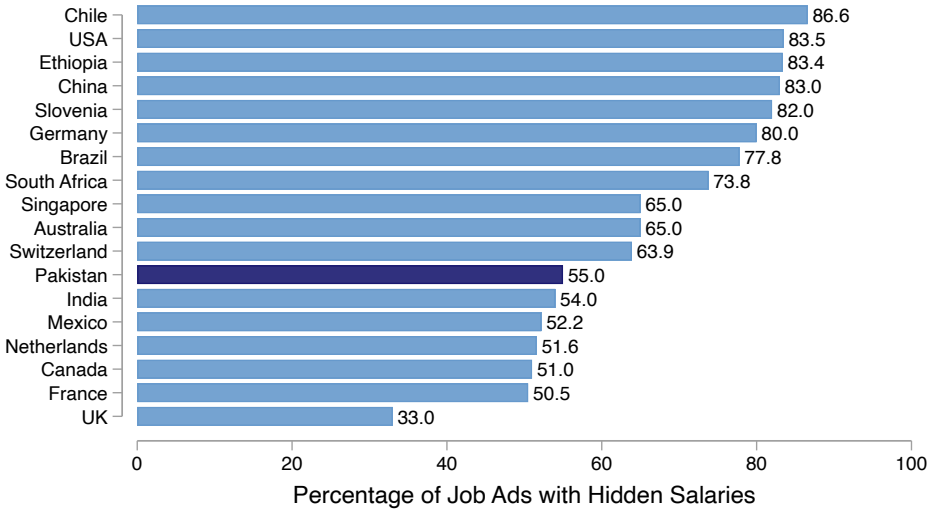
Most mechanisms identified here are likely to generalize: large firms tend to pay more but disclose less, provide fewer family-friendly amenities, and employ workers who dislike bargaining. A potential exception in Pakistan is that women's mobility constraints may make remote work especially pivotal. In other contexts, the under-provided amenity may differ – shorter work hours, more sick leave, or reduced workplace stress. Yet the broader mechanism remains the same: hidden pay amplifies the value of scarce amenities, shaping sorting patterns in gendered ways.

These findings also open new avenues for research. An important next step is to study how pay transparency influences bargaining behavior and wage trajectories, not just for applicants but also for incumbent workers who may renegotiate when salary information becomes public. Taken

together, the evidence shows that the recruitment practices of large, well-paying firms can inadvertently screen women out of promising careers. Improving the information environment does not eliminate gendered preferences for certain amenities or resolve deeper structural constraints. But by ensuring that workers choose between pay and flexibility with full information, transparency can narrow gender gaps in sorting and improve women's access to high-paying firms.

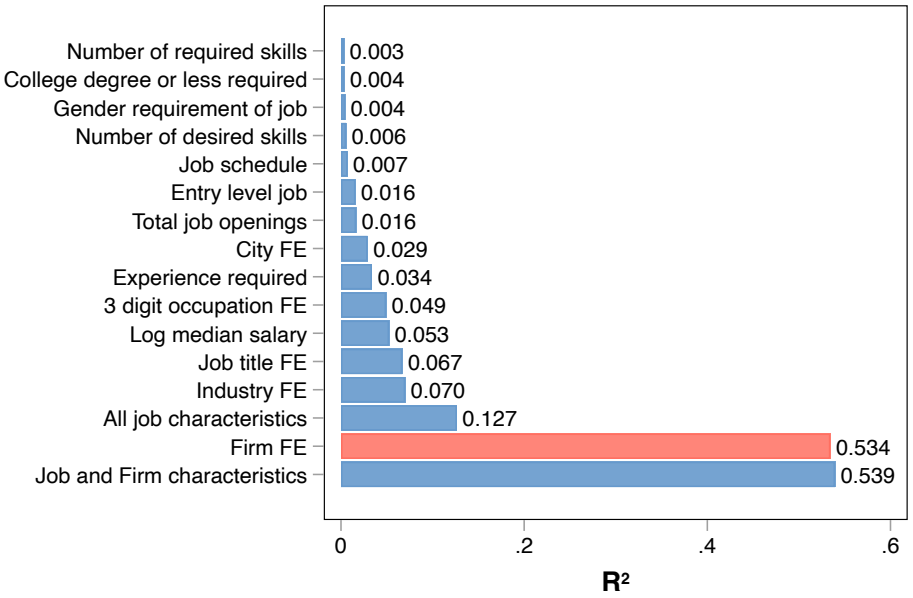
14 Figures

Figure 1: Salary non-disclosure globally



Notes: The figure shows percentage of job ads across different countries in which no salary information is stated. This information has been gathered from the following sources: Banfi and Villena-Roldán (2019); Batra et al. (2023); Brenčič (2012); Kuhn and Shen (2013); Skoda (2022); Indeed Hiring Lab (2023a, 2024, 2023b); HRD Asia (2023); People Matters (2023); AfriWorket (2024); Adzuna (2022) [Return to page 8]

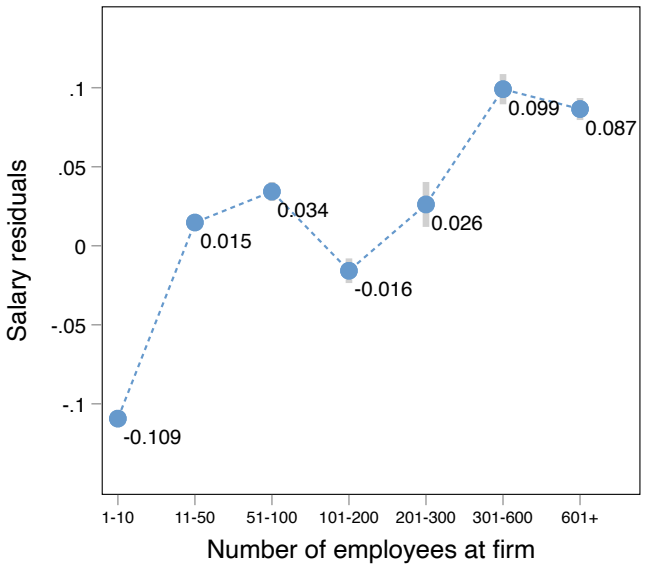
Figure 2: Explaining variation in salary non-disclosure



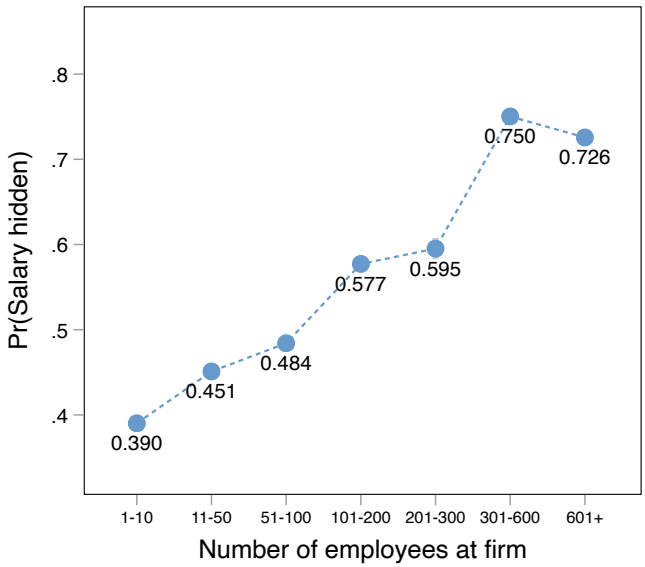
Notes: The figure shows R^2 from a job-level regression of an indicator for whether the salary of the job is hidden on each characteristic of the job, one at a time, until the third-last row in which all job characteristics are entered together. In the second to last row, firm fixed effects are controlled for by themselves, and the final row includes firm effects together with job characteristics. The sample consists of all baseline jobs. [Return to page 12]

Figure 3: Pay premiums and salary non-disclosure by firm size

(a) Firm-size pay premiums



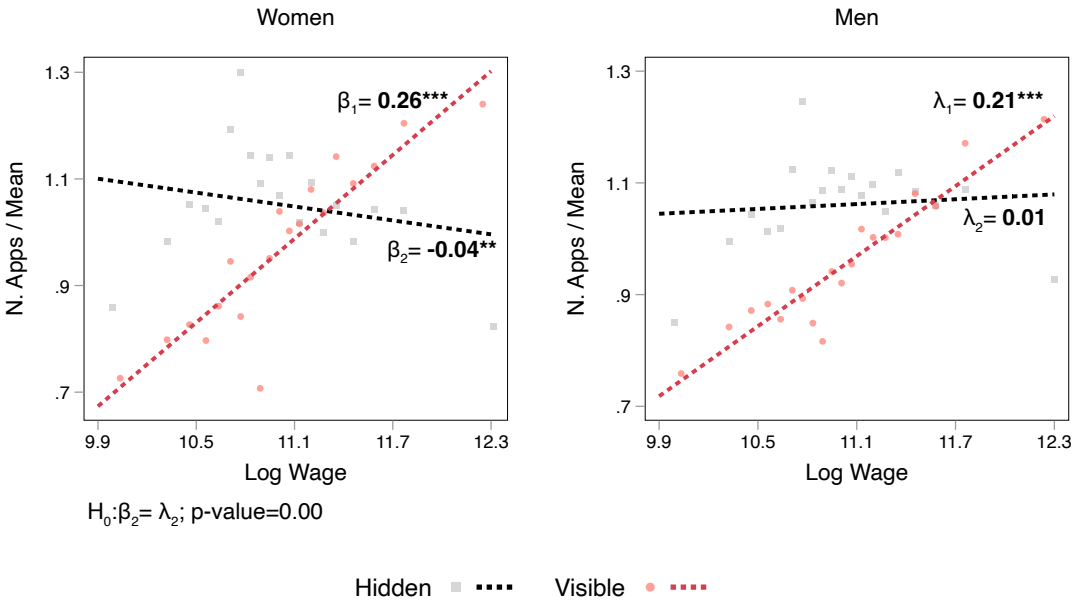
(b) Salary non-disclosure by firm size



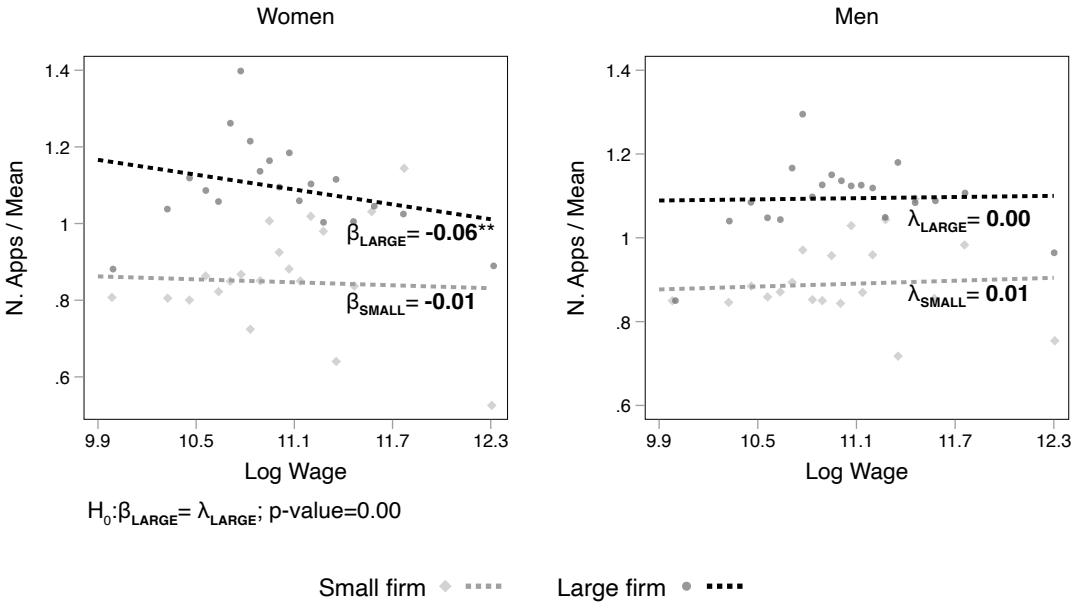
Notes: Panel A plots the relationship between firm size and salary premiums. Each point represents the predicted residual from a regression of log posted salary on detailed job-level controls (occupation, industry, city, experience and education requirements, career level, and work schedule), plotted by firm size bin. Residuals capture excess pay conditional on job characteristics. Panel B shows the predicted probability that a job posting hides salary information, estimated from a linear probability model with firm size dummies. The same job-level sample is used as in Panel A. Both panels use firm size categories defined by the number of employees, and error bars reflect 95% confidence intervals. [Return to page 11]

Figure 4: Gendered application response to disclosed and hidden salaries

(a) Application response by salary disclosure status

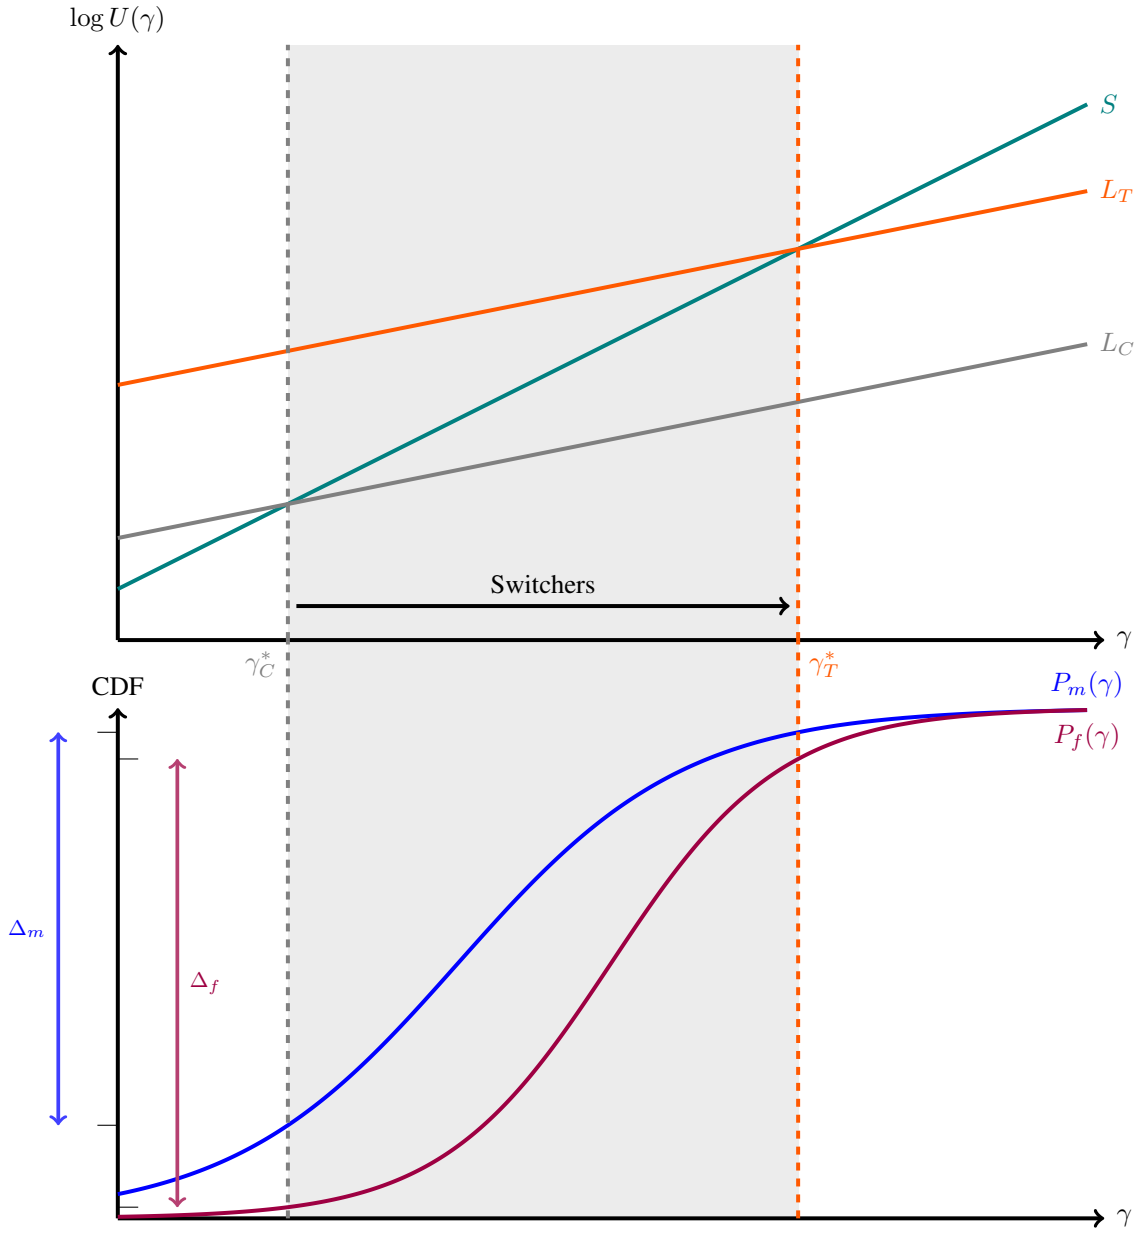


(b) Among jobs with undisclosed salaries: Application response by firm size



Notes: **Panel (a)** of the graph shows the relationship between the number of applications sent by men and women to a job, and the log maximum wage of the job, separately by whether the salary of the job is hidden (black) or visible (red) to the job-seeker. To account for differences in the number of male and female job-seekers, application counts are re-scaled by dividing with the gender-specific mean number of applications. A version of this graph with un-scaled application counts are shown in Appendix Figure A.8. The specification controls for industry, occupation and city fixed effects, and job characteristics (experience requirement, education requirement, career level of job, number of vacancies corresponding to the ad, number of required skills and job schedule). **Panel (b)** of the graph decomposes the application response to hidden salaries in Panel (a) into small- and large-firm jobs (gray and black colors, respectively). [Return to page 14 or 33.]

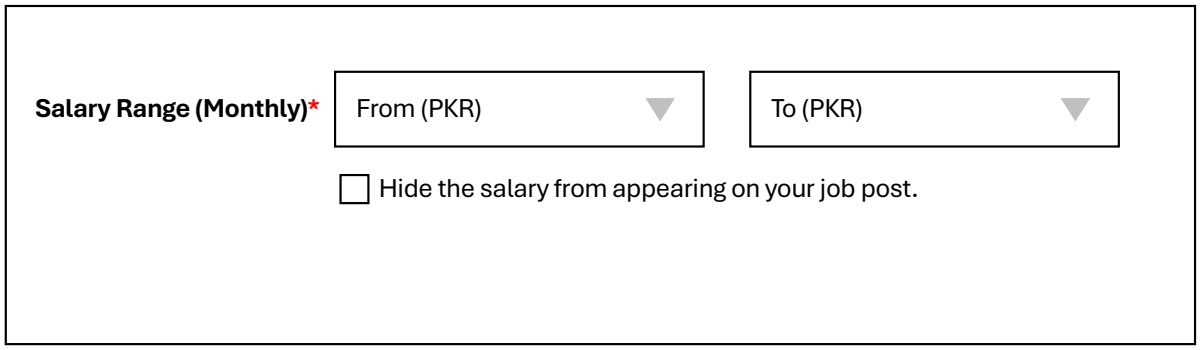
Figure 5: Theoretical framework



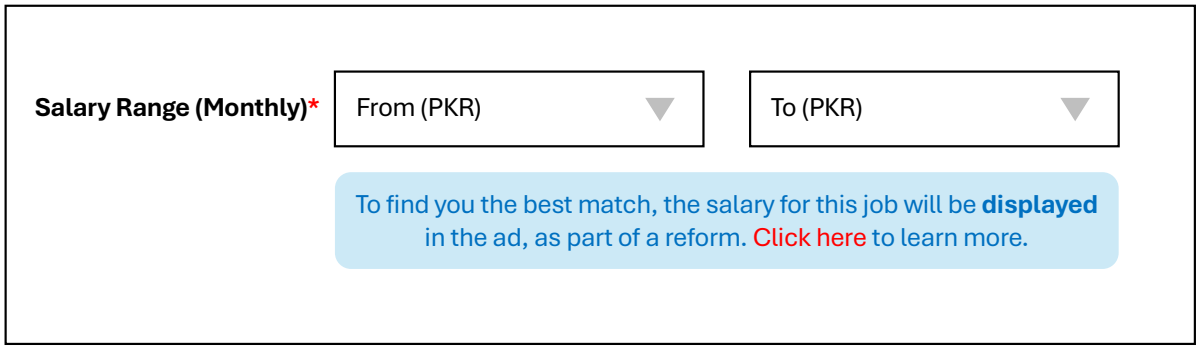
Notes: **Upper panel:** Utility functions by amenity preference γ. The blue line S represents small-firm utility. The gray line L_C shows large-firm utility under control (wage hidden), while the red line L_T shows large-firm utility under transparency (wage revealed). Transparency shifts the large-firm utility function upward, moving the indifference point from γ_C^* to γ_T^*. **Lower panel:** Cumulative distribution functions of amenity preferences by gender. Women (red curve) have stronger preferences for amenities than men (blue curve), reflected in first-order stochastic dominance. The gray shading highlights the switcher band: workers with $\gamma \in (\gamma_C^*, \gamma_T^*)$ who move from small to large firms under transparency. The vertical arrows show that more women than men fall in this switcher region ($\Delta_f > \Delta_m$). [Return to page 17]

Figure 6: Experiment Design

(a) Control Ads

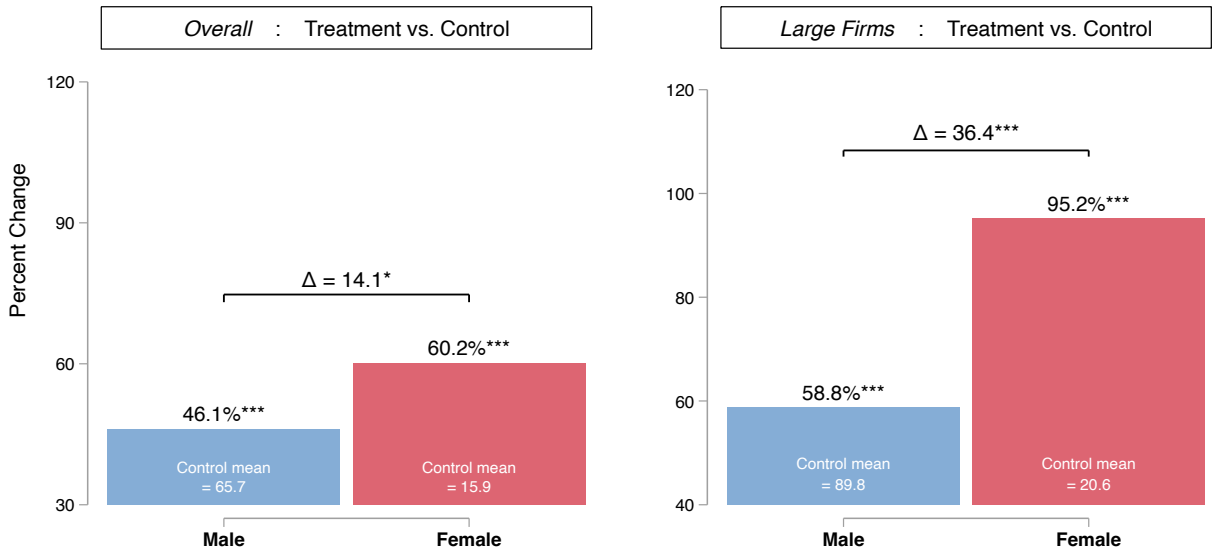


(b) Treatment Ads



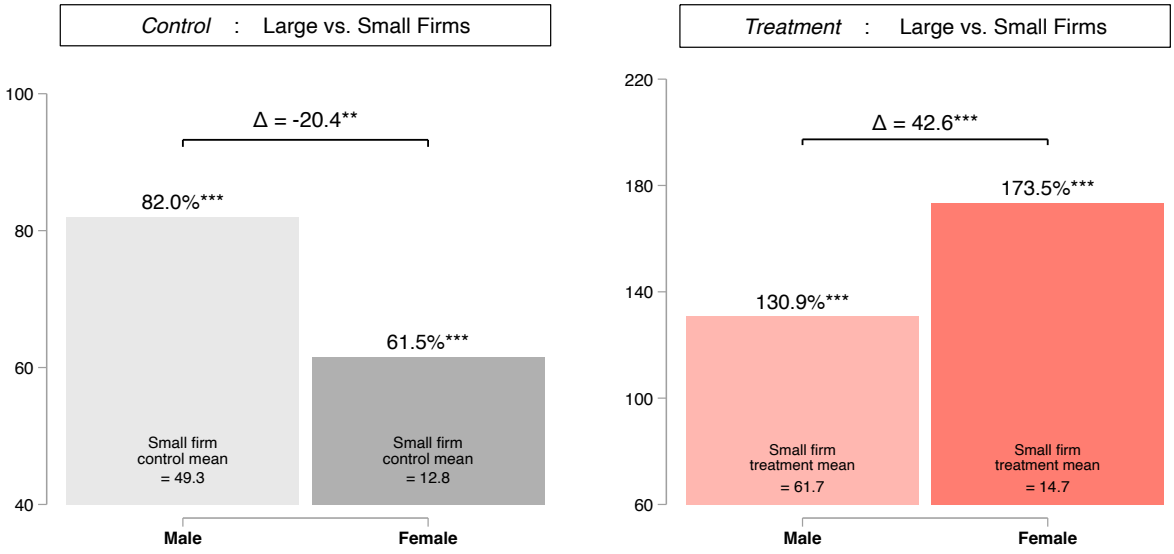
Notes: Panel (a) Control shows the standard job posting interface for specifying salary ranges. Employers can optionally check a box to hide the salary range from appearing in the job post. Panel (b) Treatment shows the modified interface for the treatment group, where salary ranges are displayed by default. Employers are informed via a notification that this transparency is part of a reform to improve job matches. Employers in this version cannot hide the salary range. [Return to page 20]

Figure 7: Treatment impacts on applications by gender



Notes: The figure shows treatment effects on application volumes by gender, both overall (left panel) and restricted to large-firm jobs (right panel). Blue bars represent male applications while red bars represent female applications. A firm is defined as large if it has more than 50 employees. The bars report percentage change estimates from a Poisson regression of application counts on treatment status, with control group means shown at the base of each bar. The brackets across bars report the gender gap, and its corresponding statistical significance. [Return to page 25]

Figure 8: Gender gap in directed search to large firms by treatment status



Notes: The figure illustrates how salary transparency affects the gender gap in applications to large versus small firms. Bars show the percent difference in application volumes to large firms relative to small firms, separately for men and women, using Poisson regression estimates. The gray bars represent the control group while the coral bars represent the treatment group. The brackets across bars report the gender gap, and its corresponding statistical significance. [Return to page 25]

15 Tables

Table 1: Valuations of Job Attributes in Discrete Choice Experiment

	Odds Ratio		Difference (Female-Male)	
Amenity	Female	Male	Percentage Points	*p*-value
Large firm	1.23**	1.65***	-6.24	0.043
	(0.12)	(0.18)	(2.55)	
Transport subsidy	1.69**	1.48*	1.50	0.681
	(0.36)	(0.32)	(4.92)	
Diversity message	1.53***	1.44***	0.73	0.653
	(0.14)	(0.15)	(2.56)	
Salary visibility	2.65***	2.14***	3.22	0.152
	(0.27)	(0.24)	(2.48)	
Remote work	1.38***	0.98	6.51	0.032
	(0.15)	(0.12)	(3.02)	
Observations	1946	1500		

Notes: This table reports conditional logit coefficients from a discrete choice experiment in which respondents choose between randomly generated pairs of stylized jobs. Columns 2-3 show odds ratios for the estimated effect of each amenity on the likelihood of a job being chosen, holding other attributes constant. Column 4 shows the gender difference in marginal effects (female minus male) in percentage points. Column 5 reports p-values from tests of whether the coefficients differ significantly between women and men. Standard errors are shown in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. [Return to page 13 or 30.]

Table 2: Treatment effects on salaries

	(1)	(2)	(3)
	Log Max Salary	**Log Min Salary**	**Salary Range: (Max-Min)/Median**
Treated Job	-0.01	0.03***	-0.03***
	(0.01)	(0.01)	(0.00)
Constant	11.24	10.77	0.43
Observations	19,890	19,890	19,890

Notes: The sample consists of all jobs in the experiment. Column 1 shows log maximum advertised salary. Column 2 shows log minimum advertised salary. Column 3 shows the width of the salary range by subtracting minimum from maximum salary and dividing by the midpoint of the range. Observations are missing for a handful of firms who do not report salaries to the platform because job ads are pulled from those firms' internal job posting sites where this information is not mandatory to report. $*p < 0.1, **p < 0.05, ***p < 0.01$. [Return to page 23]

Table 3: Treatment effects on workers' job search

Panel A: *Treatment effect on Incidence Rate Ratios (IRR) using Poisson*

	IRR: Views			**IRR: Applications**		
	(1) All	(2) Male (β_m)	(3) Female (β_f)	(4) All	(5) Male (β_m)	(6) Female (β_f)
Treated Job	1.45***	1.44***	1.52***	1.49***	1.46***	1.60***
	(0.10)	(0.09)	(0.14)	(0.11)	(0.09)	(0.17)
Observations	20,088	20,088	20,088	20,088	20,088	20,088
H_0: $\beta_f = \beta_m$; p-value			0.18			0.06

Panel B: *Treatment effect on counts using OLS*

	N. Views			**N. Applications**		
	(1) All	(2) Male (β_m)	(3) Female (β_f)	(4) All	(5) Male (β_m)	(6) Female (β_f)
Treated Job	49.91***	33.91***	13.04***	39.91***	30.30***	9.59***
	(10.58)	(6.86)	(3.36)	(8.22)	(5.75)	(2.52)
Control job mean	111.56	77.36	25.24	81.65	65.69	15.93

Notes: This table presents the treatment effects on workers' job search behavior, measured in terms of job views and applications. **Panel A** presents incidence rate ratios (IRR) estimated using a Poisson model (described in Equaltion 5), capturing the proportional change in views and applications due to treatment. An IRR greater than 1 indicates an increase in job search activity relative to control jobs, e.g., Column 1 suggests a 50% increase in views due to treatment. **Panel B** reports the estimated effects on the number of views and applications using ordinary least squares (OLS) as described in Equation 4. Columns 1 shows the treatment effect on total number of views, while Columns 2 and 3 show these effects separately for male (β_m) and female (β_f) workers. Columns 4-6 report similar estimates for the total number of applications submitted. The control job mean in each row provides the baseline views and application counts for comparison. The p-value for $H_0 : \beta_f = \beta_m$ tests whether the treatment effects from the Poisson specification differ by gender. Robust standard errors are in parentheses. For a **visual representation** of Columns 1 and 4, see Figure A.18. For similar graphs on gender specific results in Columns 2-3 and Columns 5-6, see Figure A.19. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. [Return to page 24]

Table 4: Treatment effects by firm size on workers' job search

Panel A: *Treatment effect as Incidence Rate Ratios (IRR) using Poisson*

	IRR: Views			**IRR: Applications**		
	(1) All firms	(2) Small firms (β_s)	(3) Large firms (β_l)	(4) All firms	(5) Small firms (β_s)	(6) Large firms (β_l)
Treated Job	1.45***	1.22***	1.59***	1.49***	1.23***	1.66***
	(0.10)	(0.05)	(0.17)	(0.11)	(0.07)	(0.18)
Observations	20,088	11,747	8,341	20,088	11,747	8,341
H_0: $\beta_l = \beta_s$; p-value			0.02			0.02

Panel B: *Treatment effect on counts using OLS*

	N. Views			**N. Applications**		
	(1) All firms	(2) Small firms (β_s)	(3) Large firms (β_l)	(4) All firms	(5) Small firms (β_s)	(6) Large firms (β_l)
Treated Job	49.91***	18.24***	89.88***	39.91***	14.38***	72.46***
	(10.58)	(3.91)	(24.37)	(8.22)	(3.62)	(18.75)
Control job mean	111.56	83.66	152.61	81.65	62.08	110.45

Notes: This table presents heterogeneous treatment effects by firm size on workers' job search behavior, measured in terms of job views and applications. **Panel A** presents incidence rate ratios (IRR) estimated using a Poisson model, capturing the proportional change in job views and applications due to treatment. An IRR greater than 1 indicates an increase in search activity relative to control jobs. For example, in Column 1, treatment increases job views by 50%. The *p*-value for $H_0 : \beta_l = \beta_s$ tests whether the treatment effects from the Poisson specification differ between small and large firms. **Panel B** reports the estimated effects on the number of views and applications using ordinary least squares (OLS). Column 1 shows the overall treatment effect on job views, while Columns 2 and 3 separate these effects for searches directed at small and large firms, respectively. Columns 4–6 follow the same structure for job applications. The control job mean in each row provides baseline views and application counts. A **visual representation** of the results in Columns 2-3 and 5-6 are in Figure A.20. Robust standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. [Return to page 24]

Table 5: Treatment effects on salaries by firm size

	(1) **Log Max Salary**	(2) **Log Min Salary**	(3) **Salary Range: (Max-Min)/Median**
Treated Job	-0.01	0.04***	-0.04***
	(0.01)	(0.01)	(0.01)
Large Firm	0.21***	0.32***	-0.10***
	(0.02)	(0.02)	(0.01)
Large Firm × Treated Job	-0.01	-0.02	0.01
	(0.02)	(0.02)	(0.01)
Constant	11.15	10.64	0.47
Observations	19,890	19,890	19,890

Notes: The sample consists of all jobs in the experiment. Column 1 shows log maximum advertised salary. Column 2 shows log minimum advertised salary. Column 3 shows the width of the salary range by subtracting minimum from maximum salary and dividing by the mid-point of the range. Observations are missing for a handful of firms who do not report salaries to the platform because job ads are pulled from those firms' internal job posting sites where this information is not mandatory to report. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 25]

Table 6: Amenities and firm size

Panel A: Control jobs					
	(1) Remote work	(2) Transport	(3) Flex work	(4) Safe envir.	(5) Equal opp. employer
Large firm	-0.05***	0.00	-0.02***	0.01***	0.01***
	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)
Small firm	0.10	0.02	0.05	0.02	0.02
Observations	8,745	8,745	8,745	8,745	8,745
Panel B: Treatment jobs					
	(1) Remote work	(2) Transport	(3) Flex work	(4) Safe envir.	(5) Equal opp. employer
Large firm	-0.06***	-0.00	-0.03***	0.01***	0.02***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Small firm	0.10	0.02	0.05	0.02	0.02
Observations	10,975	10,975	10,975	10,975	10,975

Notes: The table reports results from regressions examining how the likelihood of offering specific non-wage amenities varies with firm size and treatment. Amenities are detected in job descriptions via text analysis of job ads. Each coefficient represents the marginal effect of being a large firm (relative to a small firm), holding job-level controls constant: industry and occupation fixed effects, experience and education requirements, stated gender preferences for applicants, and career level of the job. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 27]

Table 7: Amenities and job search under disclosed vs. undisclosed salaries

	Control: Salary Hidden		Treatment	
	(1) N female apps	(2) N male apps	(3) N female apps	(4) N male apps
Large firm	1.13	1.28***	1.61***	1.42***
	(0.13)	(0.08)	(0.19)	(0.12)
Remote work	2.34***	1.59***	1.74***	1.51***
	(0.22)	(0.16)	(0.22)	(0.15)
Transport	1.17	1.02	0.56**	0.74**
	(0.23)	(0.21)	(0.13)	(0.11)
Flex work	1.01	0.87	0.65**	0.73***
	(0.15)	(0.11)	(0.13)	(0.09)
Safe environment	1.42	1.40*	1.03	1.01
	(0.33)	(0.27)	(0.17)	(0.12)
Equal opp. employer	1.70**	1.40**	1.27	1.24
	(0.38)	(0.22)	(0.31)	(0.19)
Dependent var. mean	18.31	72.07	25.52	95.98
Observations	3,851	3,851	10,975	10,975

Notes: The table reports Poisson regression estimates of how job-level amenities predict the number of applications submitted by women (Columns 1 and 3) and men (Columns 2 and 4), separately for control jobs with hidden salaries and treatment jobs. Coefficients reflect the relative change in applications associated with each amenity, controlling for other job characteristics. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 28]

Table 8: Preference for remote work by salary disclosure in large firms

Panel A: Women

	On-site	Remote	Difference
No salary	0.347***	0.479***	0.132***
	(0.025)	(0.034)	
Low salary	0.512***	0.526***	0.014
	(0.052)	(0.060)	
Medium salary	0.647***	0.694***	0.048
	(0.046)	(0.063)	
High salary	0.636***	0.705***	0.069
	(0.049)	(0.052)	
Observations			1042

Panel B: Men

	On-site	Remote	Difference
No salary	0.491***	0.473***	-0.018
	(0.031)	(0.038)	
Low salary	0.524***	0.532***	0.008
	(0.057)	(0.070)	
Medium salary	0.525***	0.625***	0.099
	(0.056)	(0.061)	
High salary	0.794***	0.672***	-0.122
	(0.046)	(0.069)	
Observations			820

Notes: Each table reports estimated preferences for remote versus on-site jobs within large firms, separately by gender and salary disclosure level. Panel A includes women; Panel B includes men. Each row corresponds to one of four salary levels: no salary disclosed, low, medium, or high salary. The columns show the predicted probability of choosing an on-site or remote job, as estimated from a linear probability model. The final column reports the difference in remote vs. on-site preferences at each salary level, along with significance stars. Predictions are generated from models that include interactions between salary level and a remote work indicator, controlling for other randomly assigned job attributes including diversity language and transport support. Coefficients represent the average marginal effect of each attribute on job choice. Standard errors are clustered at the choice set level. Results with a conditional logit specification are reported in Table B.7. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. [Return to page 30]

Table 9: Average applicant quality

Panel A: *Share of applicants who meet* ***experience*** *requirements*

	All Firms		Small Firms		Large Firms	
	(1) Male	(2) Female	(3) Male	(4) Female	(5) Male	(6) Female
Treated Job	-0.01**	-0.01*	-0.01**	-0.02***	-0.00	0.01
	(0.00)	(0.01)	(0.00)	(0.01)	(0.00)	(0.01)
Control Job Mean	0.71	0.60	0.74	0.64	0.68	0.55
Observations	19,879	17,889	11,666	10,467	8,213	7,422

Panel B: *Share of applicants who meet* ***education*** *requirements*

	All Firms		Small Firms		Large Firms	
	(1) Male	(2) Female	(3) Male	(4) Female	(5) Male	(6) Female
Treated Job	-0.01***	-0.01***	-0.01***	-0.00	-0.01***	-0.01***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Control Job Mean	0.88	0.91	0.89	0.92	0.88	0.91
Observations	19,883	17,973	11,670	10,532	8,213	7,441

Panel C: *Share of applicants who meet* ***skills*** *requirements*

	All Firms		Small Firms		Large Firms	
	(1) Male	(2) Female	(3) Male	(4) Female	(5) Male	(6) Female
Treated Job	-0.01*	-0.00	0.01*	0.01	-0.03***	-0.02**
	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Control Job Mean	0.55	0.50	0.56	0.51	0.54	0.49
Observations	20,009	18,501	11,707	10,828	8,302	7,673

Notes: The sample consists of all jobs in the experiment. Observations are missing in columns (1) and (2) all the applicants who applied had the relevant characteristic (experience, education or skills) missing. They may additionally be missing in columns (3)-(6) if no male or female applicant applied to that job. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 37]

Table 10: Selection into application pools of large firms

Panel A: High Ability Men

	(1) **Index Top Decile**	(2) **Education Top Decile**	(3) **GPA Top Quartile**	(4) **Years of Experience Top Decile**	(5) **Managed a Team**
Treated Job (β_m)	0.02***	0.02***	0.02***	0.01*	0.01**
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Control mean	0.92	0.93	0.90	0.85	0.93

Panel B: High Ability Women

	(1)	(2)	(3)	(4)	(5)
Treated Job (β_f)	0.07***	0.07***	0.08***	0.08***	0.07***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Control mean	0.63	0.71	0.54	0.54	0.56
Observations	8,341	8,341	8,341	8,341	8,341
H_0: $\beta_f = \beta_m$; p-value	0.00	0.00	0.00	0.00	0.00

Panel C: Low Ability Men

	(1) **Index Bottom Decile**	(2) **Education Bottom Decile**	(3) **GPA Bottom Quartile**	(4) **Years of Experience Bottom Decile**	(5) **Never managed a Team**
Treated Job (β_m)	0.01	-0.00	0.01*	0.01**	0.00
	(0.01)	(0.00)	(0.01)	(0.01)	(0.00)
Control mean	0.68	0.97	0.92	0.94	0.98

Panel D: Low Ability Women

	(1)	(2)	(3)	(4)	(5)
Treated Job (β_f)	0.03**	0.04***	0.06***	0.04***	0.04***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Control mean	0.48	0.81	0.61	0.73	0.84
Observations	8,341	8,341	8,341	8,341	8,341
H_0: $\beta_f = \beta_m$; p-value	0.12	0.00	0.00	0.00	0.00

Notes: The sample consists of job ads posted by large firms. Column 1 of Panels A and B shows the likelihood that at least one applicant in the top decile of a composite résumé index applied. This is an Anderson index of four traits: education, GPA, years of experience, and whether the applicant has ever managed a team. Panels C and D show the same for men and women in the bottom decile of expected wage. Columns 2–5 show effects for each of these traits separately. Panels A and B show results for high-ability applicants – those in the top decile (or top quartile, in the case of GPA) of the market-wide distribution for a given résumé trait. Panel A reports results for men; Panel B for women. Panels C and D report corresponding results for low-ability applicants – those in the bottom decile (or bottom quartile for GPA) of the same distributions. Panel C reports effects for men; Panel D for women. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 37]

Table 11: Characteristics of the top applicants

	(1) **Expected Salary (SDs)**	(2) **Education Level (SDs)**	(3) **Experience (SDs)**	(4) **GPA (SDs)**	(5) **Team Size Managed (SDs)**	(6) **At least 1 Female**
Panel A: All firms						
Treated Job	0.01	0.10***	0.07***	0.05***	0.13***	0.02***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Control Job Mean	-0.00	-0.00	0.00	-0.00	0.00	0.58
Observations	19,975	19,975	19,967	19,833	19,958	19,976
Panel B: Small firms						
	(1)	(2)	(3)	(4)	(5)	(6)
Treated Job	-0.00	0.07***	0.06***	0.05***	0.14***	0.01
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)
Control Job Mean	0.00	-0.00	0.00	0.00	-0.00	0.59
Observations	11,708	11,708	11,702	11,635	11,694	11,708
Panel C: Large firms						
	(1)	(2)	(3)	(4)	(5)	(6)
Treated Job	0.01	0.13***	0.07***	0.06**	0.09***	0.04***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)
Control Job Mean	0.00	0.00	-0.00	0.00	-0.00	0.56
Observations	8,267	8,267	8,265	8,198	8,264	8,268

Notes: The table presents the impact of treatment on the quality of the top applicants for a job. To identify top candidates, the top three values in each applicant pool are first determined for five key attributes: expected salary, education, experience, GPA, and team size managed. Applicants who rank in the top three for any of these attributes are flagged accordingly. Next, a score is constructed for each applicant, representing the proportion of these five attributes in which they rank among the top three. Finally, applicants are classified as the best in the pool if their score is among the top three in their applicant pool. Column 1 shows the average expected salary of these top candidates, Column 2 shows the average years of education, Column 3 shows average years of experience, Column 4 shows average GPA, and Column 5 shows size of the team managed. Each of these are normalized by the mean and standard deviation of the control group. Column 6 shows whether than at least one of the top candidates is a woman. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 38]

Table 12: Post experiment: salary disclosure by treatment exposure

	Firm ever hides salary post experiment		
	(1) **All firms**	(2) **Small firms**	(3) **Large firms**
Share of firm jobs treated × 10	-0.011*	-0.007	-0.025**
	(0.006)	(0.007)	(0.011)
Control Job Mean	0.63	0.54	0.81
Observations	642	382	260

Notes: This table reports the effect of firms' treatment exposure during the experiment on their post-experiment salary disclosure behavior. The outcome is an indicator equal to 1 if the firm hid salaries in any job posting during the month after the experiment, and 0 otherwise. The main independent variable is the share of the firm's jobs that were randomly assigned to the salary disclosure treatment during the experiment, scaled by 10 for interpretability. Coefficients therefore reflect the effect of a 10 percentage point increase in treatment exposure. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. [Return to page 39]

References

Adzuna, "Adzuna Salary Transparency White Paper," https://www.adzuna.com/blog/wp-content/uploads/2022/11/Adzuna-Salary-Transparency-White-Paper.pdf 2022. Accessed: 2024-12-23. [Return to page 44]

AfriWorket, "AfriWorket: Empowering Africa's Workforce," https://afriworket.com 2024. Accessed: 2024-12-23. [Return to page 44]

Arnold, David, Simon Quach, and Bledi Taska, "The Impact of Pay Transparency in Job Postings on the Labor Market," *SSRN Electronic Journal*, August 2022. [Referenced on pages 6, 21, and 24]

Ashraf, Nava, Oriana Bandiera, Edward Davenport, and Scott S. Lee, "Losing Prosociality in the Quest for Talent? Sorting, Selection, and Productivity in the Delivery of Public Services," *American Economic Review*, May 2020, *110* (5), 1355–1394. [Return to page 103]

_ , Xavier Giné, and Dean Karlan, "Finding Missing Markets (and a Disturbing Epilogue): Evidence from an Export Crop Adoption and Marketing Intervention in Kenya," *American Journal of Agricultural Economics*, 2009, *91* (4), 973–990. [Referenced on pages 10, 34, and 101]

Bandiera, Oriana, Amen Jalal, and Nina Roussille, "The Illusion of Time: Gender Gaps in Job Search and Employment," Working Paper 34051, National Bureau of Economic Research July 2025. [Return to page 7]

_ , Ananya Kotia, Ilse Lindenlaub, Christian Moser, and Andrea Prat, "Meritocracy across Countries," April 2024. 82 pages. Posted: 18 Apr 2024. Multiple versions available. [Return to page 36]

Banfi, Stefano and Benjamín Villena-Roldán, "Do High-Wage Jobs Attract More Applicants? Directed Search Evidence from the Online Labor Market," *Journal of Labor Economics*, 2019, *37* (3), 715–746. [Referenced on pages 6, 12, and 44]

Bassier, Ihsaan and Leila Gautham, "The firm-pay gender gap and formal sector churn over the life cycle," *Journal of Development Economics*, 2025, *176*, 103498. Short communication. [Return to page 1]

Batra, Honey, Amanda Michaud, and Simon Mongey, "Online Job Posts Contain Very Little Wage Information," Technical Report 31984, National Bureau of Economic Research December 2023. [Referenced on pages **6, 12, and 44**]

Benson, Alan, Danielle Li, and Kelly Shue, "Potential" and the gender promotion gap," *Unpublished working paper*, 2021. [Return to page **6**]

Bertrand, Marianne, "New Perspectives on Gender," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 4B, Amsterdam: North-Holland, 2011, pp. 1543–1590. [Return to page **35**]

_ **, Bruno Crépon, and Stefano Caria**, "Improving job seekers' employment and earnings through credible skills signals," 2024. Accessed: 2024-12-23. [Return to page **36**]

Biasi, Barbara and Heather Sarsons, "Flexible wages, bargaining, and the gender gap," *The Quarterly Journal of Economics*, 2022, *137* (1), 215–266. [Return to page **6**]

Blau, Francine D. and Lawrence M. Kahn, "The Gender Wage Gap: Extent, Trends, and Sources," *Journal of Economic Literature*, 2017, *55* (3), 789–865. [Referenced on pages **6 and 13**]

Brenčič, Vera, "Wage posting: evidence from job ads," *Canadian Journal of Economics/Revue canadienne d'économique*, 2012, *45* (4), 1529–1559. [Return to page **44**]

Breza, Emily, Supreet Kaur, and Yogita Shamdasani, "Labor Rationing," *American Economic Review*, 2021, *111* (10), 3184–3224. [Return to page **7**]

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez, "Inequality at work: The effect of peer salaries on job satisfaction," *American Economic Review*, 2012, *102* (6), 2981–3003. [Referenced on pages **5 and 16**]

_ **, Ana Rute Cardoso, and Patrick Kline**, "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *The Quarterly Journal of Economics*, 2016, *131* (2), 633–686. [Referenced on pages **1 and 5**]

_ **,** _ **, Joerg Heining, and Patrick Kline**, "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 2018, *36* (1), S13–S70. [Return to page **11**]

_ , **Jörg Heining, and Patrick Kline**, "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, 2013, *128* (3), 967–1015. [Return to page 11]

Chandrasekhar, Arun G, Melanie Morten, and Alessandra Peter, "Network-based hiring: Local benefits; global costs," Technical Report, National Bureau of Economic Research 2020. [Return to page 37]

Colonnelli, Emanuele, Joacim Tåg, Michael Webb, and Stefanie Wolter, "A Cross-Country Comparison of Dynamics in the Large Firm Wage Premium," *AEA Papers and Proceedings*, May 2018, *108*, 323–327. [Return to page 11]

Cullen, Zoë and Ricardo Perez-Truglia, "How much does your boss make? The effects of salary comparisons," *Journal of Political Economy*, 2022, *130* (3), 766–822. [Referenced on pages 5 and 16]

Cullen, Zoë B, "Is Pay Transparency Good?," *Journal of Economic Perspectives*, 2024, *38* (1), 153–180. [Return to page 8]

Cullen, Zoë B. and Bobak Pakzad-Hurson, "Equilibrium Effects of Pay Transparency," *Econometrica*, 2023, *91* (3), 765–802. [Referenced on pages 6 and 13]

Frimmel, Wolfgang, Bernhard Schmidpeter, Rene Wiesinger, and Rudolf Winter-Ebmer, "Mandatory wage posting, bargaining and the gender wage gap," Working Paper 2202, Linz 2022. [Return to page 6]

Garlick, Robert, Erica Field, and Kate Vyborny, "Women's Mobility and Labor Supply: Experimental Evidence from Pakistan," 2025. [Return to page 6]

Goldin, Claudia, "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 2014, *104* (4), 1091–1119. [Referenced on pages 6 and 13]

Goux, Dominique and Eric Maurin, "Persistence of interindustry wage differentials: a reexamination using matched worker-firm panel data," *Journal of labor Economics*, 1999, *17* (3), 492–533. [Return to page 11]

Gruetter, Max and Rafael Lalive, "The importance of firms in wage determination," *Labour Economics*, 2009, *16* (2), 149–160. [Return to page 11]

Hansen, Benjamin and Drew McNichols, "Information and the persistence of the gender wage gap: early evidence from California's salary history ban," Technical Report, National Bureau of Economic Research 2020. [Return to page 6]

Ho, Lisa, Suhani Jalota, and Anahita Karandikar, "Bringing work home: Flexible work arrangements as gateway jobs for women in West Bengal," Technical Report, Technical Report 2024. [Return to page 6]

HRD Asia, "Pay transparency: More Singapore job postings including salary offers," June 2023. Accessed: 2024-04-15. [Return to page 44]

ILO, *Global Wage Report 2018/19: What Lies Behind Gender Pay Gaps*, Geneva: International Labour Office, 2018. [Return to page 7]

_ , "Labour force participation rate by sex and age (%) – Annual," https://ilostat.ilo.org 2025. Based on the 13th ICLS definitions; includes countries using the 19th ICLS standards for comparability. [Return to page 7]

Indeed Hiring Lab, "Pay transparency: More Singapore job postings including salary offers," July 2023. Accessed: 2024-04-15. [Return to page 44]

_ , "Salaires dans les offres d'emploi : l'importance croissante de la transparence en Europe," April 2023. Accessed: 2024-04-15. [Return to page 44]

_ , "Salary Transparency Increasingly the Norm in Canadian Job Postings," March 2024. Accessed: 2024-04-15. [Return to page 44]

Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer, "Worker beliefs about outside options," *The Quarterly Journal of Economics*, 2024, *139* (3), 1505–1556. [Referenced on pages 5 **and** 16]

Jayachandran, Seema, "Microentrepreneurship in developing countries," *Handbook of labor, human resources and population economics*, 2021, pp. 1–31. [Return to page 37]

Kuhn, Peter and Kailing Shen, "GENDER DISCRIMINATION IN JOB ADS: EVIDENCE FROM CHINA," *The Quarterly Journal of Economics*, 2013, *128* (1), 287–336. [Referenced on pages **6** **and** **44**]

Lavetti, Kurt, "Compensating wage differentials in labor markets: Empirical challenges and applications," *Journal of Economic Perspectives*, 2023, *37* (3), 189–212. [Referenced on pages **5** **and** **16**]

Leibbrandt, Andreas and John A List, "Do women avoid salary negotiations? Evidence from a large-scale natural field experiment," *Management Science*, 2015, *61* (9), 2016–2024. [Return to page **5**]

Lucas, Robert E., "On the Size Distribution of Business Firms," *Bell Journal of Economics*, 1978, *9* (2), 508–523. [Referenced on pages **6** **and** **11**]

Marinescu, Ioana and Ronald Wolthoff, "Opening the Black Box of the Matching Function: The Power of Words," *Journal of Labor Economics*, 2020, *38* (2), 535–568. [Return to page **71**]

Mas, Alexandre and Amanda Pallais, "Valuing Alternative Work Arrangements," *American Economic Review*, 2017, *107* (12), 3722–3759. [Referenced on pages **5** **and** **13**]

_ **and** _ , "Alternative work arrangements," *Annual Review of Economics*, 2020, *12* (1), 631–658. [Return to page **5**]

Michelacci, Claudio and Javier Suarez, "Incomplete Wage Posting," *Journal of Political Economy*, 2006, *114* (6), 1098–1123. [Return to page **37**]

Moore, Henry Ludwell, *Laws of Wages: An Essay in Statistical Economics*, New York: Macmillan Company, 1991. [Referenced on pages **6** **and** **11**]

Morchio, Iacopo and Christian Moser, "The Gender Pay Gap: Micro Sources and Macro Consequences," CEPR Discussion Paper DP16383, Centre for Economic Policy Research 2021. [Return to page **1**]

Muralidharan, Karthik and Abhijeet Singh, "India's New National Education Policy: Evidence and Challenges," 2023. Accessed via RISE. [Return to page **36**]

Niederle, Muriel and Lise Vesterlund, "Do women shy away from competition? Do men compete too much?," *The quarterly journal of economics*, 2007, *122* (3), 1067–1101. [Referenced on pages **6** **and** **13**]

_ , Hessel Oosterbeek et al., "Gender, competitiveness, and career choices," *The Quarterly Journal of Economics*, 2014, *129* (3), 1409–1447. [Referenced on pages **6** **and** **13**]

Oi, Walter Y. and Todd L. Idson, "Firm Size and Wages," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 3, Amsterdam: Elsevier, 1999, chapter 33, pp. 2165–2214. [Return to page **11**]

People Matters, "Pay transparency on the rise in India, 50postings include salary details," May 2023. Accessed: 2024-04-15. [Return to page **44**]

Roussille, Nina, "The Role of the Ask Gap in Gender Pay Inequality," *Working Paper*, August 2023. [Return to page **6**]

Skoda, Samuel, "Directing Job Search in Practice: Mandating Pay Information in Job Ads," *Job Market Paper*, November 2022. [Referenced on pages **4**, **6**, **8**, **24**, **37**, **and** **44**]

Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter, "Firming up inequality," *The Quarterly journal of economics*, 2019, *134* (1), 1–50. [Referenced on pages **5** **and** **16**]

Sorkin, Isaac, "The Role of Firms in Gender Earnings Inequality: Evidence from the United States," *American Economic Review: Papers & Proceedings*, 2017, *107* (5), 384–387. [Return to page **1**]

Troske, Kenneth R, "Evidence on the employer size-wage premium from worker-establishment matched data," *Review of Economics and Statistics*, 1999, *81* (1), 15–26. [Referenced on pages **6** **and** **11**]

Vattuone, Giulia, "Worker Sorting and the Gender Wage Gap," 2023. Job Market Paper. [Return to page **1**]

Wiswall, Matthew and Basit Zafar, "Preference for the Workplace, Investment in Human Capital, and Gender," *The Quarterly Journal of Economics*, 2018, *133* (1), 457–507. [Referenced on pages **5** **and** **13**]

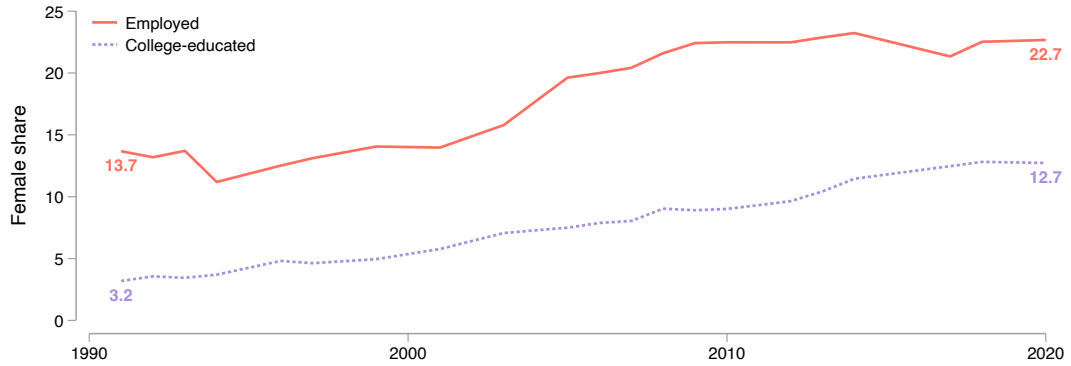
World Economic Forum, "Global Gender Gap Report 2025," Technical Report, World Economic Forum June 2025. 19th annual Global Gender Gap Index benchmarking 148 economies. [Return to page 7]

A Appendix

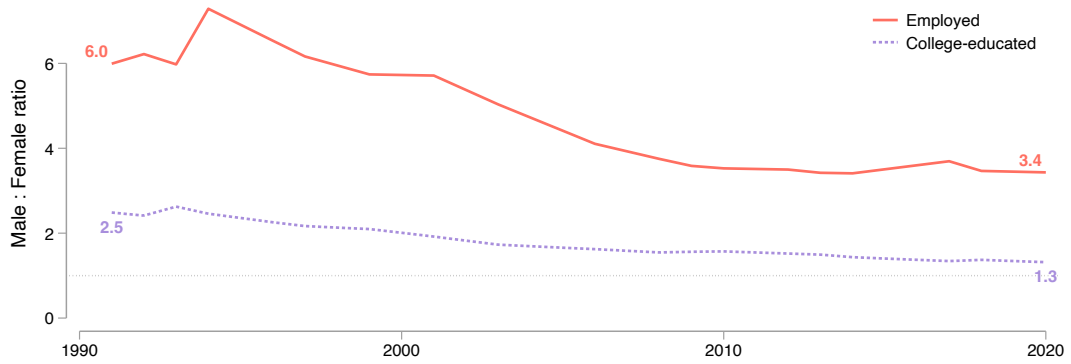
A.1 Appendix Figures

Figure A.1: Trends in women's education and employment in Pakistan

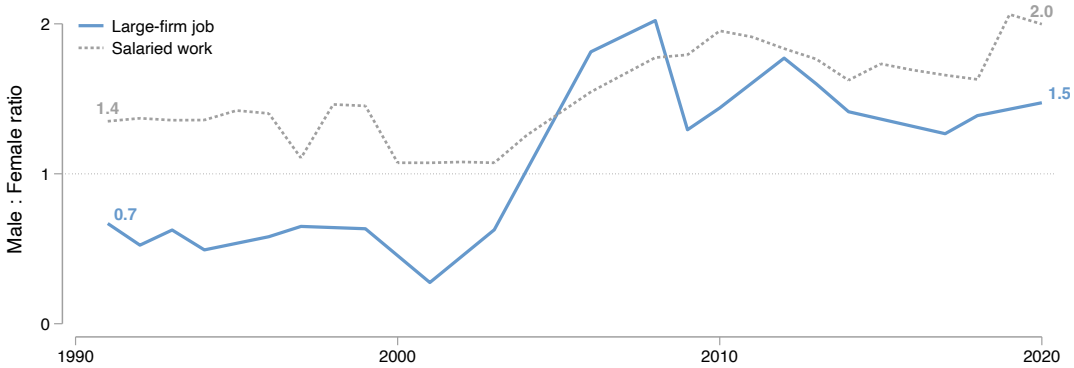
(a) Female share employed and college-educated



(b) Gender gaps in employment and college-education

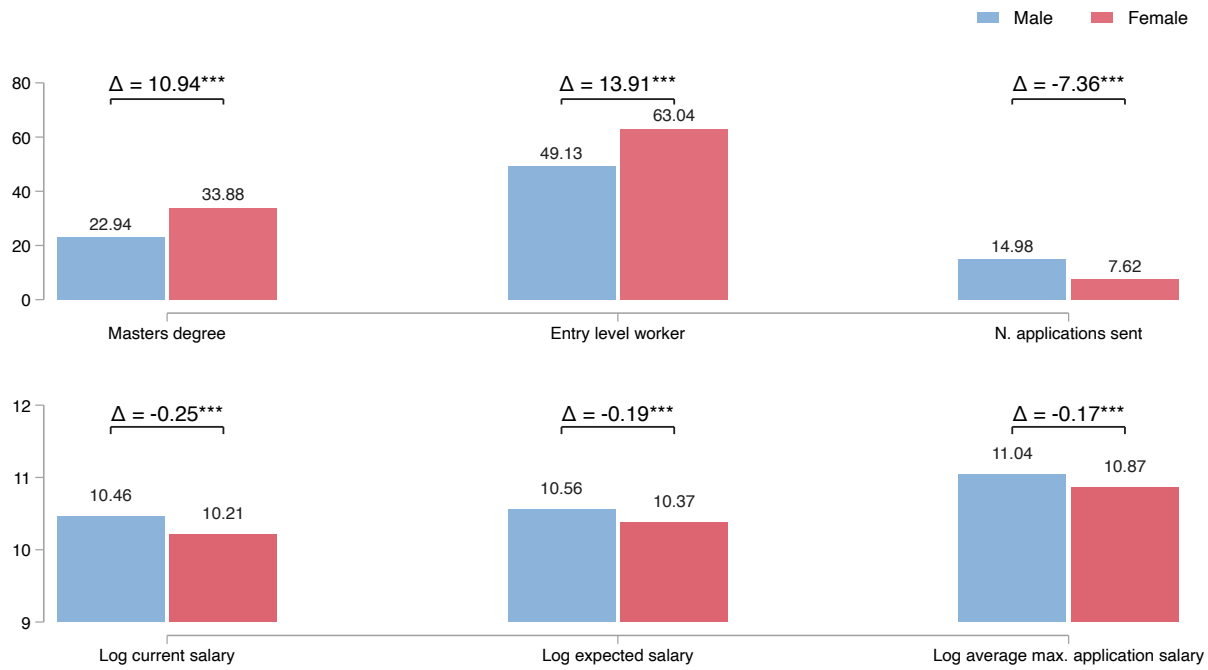


(c) Gender gaps in large-firm and salaried employment



Notes: Panel (a) plots the share of adult women that employed or have college-education. Panel (b) plots the male-to-female ratio of share of adults employed and college-educated. Panel (c) plots the male-to-female ratio of share of employed adults in large-firm and salaried work. Salaried work statistics are from ILO; all other statistics are from the Pakistan Labor Force Surveys. [Return to page 7]

Figure A.2: Gender differences in worker characteristics on the platform



Notes: The figure shows average differences in characteristics between male and female job-seekers. The blue bars represent male while red bars represent female job-seekers. The sample consists of 1.9 million job-seekers registered on the job platform at baseline. In the first row of the figure, the first two bars show the likelihood that a job-seeker has a Master's degree, the next two bars show differences in career levels of the job-seekers (specifically whether they self-report their career level to be entry level versus a more senior role), and the final two bars show differences in number of job applications sent by the average male and female applicant. The bottom row shows gender differences across three salary dimensions, accounting for occupation fixed effects: current salary, expected salary and the maximum salary of the average job candidates applied to. [Return to page 9]

Figure A.3: First Word of Job Title

(a) Jobs with Hidden Salaries

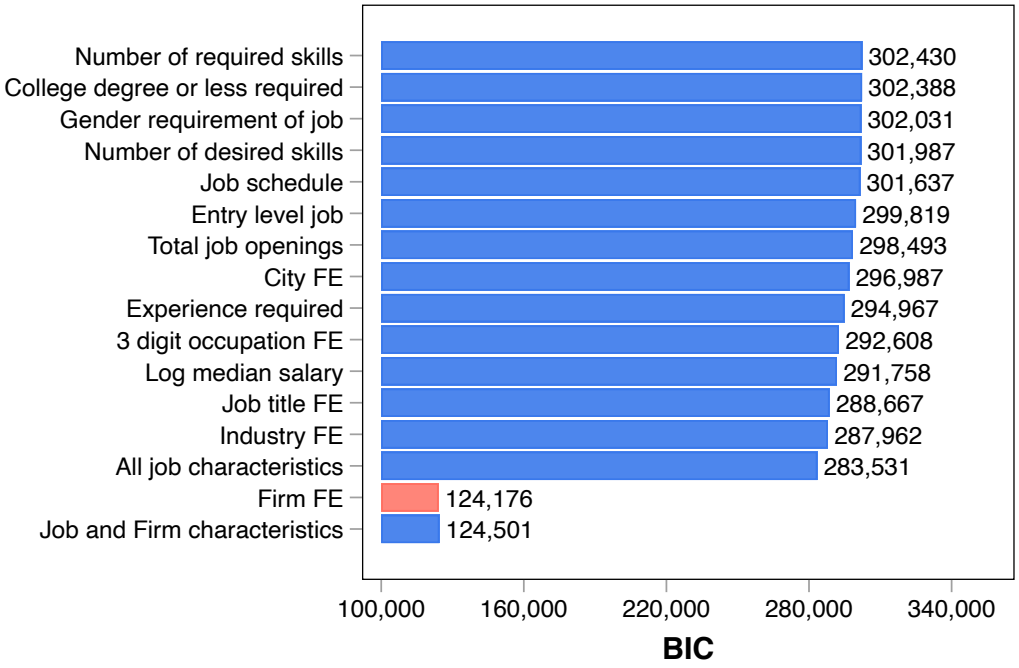


(b) Jobs with Visible Salaries



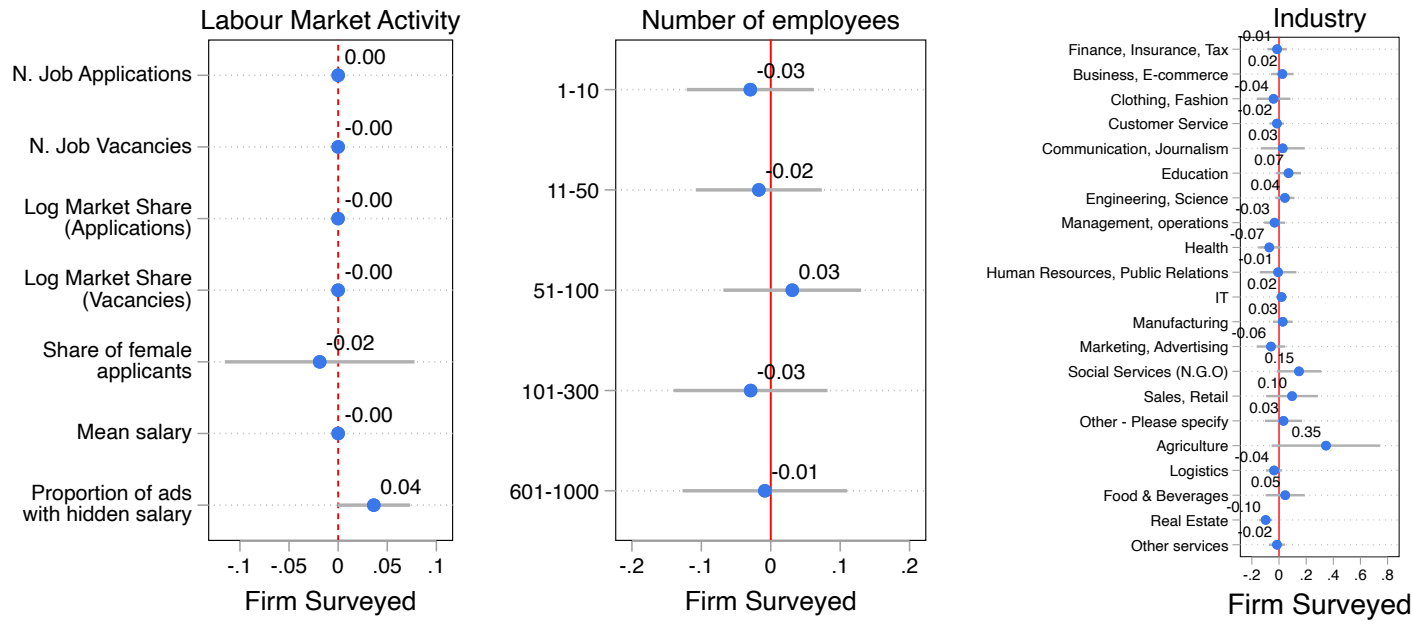
Notes: The figure shows the forst word of job titles separately for jobs with hidden salaries (panel a) and jobs with visible salaries (panel b). The job titles are cleaned following the approach in Marinescu and Wolthoff (2020). [Return to page 8]

Figure A.4: Firm effects explain most of the variation in hiding



Notes: The figure shows the Bayesian Information Criterion (BIC) from a job-level regression of an indicator for whether the salary of the job is hidden on each characteristic of the job, one at a time, until the third-last row in which all job characteristics are entered together. In the second to last row, firm fixed effects are controlled for by themselves, and the final row includes firm effects together with job characteristics. The sample consists of all baseline jobs. [Return to page 12]

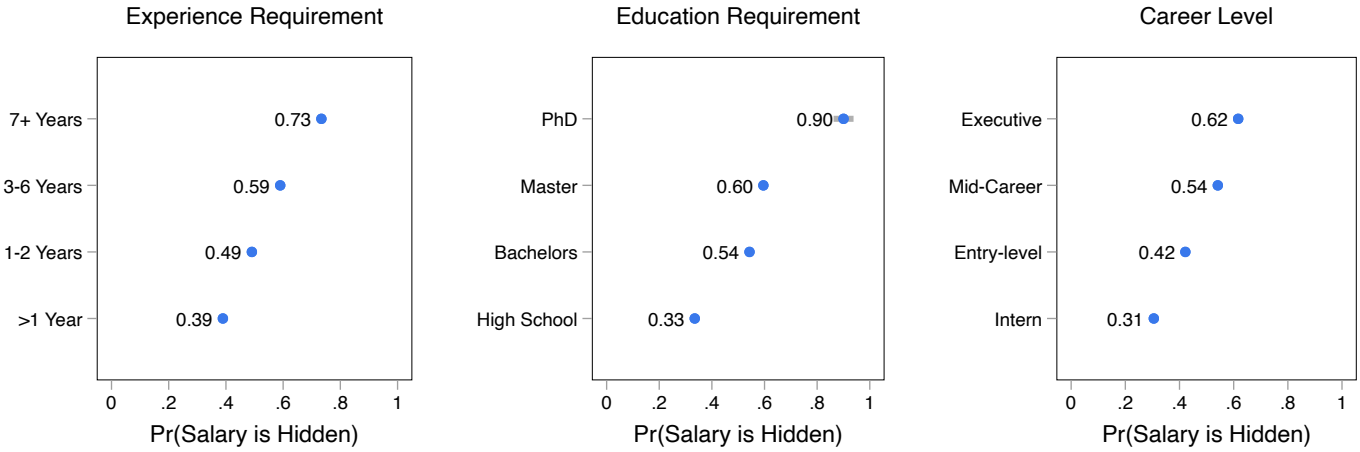
Figure A.5: Selection into firm survey



F-test of joint significance: 1.37; p-val=0.07

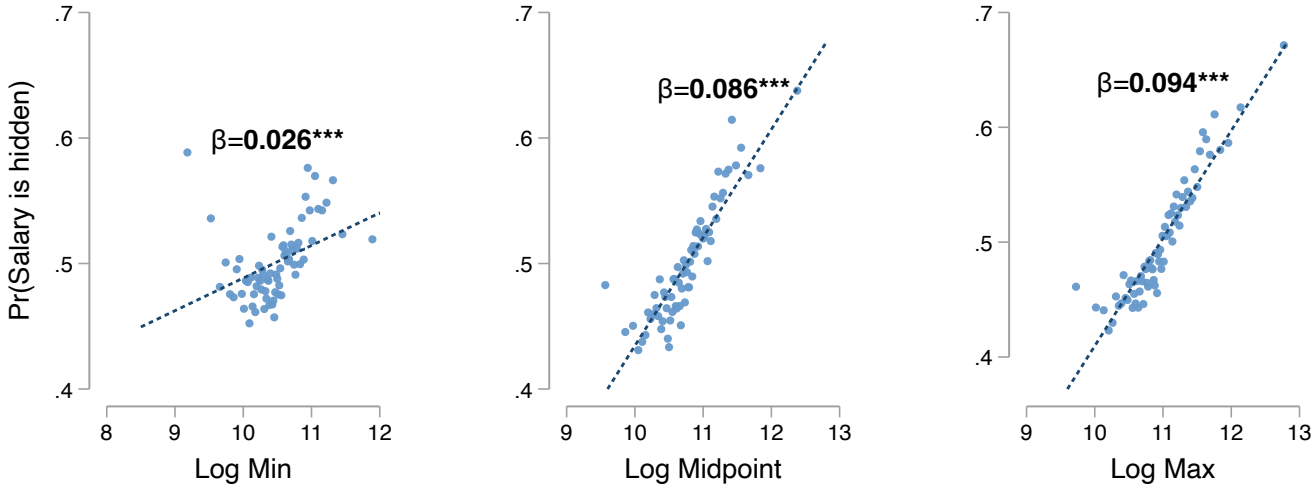
Notes: The figure tests whether there was selection in the firms who consented to participate in the survey. The first subfigure tests whether there is differential selection into the survey by the number of job applications received by the firm, the number of job vacancies, the log market share of applications and vacancies, the share of female applicants, average salary across jobs posted by the firm, and the proportion of jobs at the firm with salary hidden. The second subfigure tests whether response rate varies by firm size. The third sub figure tests whether specific industries are more or less likely to participate in the survey. The figure also reports an F-test of joint significance. [Return to page 10]

Figure A.6: Job complexity and salary non-disclosure



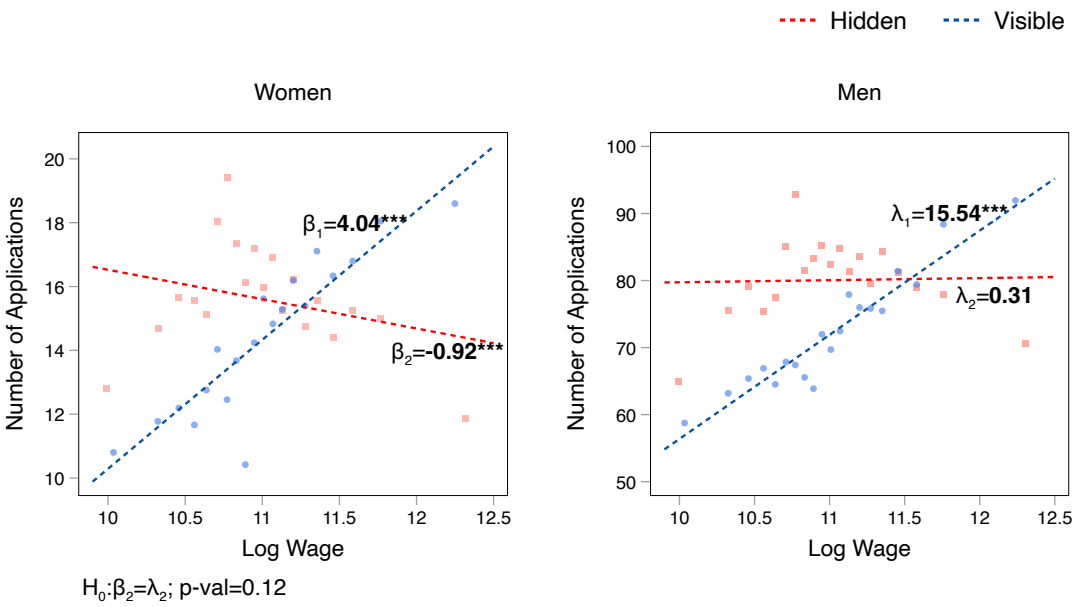
Notes: The figure shows correlations at the job level between an indicator for whether the salary is hidden, and various job characteristics. The sample consists of all baseline jobs. [Return to page 12]

Figure A.7: Salary levels and non-disclosure



Notes: The figure shows the job-level relationship between an indicator for whether the salary of the job is hidden, and the log minimum, median and maximum salary of the job, residualized on industry, occupation and city fixed effects, and job characteristics (experience requirement, education requirement, career level of job, number of vacancies corresponding to the ad, number of required skills and job schedule). The median salary corresponds to the midpoint of the range. The sample consists of all baseline jobs. [Return to page 12]

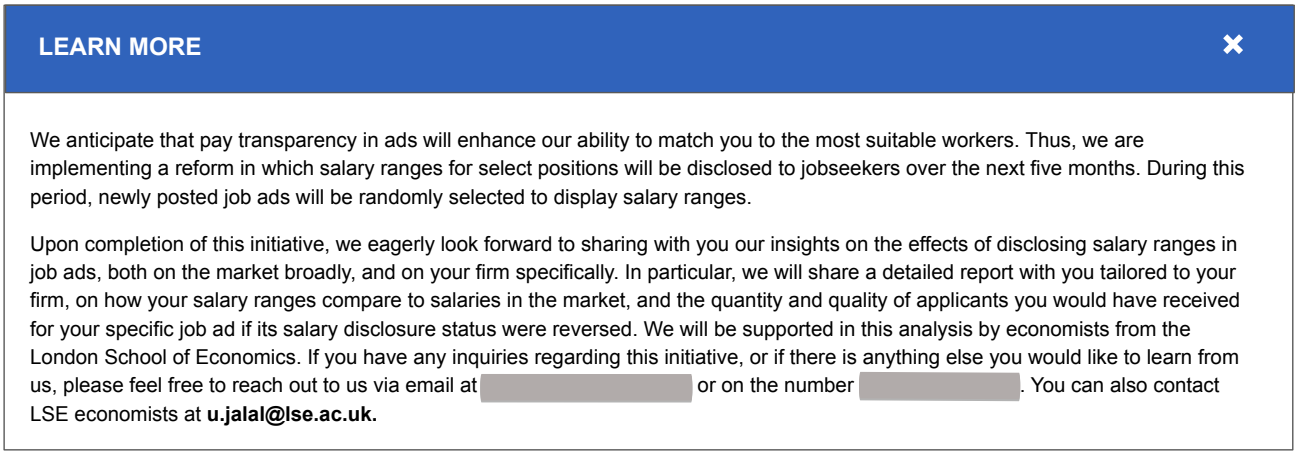
Figure A.8: Gendered application response to disclosed and hidden salaries (application counts)



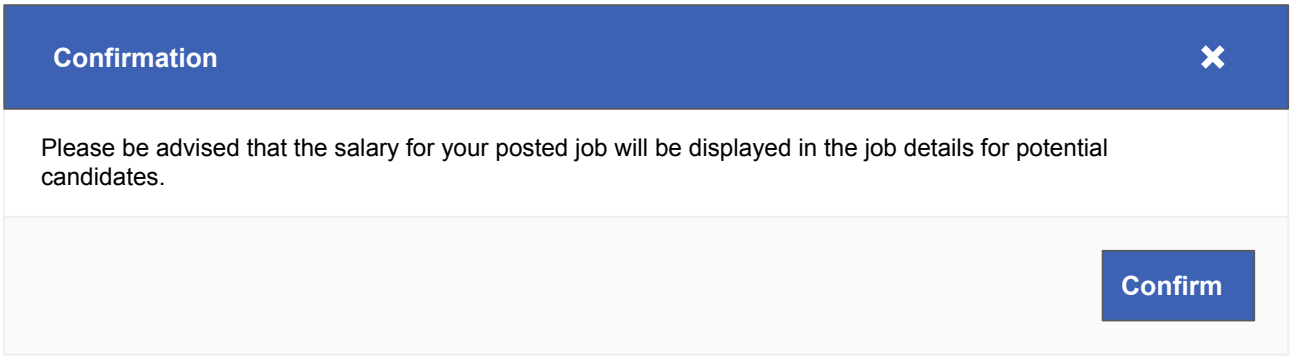
Notes: The graph shows the relationship between the number of applications sent by men and women to a job, and the log maximum wage of the job, separately by whether the salary of the job is hidden or visible to the job-seeker. It controls for industry, occupation and city fixed effects, and job characteristics (experience requirement, education requirement, career level of job, number of vacancies corresponding to the ad, number of required skills and job schedule). [Return to page 46]

Figure A.9: Further information and inattention

(a) Further information for treated ads

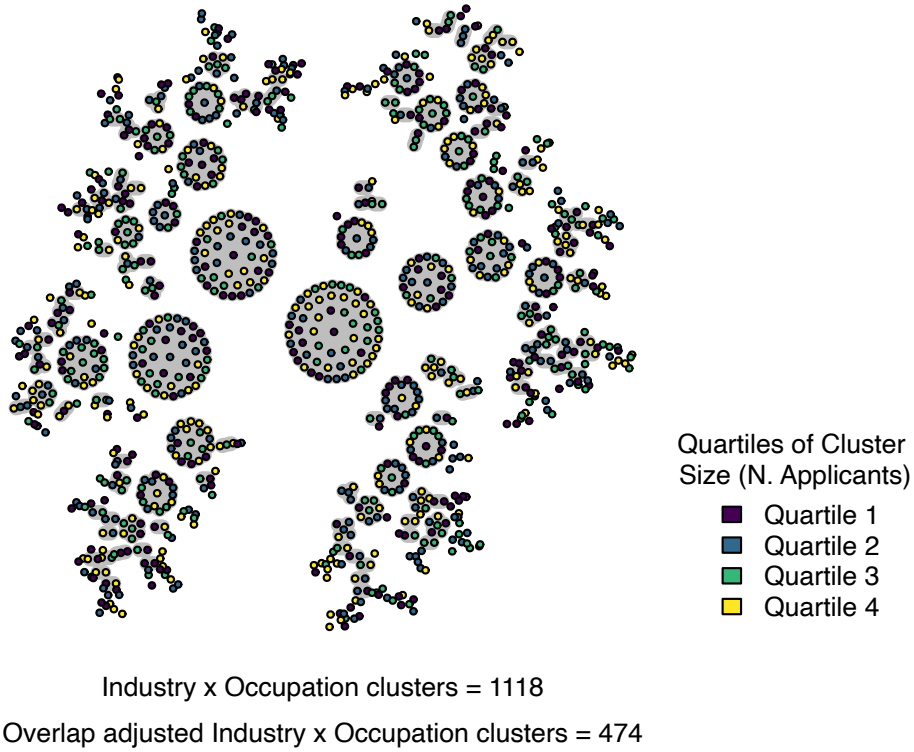


(b) Inattention



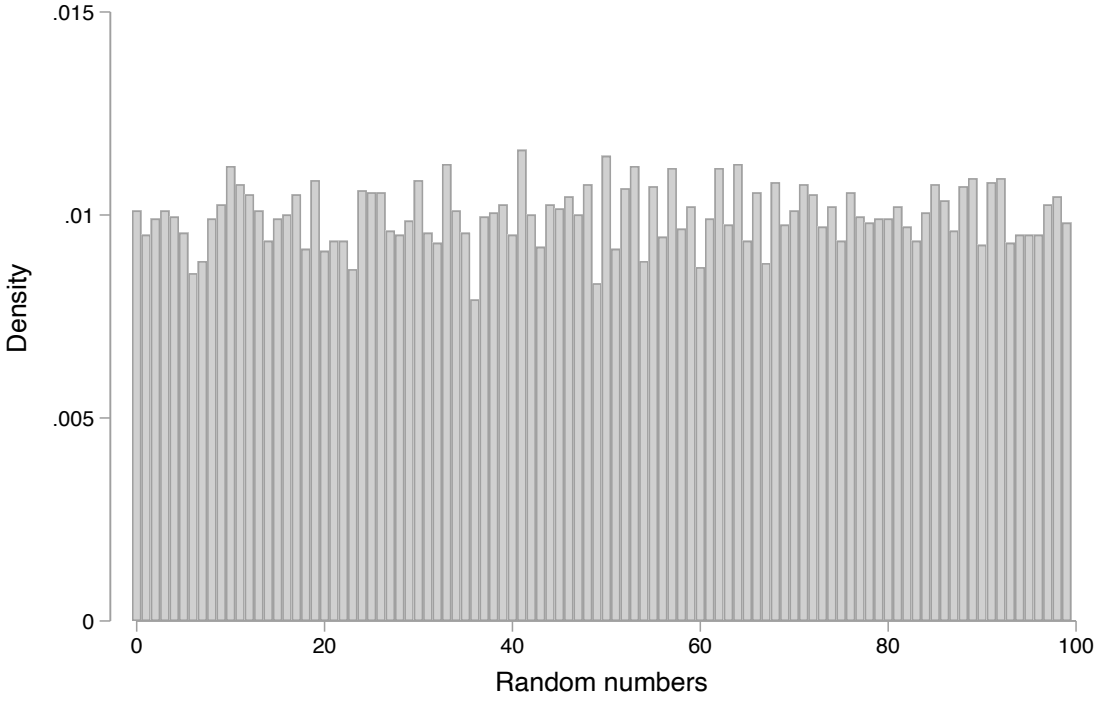
Notes: Panel (a) shows the "Learn More" pop-up provided to employers posting treated ads. It details the reform's purpose, emphasizing improved candidate matching and transparency. The pop-up invites employers to reach out with inquiries and mentions the involvement of researchers from the London School of Economics. Panel (b) illustrates an inattention check, where employers in the treatment group receive a confirmation prompt explicitly stating that the salary range will be visible to job-seekers. This ensures employers understand the implications before finalizing their job postings. [Return to page 20]

Figure A.10: Merged industry-occupation clusters based on historical overlap of applicants



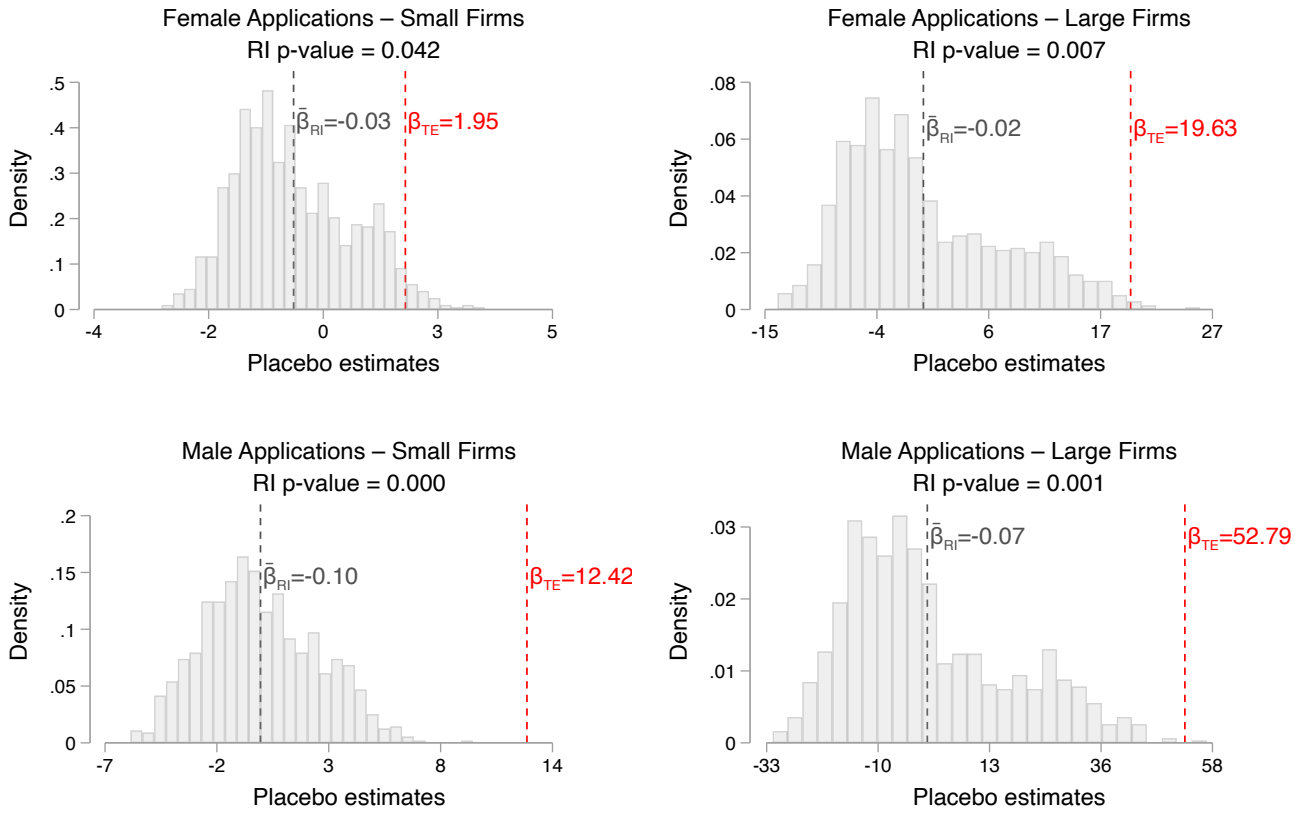
Notes: This figure visualizes the 474 labor market clusters used in the saturation design, formed by merging 1,118 initial industry–occupation cells based on historical application data. Industry-occupation cells with more than 10% mutual overlap in their applicant pools were merged to ensure that clusters are relatively self-contained in terms of job-seeker flows. Each node represents a unique industry-occupation cluster. The grey circles represent all the industry-occupation clusters that were combined into one. Nodes are color-coded by quartile of applicant volume to illustrate variation in collapsed clusters. [Return to page 21]

Figure A.11: Uniformity of Random Numbers



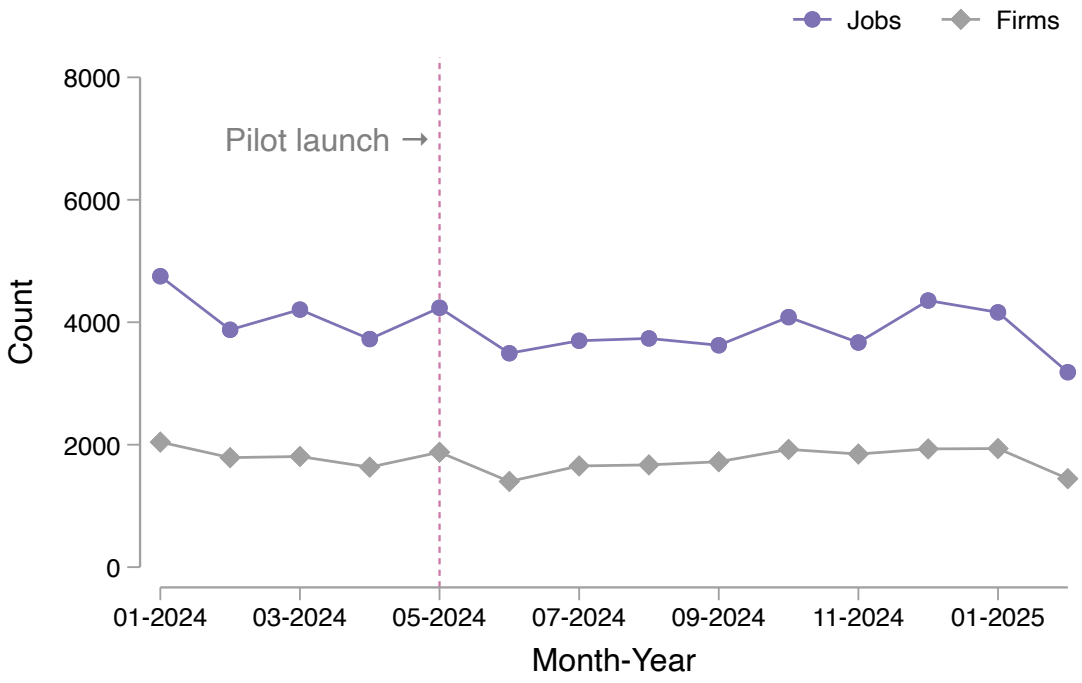
Notes: Distribution of the random numbers used to assign treatment across all job postings. Jobs were assigned to treatment if their random number—drawn from a continuous uniform distribution over $[0, 100]$—was less than or equal to a pre-specified threshold: 75 in high-saturation clusters and 25 in low-saturation clusters. These thresholds were determined by cluster-level treatment saturation targets defined in the pre-analysis plan. [Return to page 20]

Figure A.12: Randomization Inference



Notes: Each panel shows the distribution of placebo estimates generated by 1,000 re-randomizations of the treatment assignment within firm size. The vertical black dashed line indicates the mean placebo effect ($\bar{\beta}_{RI}$); the red dashed line indicates the actual treatment effect (β_{TE}). RI p-values indicate the proportion of placebo estimates with absolute values greater than or equal to the observed effect. [Return to page 26]

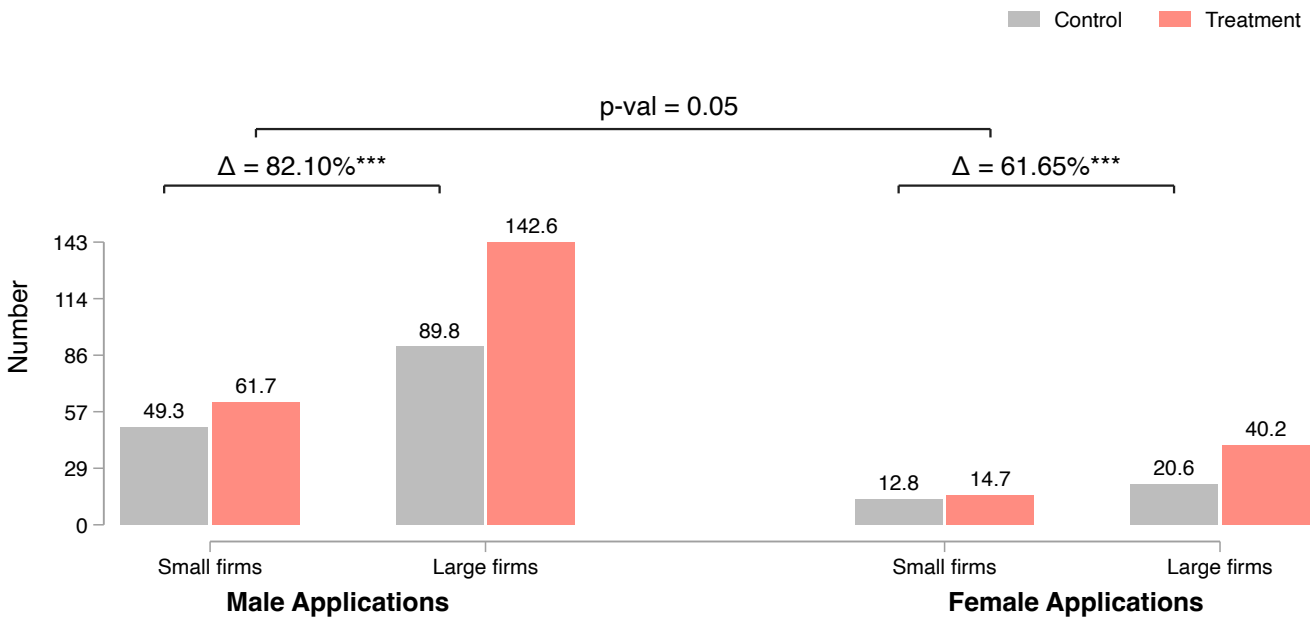
Figure A.13: No evidence of job posting or firm exits following experiment launch



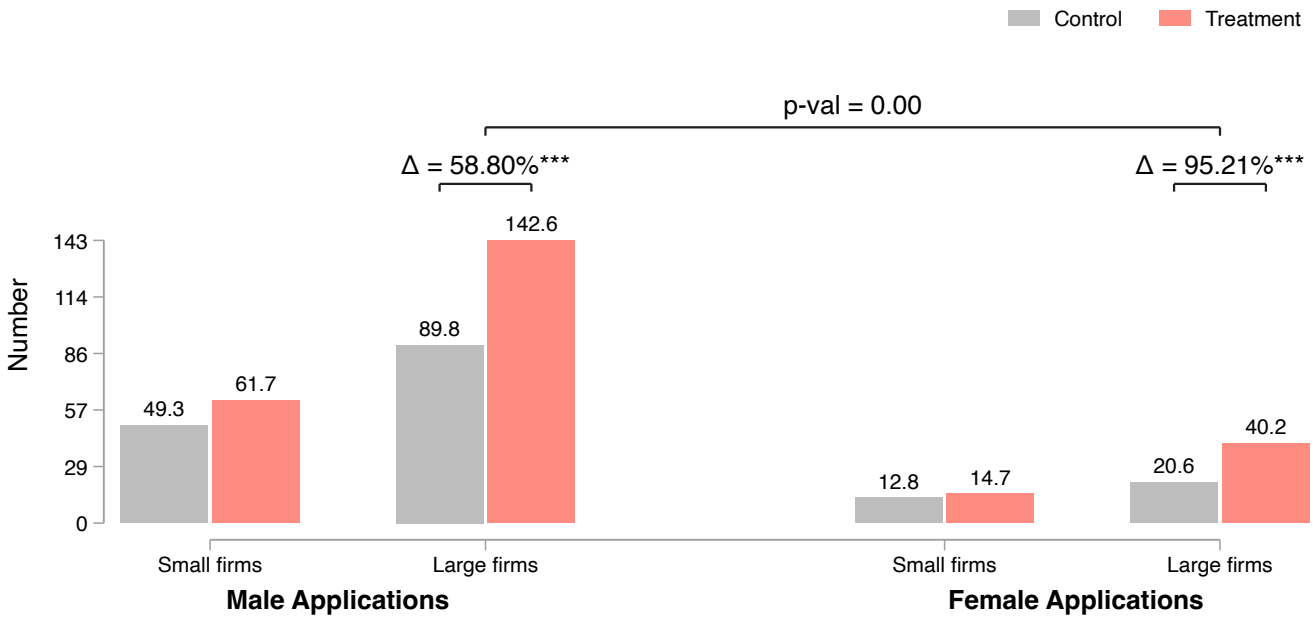
Notes: The figure shows the monthly counts of job postings (purple line) and active firms (orange line) on the platform from January 2024 to November 2024. The vertical dashed line marks the launch of the pilot intervention in May 2024. There is no evidence of a significant decline in the number of job postings or firms following the intervention, indicating no discernible exit of participants from the platform. [Return to page 23]

Figure A.14: Treatment impacts on job applications by gender (corresponding Table: B.6)

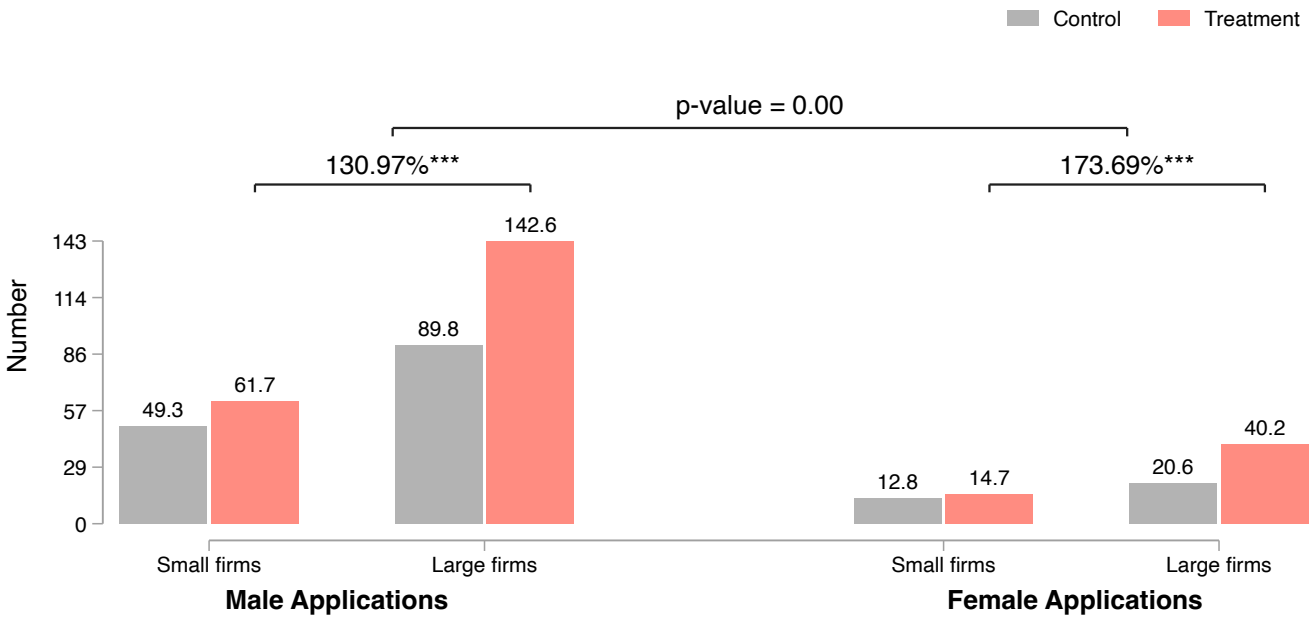
(a) Women's directed search to large firms is weaker than men's



(b) Reallocation of search is larger for women

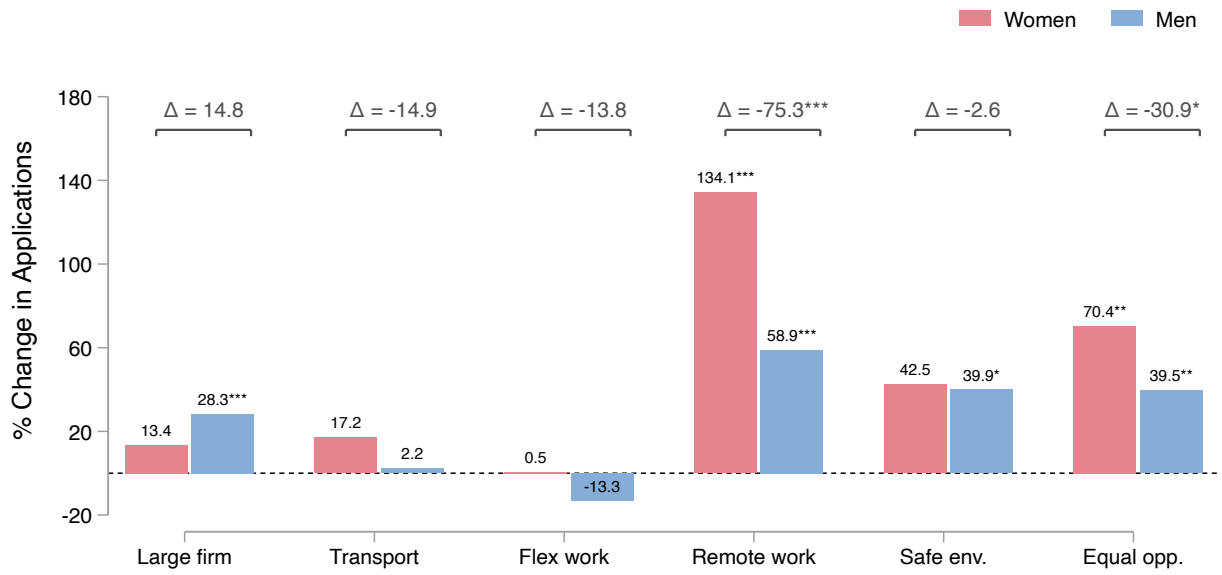


(c) Treatment closes gender gaps



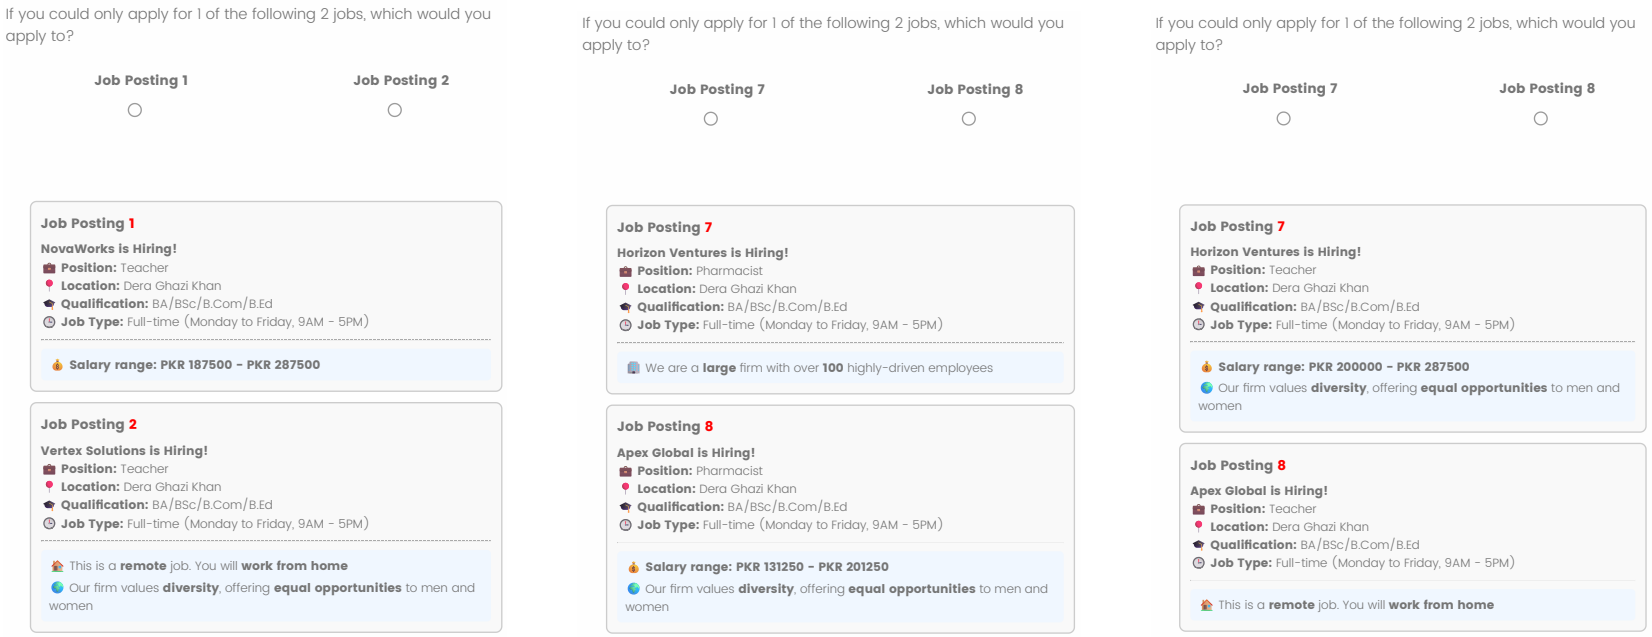
Notes: The figure shows treatment effects on applications with heterogeneity by firm size and gender. A tabular version of these results can be found in Appendix **Table** B.6. A firm is defined as large if it has more than 50 employees. The bars represent control and treatment group means. The lower brackets represent percentage change while the top brackets test whether the gender difference in percentages is significantly different. Panel (a) compares control group differences between men and women in directing search to large firms, using the reduced form specification. Panel (b) shows whether the number of applications to treated jobs at large firms increases relative to control jobs at large firms, and whether the relative gender differences are significant. Panel (c) shows whether treatment induces a reallocation of search within the treatment group from small to large firms, and whether the magnitudes of reallocation are significantly different by gender. [Return to page 25 or 92]

Figure A.15: Gender differences in appeal of amenities in the control group



Notes: This figure displays the gender gap in application responses to job-level amenities when salaries are not disclosed in the control group. Bars show the percentage change in applications associated with each amenity, estimated separately for women (red) and men (blue) using Poisson regressions, controlling for job characteristics and the other amenities shown in the graph. The gap above each pair of bars reflects the difference in responsiveness between women and men, and the corresponding significance. [Return to page 28]

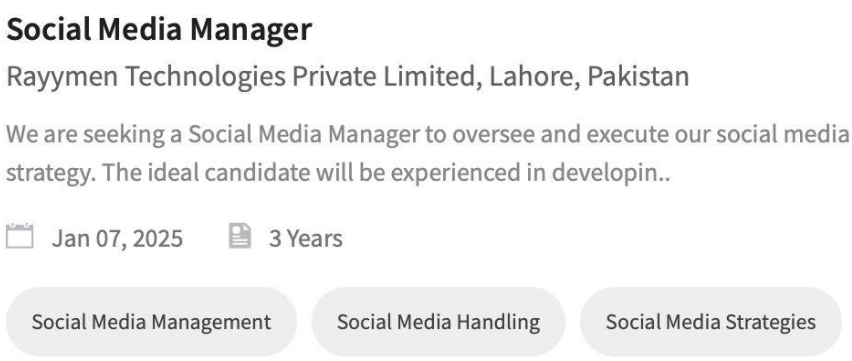
Figure A.16: Examples of vignettes from discrete choice experiment



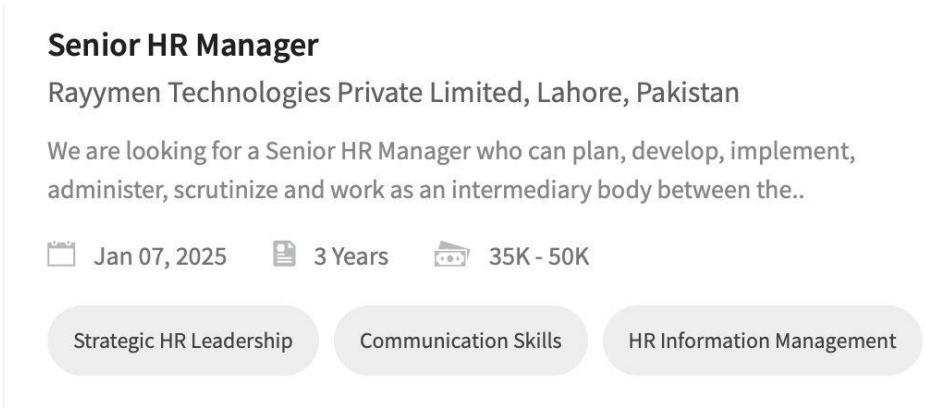
Notes: Each panel presents an example vignette from the discrete choice experiment shown to job-seekers, who were asked to choose between two job postings. Vignettes randomly vary salary visibility, non-wage amenities (remote work, transport support, and equal opportunity language), and a large firm label. When visible, salaries are randomly drawn between 110–120% of the respondent's expected wage. These randomized traits are always displayed in the blue shaded region. The grey region above this contains the job title, location, required education level, and work schedule of the job – attributes held constant across the two job profiles in each choice. The first three are tailored to each respondent based on their current or preferred job characteristics. Additionally, each job is attributed to a fictitious firm name designed to be neutral and uninformative about the firm's industry or other characteristics. The left panel presents a trade-off between a displayed salary range and a job offering remote work with equal opportunity language. The center panel shows a case where only one job discloses salary, contrasting a large-firm signal ("over 100 highly-driven employees") with a equal opportunity language. The right panel displays a choice between two jobs where the first discloses salary and includes equal opportunity language, while the second offers remote work. [Return to page 29]

Figure A.17: Short previews of jobs

(a) Job with hidden salaries

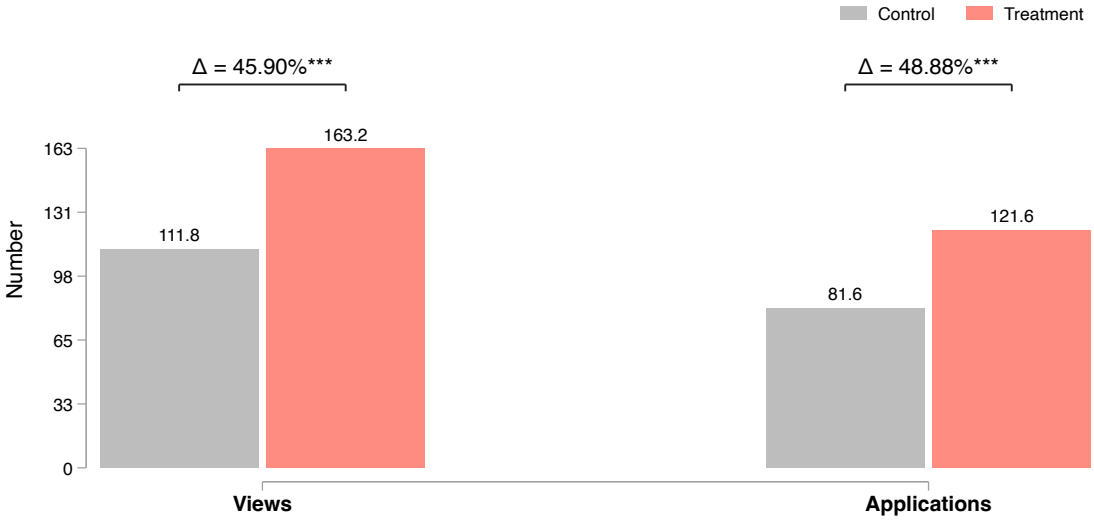


(b) Job with visible salaries



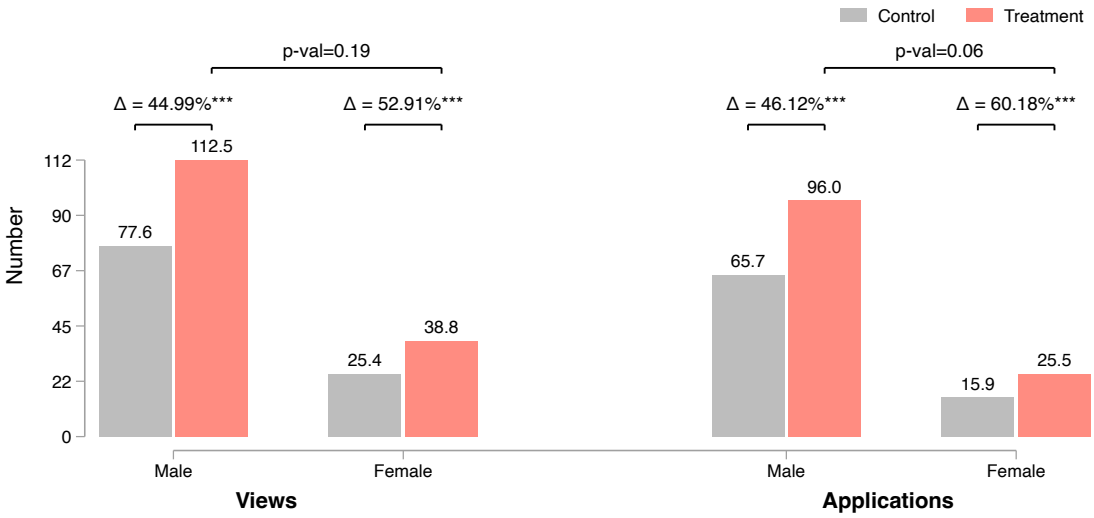
Notes: The figure presents initial previews of job postings on the platform that are visible to workers when they browse for jobs. Panel (a) illustrates a job posting where salary information is hidden, while Panel (b) shows a job posting with visible salary details. Specifically, it shows that the salary range is PKR 35,000 to 50,000. Both job postings are from the same firm and feature the job title, firm name and location, limited description of the job characteristics, application deadline, as well as experience and skills requirements. The placement of each of these details is the same for each job. Workers can click on these vignettes to see full job details which usually span 1-2 pages. [Return to page 24]

Figure A.18: Treatment effect on views and applications



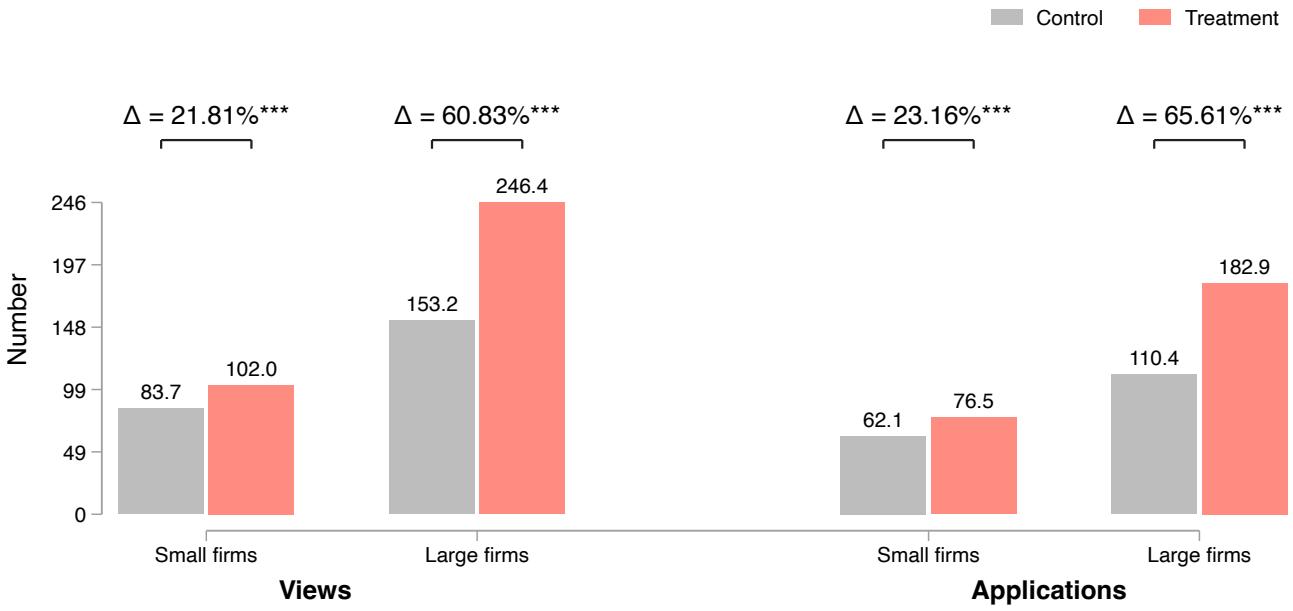
Notes: The figure shows treatment effect on job ad views and applications. The bars represent control and treatment group means, while the brackets represent percentage change from a poisson regression. [Return to page 52]

Figure A.19: Treatment effects by gender on views and applications



Notes: The figure shows treatment effect on job ad views and applications. The bars represent control and treatment group means, while the brackets represent percentage change from a poisson regression. [Return to page 52]

Figure A.20: Treatment effect on views and applications by firm size



Notes: The figure shows treatment effect on job ad views and applications with heterogeneity by firm size. A firm is defined as large if it has more than 50 employees. The bars represent control and treatment group means, while the brackets represent percentage change from a poisson regression. [Return to page 53]

A.2 Appendix Tables

Table B.1: Descriptive Statistics: Jobs

	Mean	SD
Salary		
Min Salary - Visible (PKR)	42,517	43,895
Max Salary - Visible (PKR)	68,183	67,561
Min Salary - Hidden (PKR)	56,152	61,529
Max Salary - Hidden (PKR)	98,143	107,424
Years of Experience Required	0.90	1.89
N. Applications Received	75.45	270.02
N. Requred Skills	3.25	2.44
N. Desired Skills	0.61	1.45
Full-time Job	0.94	
Schedule: Day time job	0.79	
Screening Question Required	0.05	
Career Level		
Intern/Student	0.04	
Entry Level	0.29	
Mid-Career/Professional	0.65	
Dept. Head	0.01	
Executive/Top Management	0.00	
Career Level Not Reported	0.01	
Education Requirement		
Matriculation or Less	0.07	
Intermediate	0.16	
Certificate/Diploma	0.03	
Bachelor	0.65	
Master	0.09	
PhD	0.00	
Education Requirement Not Reported	0.01	
Firm Size: N. Employees		
1-10	0.22	
11-50	0.32	
51-100	0.17	
101-200	0.10	
201-300	0.03	
301-600	0.05	
600+	0.12	
Observations	326,145	

Notes: The table shows baseline descriptive statistics of jobs covering the period January 2019-April 2024. Means and standard deviations are reported for averages in Columns 1 and 2, respectively. Standard deviations are omitted where the statistic represents share of jobs. [Return to page 24]

Table B.2: Top 10 Occupations and Industries on the Platform

Rank	Industry	Industry Percent	Occupation	Occupation Percent
1	Information Technology	22.69	Sales and Business Development	15.08
2	Services	7.18	Software and Web Development	12.26
3	N.G.O./Social Services	5.97	Accounts, Finance and Financial Services	8.89
4	Call Center	5.91	Client Services and Customer Support	7.69
5	Education/Training	5.54	Marketing	5.06
6	Manufacturing	5.21	Telemarketing	4.73
7	Real Estate/Property	4.40	Human Resources	4.09
8	Recruitment / Employment Firms	4.19	Creative Design	3.66
9	Business Development	3.66	Teachers/Education, Training and Development	3.61
10	BPO	3.28	Computer Networking	2.85

Notes: The table shows the top 10 industries and occupations on the platform. Column 1 ranks top 10 industries and occupations in order of most to least common. Column 2 lists the industries while column 3 lists the percentage of total jobs the industry corresponds to. Column 4 lists the occupations while column 3 lists the percentage of total jobs the occupation corresponds to. [Return to page 8]

Table B.3: Descriptive Statistics: Job-seekers

	Mean	SD
Age	28.13	6.89
Years of Experience	3.27	5.01
GPA	2.85	0.80
N. Applications Sent	8.84	46.14
Share Female	0.27	
Current Salary (PKR)	45,998	52,795
Expected Salary (PKR)	49,441	54,151
Career Level		
Intern/Student	0.26	
Entry Level	0.29	
Mid-Career/Professional	0.41	
Dept. Head	0.03	
Executive/Top Management	0.01	
Education Level		
High School or Less	0.02	
Intermediate	0.11	
Certificate/Diploma	0.07	
Bachelor	0.53	
Master	0.26	
PhD	0.01	
Observations	3,158,049	

Notes: The table shows baseline descriptive statistics of job-seekers covering the period January 2019-December 2024. Means and standard deviations are reported for averages in Columns 1 and 2, respectively. Standard deviations are omitted where the statistic represents share of job-seekers. [Return to page 9]

Table B.4: Female-friendly amenities and firm size - baseline data

	(1) Flex work	(2) Remote work	(3) Transport
Large firm	-0.26***	-0.08***	-0.07***
	(0.01)	(0.01)	(0.01)
Small firm	0.45	0.14	0.18
Observations	6,746	6,998	6,880

Notes: The sample consists of job advertisements for which the platform successfully collected survey data at baseline on a number of non-wage amenities. Column 1 shows whether the job ad offered flexible work hours, Column 2 remote work, and Column 3 transport support for commuting. Each column reports the association between firm size and the availability of these amenities, controlling for industry and occupation fixed effects, education and experience requirements, stated gender preferences, and the career level of the job. Large firms are defined as firms with more than 50 employees. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 23]

Table B.5: First stage by firm size

	Visible Salary	
	(1)	(2)
Treatment	0.401***	0.358***
	(0.006)	(0.007)
Large Firm		-0.192***
		(0.011)
Large Firm × Treated Job		0.112***
		(0.011)
Constant	0.562	0.640
Non-compliance rate in treatment	0.037	
Non-compliance rate in small firms		0.003
Non-compliance rate in large firms		0.083
Observations	20,088	20,088

Notes: This table presents first-stage estimates of salary visibility by firm size. **Column (1)** estimates the impact of treatment on salary visibility, while **Column (2)** accounts for firm size by including an interaction between treatment and large firms. Non-compliance rates indicate the share of firms that were assigned to treatment but did not comply by making salary information visible. Robust standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. [Return to page 25]

Table B.6: Treatment impacts by gender and firm size (Corresponding graph: Figure A.14)

Panel A: *Treatment effect as Incidence Rate Ratios (IRR) using Poisson*

	IRR: Views		**IRR: Applications**	
	(1) Male (m)	(2) Female (f)	(3) Male (m)	(4) Female (f)
Treated Job (β_1)	1.27***	1.17***	1.25***	1.15*
	(0.04)	(0.06)	(0.07)	(0.09)
Large Firm (β_2)	1.82***	1.72***	1.82***	1.62***
	(0.08)	(0.14)	(0.11)	(0.17)
Large Firm × Treated Job (β_3)	1.22*	1.53***	1.27**	1.69***
	(0.14)	(0.25)	(0.14)	(0.30)
Treatment effect on large firms ($\beta_1 \times \beta_3$)	1.55***	1.79***	1.59***	1.95***
	(0.18)	(0.27)	(0.15)	(0.31)
Treatment reallocation from small to large firms ($\beta_2 \times \beta_3$)	2.21***	2.64***	2.31***	2.74***
	(0.24)	(0.37)	(0.21)	(0.39)
H_0: $\beta_{2,f} = \beta_{2,m}$; p-value		0.39		0.05
H_0: $(\beta_{1,f} \times \beta_{3,f}) = (\beta_{1,m} \times \beta_{3,m})$; p-value		0.01		0.00
H_0: $(\beta_{2,f} \times \beta_{3,f}) = (\beta_{2,m} \times \beta_{3,m})$; p-value		0.00		0.00

Panel B: *Treatment effect as counts using OLS*

	N. Views		**N. Applications**	
	(1) Male (m)	(2) Female (f)	(3) Male (m)	(4) Female (f)
Treated Job (β_1)	15.99***	3.26***	12.42***	1.95*
	(1.85)	(1.12)	(2.69)	(1.01)
Large Firm (β_2)	47.60***	14.23***	40.48***	7.86***
	(4.17)	(2.28)	(4.37)	(1.87)
Large Firm × Treated Job (β_3)	42.50**	23.39***	40.37***	17.67***
	(18.28)	(8.74)	(13.33)	(5.87)
Treatment effect on large firms ($\beta_1 + \beta_3$)	58.49***	26.65***	52.79***	19.63***
	(18.18)	(8.67)	(13.06)	(5.79)
Treatment reallocation from small to large firms ($\beta_2 + \beta_3$)	90.10***	37.62***	80.85***	25.54***
	(17.79)	(8.43)	(12.60)	(5.57)
Constant	77.57	25.39	65.69	15.93
Observations	20,088	20,088	20,088	20,088

Notes: This table presents the treatment effects on job search behavior, focusing on gender differences in how workers direct search toward small and large firms. **Panel A** reports incidence rate ratios (IRR) estimated using a Poisson model, capturing the proportional change in job views and applications due to treatment. Columns (1)-(2) show the effects on male and female job views, while Columns (3)-(4) report similar estimates for job applications. **Panel B** presents the same analysis using OLS to estimate the treatment effects in levels, showing the absolute increase in job views and applications due to treatment. The corresponding **graphical representation** of these results is provided in Figure A.18 for views, and Figure A.14 for applications. Robust standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. [Return to page 25 or 81]

Table B.7: Preference for remote work by salary disclosure in large firms - Conditional Logit

Panel A: Women

	On-site	Remote	Difference
No salary	0.574***	0.723***	0.149**
	(0.022)	(0.062)	
Low salary	0.827***	0.785***	-0.041
	(0.051)	(0.067)	
Medium salary	0.820***	0.831***	0.011
	(0.052)	(0.060)	
High salary	0.810***	0.774***	-0.036
	(0.057)	(0.064)	
Observations			510

Panel B: Men

	On-site	Remote	Difference
No salary	0.562***	0.515***	-0.047
	(0.028)	(0.083)	
Low salary	0.664***	0.593***	-0.072
	(0.084)	(0.114)	
Medium salary	0.659***	0.672***	0.013
	(0.077)	(0.091)	
High salary	0.831***	0.720***	-0.111
	(0.060)	(0.089)	
Observations			412

Notes: Each table reports estimated preferences for remote versus on-site jobs within large firms, separately by gender and salary disclosure level. Panel A includes women; Panel B includes men. Each row corresponds to one of four salary levels: no salary disclosed, low, medium, or high salary. The columns show the predicted probability of choosing an on-site or remote job, as estimated from a conditional logit model, which includes respondent-choice-set fixed effects, leveraging within choice variation. As a result, choice sets are omitted if the interaction between remote work and salary level does not vary within the pair (i.e., choices are included when one job is remote and the other is on-site within the same salary condition). Thus, the number of observations here are lower than in the linear probability model reported in Table 8. The final column reports the difference in remote vs. on-site preferences at each salary level, along with significance stars. Predictions are generated from models that include interactions between salary level and a remote work indicator, controlling for other randomly assigned job attributes including diversity language and transport support. Coefficients represent the average marginal effect of each attribute on job choice. Standard errors are clustered at the choice set level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.8: Selection in applications to small firms by high and low ability candidates

Panel A: High Ability Men

	(1) **Index Top Decile**	(2) **Education Top Decile**	(3) **GPA Top Quartile**	(4) **Years of Experience Top Decile**	(5) **Managed a Team**
Treated Job (β_m)	0.01	0.00	0.02***	0.02***	0.01*
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Control mean	0.68	0.91	0.87	0.75	0.90

Panel B: High Ability Women

	(1)	(2)	(3)	(4)	(5)
Treated Job (β_f)	0.03**	0.04***	0.04***	0.04***	0.05***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Control mean	0.48	0.65	0.49	0.44	0.50
Observations	8,341	11,747	11,747	11,747	11,747
H_0: $\beta_f = \beta_m$; p-value	0.12	0.00	0.08	0.07	0.00

Panel C: Low Ability Men

	(1) **Index Bottom Decile**	(2) **Education Bottom Decile**	(3) **GPA Bottom Quartile**	(4) **Years of Experience Bottom Decile**	(5) **Never managed a Team**
Treated Job (β_m)	0.01	0.00	0.01***	0.01***	0.00
	(0.01)	(0.00)	(0.01)	(0.00)	(0.00)
Control mean	0.68	0.99	0.91	0.95	0.99

Panel D: Low Ability Women

	(1)	(2)	(3)	(4)	(5)
Treated Job (β_f)	0.03**	0.02**	0.05***	0.03***	0.02***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Control mean	0.48	0.85	0.56	0.75	0.84
Observations	8,341	11,747	11,747	11,747	11,747
H_0: $\beta_f = \beta_m$; p-value	0.12	0.04	0.00	0.01	0.00

Notes: The sample consists of job ads posted by small firms. Column 1 of Panels A and B shows likelihood men and women in the top decile of their expected wage applied. Panels C and D show the same for men and women in the bottom decile of expected wage. Column 2 of Panels A and B shows likelihood men and women in the top quartile of GPA applied, while in Panels C and D, the column corresponds to the bottom quartile. Column 3 of Panels A and B shows likelihood men and women in the top decile of years of experience applied, while Panels C and D, the show the same for the bottom decile. Column 4 of Panels A and B shows likelihood men and women who have any management experience applied, while Panels C and D, the show the same for men and women with no management experience. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 37]

Table B.9: Gender gaps in search among high ability job-seekers

	(1) **Expected Salary Top Decile**	(2) **GPA Top Quartile**	(3) **Years of Experience Top Decile**	(4) **Managed a Team**
Female Applications	-3.91***	-4.29***	-4.77***	-4.84***
	(0.25)	(0.19)	(0.26)	(0.17)
Male Applications	13.70	12.30	15.99	12.59
Observations	89,715	108,603	97,876	147,364

Notes: The table considers gender gaps in applications among high ability workers. It shows job-seeker level regressions of number of applications sent by high ability workers on the gender dummy. The row displaying "Male Applications" shows the number of applications sent by men. The row of coefficients pertaining to "Female Applications" shows the difference relative to men in the number of applications sent by female job-seekers. Column 1 shows these gender differences among job-seekers in the top decile of expected salary. Column 2 shows these gender differences among job-seekers in the top quartile of GPA. Column 3 shows these gender differences among job-seekers in the top decile of years of experience. Column 4 of shows these gender differences among job-seekers with management experience. An observation is missing if that characteristic is not observed for a worker since these are not mandatory to report. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page **??**]

Table B.10: Spillovers test: Salary disclosure

	N. Female Apps		**N. Male Apps**		**Salary visible**
	(1)	(2)	(3)	(4)	(5)
High saturation cluster	1.31	-0.93	4.06	-0.38	-0.00
	(1.07)	(1.06)	(3.04)	(2.90)	(0.02)
Large firm		1.72		23.59***	
		(1.13)		(3.31)	
Constant	12.82	12.82	54.81	54.81	0.71
Observations	8,462	8,060	8,462	8,060	8,462
Cluster FE	✓		✓		✓
Drop 3 high traffic clusters in treatment		✓		✓	

Notes: The table tests for spillover effects from neighboring jobs' treatment status on untreated (control) jobs. Each observation is a control job. Columns 1–4 report effects on the number of female and male applications; Column 5 reports effects on whether the job disclosed a salary. The key independent variable is a binary indicator for whether the job is located in a "high saturation cluster," defined as a cluster where treatment probability is 75% (versus 25% in low saturation clusters). Columns 1, 2 and 5 include cluster fixed effects and Columns 2 and 4 exclude the three highest-traffic clusters. "Large firm" is included as a control in Columns 3 and 4. Standard errors are clustered at the cluster level. $*p < 0.1, **p < 0.05, ***p < 0.01$. [Return to page 26].

Table B.11: Spillovers test: Conditioning on neighbors' exposure

	N. Female Apps			N. Male Apps		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated Job	1.95*	-3.69***	-1.11	12.42***	-6.27**	-1.66
	(1.01)	(1.18)	(1.10)	(2.69)	(2.98)	(2.62)
Large firm	7.86***	7.20***	-3.11	40.48***	38.28***	0.71
	(1.87)	(1.82)	(2.37)	(4.37)	(4.26)	(5.33)
Treated Job × Large firm	17.67***	17.85***	13.05**	40.37***	40.96***	30.80**
	(5.87)	(5.88)	(6.48)	(13.33)	(13.35)	(14.60)
Constant	15.93	15.93	15.93	65.69	65.69	65.69
Observations	20,088	20,088	19,533	20,088	20,088	19,533
Control: high saturation cluster		✓			✓	
Cluster FE			✓			✓

Notes: The table re-estimates the treatment effects of salary transparency on application counts conditional on whether the job is located in a high-saturation cluster. The sample includes both treated and control jobs. Columns 1–3 report treatment effects on female applications; Columns 4–6 on male applications. The second specification (Columns 2 and 5) control for whether the job is in a high-saturation cluster, defined as a cluster where treatment probability is 75% (versus 25% in low saturation clusters). The third specification (Columns 3 and 6) include cluster fixed effects. Standard errors are clustered at the cluster level. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 26].

Table B.12: Accuracy of worker beliefs about hidden salary

	Guess within 10% of min salary		**Guess within 10% of max salary**	
	(1) Small firm	(2) Large firm	(3) Small firm	(4) Large firm
Female	0.018	0.028	0.039	-0.044
	(0.023)	(0.028)	(0.024)	(0.030)
Male Mean	0.069	0.088	0.065	0.137
Observations	531	464	531	464

Notes: This table reports accuracy in guessing hidden salaries by gender. Respondents were asked to guess the salary range for two job ads (one from a small firm and one from a large firm) that did not disclose pay, drawn from the respondent's own occupation. They received a financial reward if their guess fell within 10% of the true salary reported internally by the firm to the platform. Columns (1)–(2) report accuracy relative to the minimum of the salary range; Columns (3)–(4) relative to the maximum. Male means are reported in the row labeled "Male Mean." $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page **??** or 31]

Table B.13: Accuracy of worker beliefs about competition

	Guess within 10% of female apps		**Guess within 10% of male apps**	
	(1) Small firm	(2) Large firm	(3) Small firm	(4) Large firm
Female	-0.010	-0.005	0.003	-0.009
	(0.018)	(0.018)	(0.017)	(0.018)
Male Mean	0.046	0.040	0.037	0.040
Observations	497	456	497	456

Notes: This table reports accuracy in guessing the number of male and female applications by gender of the respondent. Job seekers were shown two job ads (one from a small firm and one from a large firm). Respondents were then asked to guess how many male and female applicants each ad received, and received a financial reward if their guess fell within 10% of the true number of applications recorded on the platform. Columns (1)–(2) report accuracy relative to the number of female applications; Columns (3)–(4) relative to the number of male applications. Standard errors are in parentheses. Male means are reported in the row labeled "Male Mean." $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 32]

Table B.14: Accuracy of worker beliefs about firm size

	Guess firm size correctly:	
	(1) Small firm	(2) Large firm
Female	0.042	-0.017
	(0.047)	(0.045)
Male Mean	0.356	0.728
Observations	440	402

Notes: This table reports the probability that jobseekers correctly classify a firm as small ($\leq$50 employees) or large ($>$50 employees) based on the job ad. Respondents were asked to guess the number of employees at the firm, and a response is coded as correct if it places the firm on the correct side of the large-firm threshold, using administrative data as a benchmark. Columns (1)–(2) show results for ads from small and large firms, respectively. Standard errors are in parentheses. Male means are reported in the row labeled "Male Mean." $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 16 or 32]

Table B.15: Accuracy of worker beliefs about firm size

	(1) Network size	(2) Share of network that provides job advice	(3) Share of network employed	(4) Share of employed network that earns more	(5) Share of workers' apps to treated jobs
Female	-0.050	-0.028	-0.088***	-0.012	0.073
	(0.647)	(0.031)	(0.029)	(0.032)	(0.099)
Female × Share of network employed					-0.109
					(0.123)
Share of network employed					0.038
					(0.096)
Male Mean	3.945	0.720	0.786	0.615	0.682
Observations	422	421	421	397	421

Notes: This table reports gender differences in self-reported network characteristics. Respondents were asked to list the number of close friends or relatives with whom they interact at least every two weeks (Column 1). They then reported how many of these contacts provide job search advice (Column 2), are employed (Column 3), and, conditional on being employed, earn more than the respondent (Column 4). Responses are normalized into shares of the respondent's reported network size, except for Column (1), which is the raw count. Male means are reported in the row labeled "Male Mean." $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 32]

Table B.16: Beliefs about discrimination by large firms

	Own salary offer		Salary offer to others
	(1) Within 10% of max	(2) Within 20% of max	(3) Believe men paid more
Female	-0.038	-0.054	0.054
	(0.044)	(0.047)	(0.041)
Male Mean	0.318	0.576	0.224
Observations	451	451	444

Notes: This table reports gender differences in beliefs about salaries and discrimination at large firms. After viewing a large-firm job ad, respondents were first asked to guess the minimum and maximum salary for the job, and then to state the monthly salary they thought they themselves would be offered. Columns (1)–(2) report whether the respondent's expected own offer fell within 10% or 20% of their previously stated maximum, respectively. Column (3) elicits second-order beliefs: respondents were shown two otherwise identical candidate profiles, one male and one female, and asked how much each candidate would be offered. Column (3) is coded as one if their answers implied they expect a positive gender pay gap between male and female salary offers. Male means are reported in the row labeled "Male Mean." $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 33]

Table B.17: Need to seek permission for job search by age and gender

	All Ages (%)		Below Median (%)		Above Median (%)	
Category	**Male**	**Female**	**Male**	**Female**	**Male**	**Female**
No one	50.8	31.5	45.3	31.1	55.8	32.0
Father	26.5	32.8	31.4	37.9	22.1	26.2
Mother	7.7	13.6	9.3	16.7	6.3	9.7
Spouse	4.4	12.3	3.5	4.5	5.3	22.3
Others	10.5	9.8	10.5	9.8	10.5	9.7
Observations	181	235	86	132	95	103
Male median age: 28; Female median age: 26						

Notes: This table reports the share of respondents in the job-seekers' survey who say they need to seek permission from household members before applying for jobs, by gender and by age relative to the sample median. Categories reflect the household member from whom permission is sought: father, mother, spouse, or others. "No one" indicates that the respondent does not seek permission from anyone. Male and female median ages are 28 and 26, respectively. [Return to page 34]

Table B.18: Intra-household bargaining: Seeking permission for applications

	(1) Seeks permission before applying	(2) Share of workers' apps to treated jobs
Female	0.193***	-0.006
	(0.048)	(0.055)
Female × Seeks permission before applying		-0.004
		(0.072)
Seeks permission before applying		0.021
		(0.053)
Male Mean	0.497	0.497
Observations	422	422

Notes: This table reports gender differences in permission-seeking behavior and whether such behavior predicts applications to treated jobs (jobs with disclosed salaries). Column (1) shows the probability that respondents report seeking permission from a household member before applying. Column (2) reports the share of each respondent's applications that went to treated jobs. The interaction term tests whether permission-seeking predicts differential application to treated jobs by gender. Male means are reported in the row labeled "Male Mean." $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 34]

Table B.19: Ambiguity, risk and bargaining preferences

	Preferences			**Share apps to treated jobs**	
	(1) Ambiguity averse	(2) High risk preference	(3) Bargain averse	(4) By risk	(5) By bargaining
Female	0.034	-0.632**	0.182***	-0.013	-0.083
	(0.048)	(0.265)	(0.049)	(0.037)	(0.054)
Female × Risk preference				0.024	
				(0.076)	
Risk preference				0.001	
				(0.007)	
The firm with the fixed salary offer=1 × Female					0.122*
					(0.073)
Bargain averse					-0.011
					(0.054)
Male Mean	0.587	5.940	0.434	0.677	0.677
Observations	424	419	419	419	419

Notes: This table reports gender differences in ambiguity, risk, and bargaining preferences, and examines whether these preferences predict application behavior when salaries are disclosed. Column (1) reports whether respondents are ambiguity averse in an incentivized urn game following Ashraf et al. (2009). Column (2) reports self-assessed risk preferences on a 1–10 scale, with higher values indicating greater willingness to take risks. Column (3) reports whether respondents prefer a non-negotiable job offer over a negotiable one, coded as bargaining aversion. Columns (4)–(5) report whether risk preference and bargaining aversion predict the share of applications to treated jobs (with salary disclosure), including their interactions with gender. Male means are reported in the row labeled "Male Mean." $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 35]

Table B.20: Firm Survey: Hiring outcomes

	N. Women			N. Men		
	(1) All firms	(2) Large firms	(3) Small firms	(4) All firms	(5) Large firms	(6) Small firms
Treated Job	-0.01	-0.01	-0.00	0.05	-0.04	0.08
	(0.03)	(0.03)	(0.04)	(0.06)	(0.09)	(0.07)
Control Job Mean	0.49	0.32	0.55	1.36	1.13	1.44
Observations	5,685	1,471	4,214	5,685	1,471	4,214

Notes: This table assesses differences in hiring outcomes for both genders by treatment status. Each observation represents a unique job included in the endline firm survey. The survey was administered after the experiment to 4,146 firms (46.6% of experimental firms), covering a random subset of up to five jobs per firm. In total, the survey includes responses for 5,685 jobs (28.3% of all experimental jobs). Respondents were firm representatives registered on the platform and familiar with hiring decisions. The survey asked only about the number of total hires, number of women hired, and time to fill the vacancy to limit respondent burden. Treatment status does not predict survey response. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 41].

B Theoretical framework: Why large firms hide attractive salaries

Advertising high salaries plausibly draws the best talent to a job. Why then do large firms hide their appealing salaries from job ads? To understand this, I develop a job search model that highlights the potential costs to firms of advertising high salaries. The model considers how a higher advertised wage changes the composition of the applicant pool, drawing on insights from Ashraf et al. (2020). Later, I extend it to consider why hiding these higher wages may be rational for a firm.

Model setup. Consider a one-sector, static world where a firm is seeking to recruit a worker for a job. Workers have job-specific ability θ, drawn from a common distribution with CDF $F(\theta)$ and PDF $f(\theta)$ that is known to firms and workers. Note that the same worker may have different θ for different types of jobs, but here job type is held fixed. Workers' ability in a job is thus considered by firms to be a measure of their 'fit' for the job. However, workers' θ is their private information initially, and it is revealed to firms when they invite workers for an interview.

Firms make wage announcements and receive applications from workers. In turn, they face a cost $C(n)$ for advertising a vacancy and processing applications, which is increasing in n, the expected number of applications. Firms post a wage w, and workers earn this wage if hired, regardless of their θ. The match produces surplus θ if formed, and firms get a pay-off $\theta - w$.

Hiring a worker with $\theta < \theta_{\min}$ results in a negative expected payoff. Thus, a firm only hires an applicant if their ability (θ) meets or exceeds a threshold $\theta_{\min}$. The firm selects the best applicant from its pool of n applicants, where the best applicant's ability is $\bar{\theta} \equiv \max\{\theta_1, \ldots, \theta_n\}$. Thus, the firm's hiring probability is the probability that at least one applicant meets the threshold, which is equivalent to the probability that the best applicant satisfies $\bar{\theta} \geq \theta_{\min}$. The probability that a single applicant has $\theta < \theta_{\min}$ is $Pr(\theta < \theta_{\min}) = F(\theta_{\min})$. Since applicants are independent, the probability that all n applicants have $\theta < \theta_{\min}$ is $Pr(\bar{\theta} < \theta_{\min}) = [F(\theta_{\min})]^n$. Therefore, the probability that at least one applicant meets the hiring threshold-and hence that the firm successfully hires the best worker-is:

$$Q(n, \theta_{min}) \equiv Pr(\bar{\theta} \geq \theta_{\min}) = 1 - [F(\theta_{\min})]^n$$

This matching probability is increasing in n and $\bar{\theta}$ as shown in Appendix Section B.1.

Workers' matching probability contains two components: the probability that no other worker with higher ability applied, and the probability that worker's own ability exceeds the firm's cut-off ($\theta \geq \theta_{min}$). The

probability that any given worker has ability at most θ is $F(\theta)$. Since there are $n-1$ other applicants, the probability that all of them have $\theta' < \theta$ is $[F(\theta)]^{n-1}$. The probability that the worker's own ability meets the hiring threshold is $1 - F(\theta_{min})$. The workers' matching probability is then:

$$P(n, \theta) \equiv [F(\theta)^{n-1}] \cdot (1 - F(\theta_{min}))$$

which decreases with n and increases with θ as shown in Appendix Section B.2.

Both firms and workers face uncertainty over matching. Firms also face uncertainty over workers' θ. As a result, both sides make entry decisions based on rational forecasts over $n(w)$, the number of applications a job will draw for each w (increasing in w), and $\bar{\theta}(w)$, the best applicant's expected ability in the applicant pool. The latter also increases in w, as higher wages compensate workers for their outside options.

Firms. The firm chooses w to maximize profits:

$$\pi(w) = [1 - F(\theta_{\min})^{n(w)}] \left(\bar{\theta}(w) - w\right) - c(n(w)) \tag{6}$$

where $\theta_{min} = \frac{c(n(w))+w}{1-F(\theta_{\min})^{n(w)}}$ ensures non-negative profits.

Workers. Workers apply if their expected utility from applying is weakly positive:

$$U_\theta = F(\theta)^{n-1}[1 - F(\theta_{min})] \left(\lambda\theta_i + w\right) - V(\theta) - s \geq 0 \tag{7}$$

where workers receive a wage w if employed. They have an outside option $V(\theta)$, increasing in θ, and pay a search cost s. Workers derive non-pecuniary enjoyment from working, which depends on how well-matched they are with the job. Specifically, they get $\lambda\theta$ from working, where $0 \leq \lambda < 1$.

Matching. Matching occurs in three stages. In Stage 1, firm choose w to maximize profits. In Stage 2, workers make application decisions given w. In the third stage, not modeled explicitly but anticipated in stages 1 and 2, firms pick the best worker.

Workers apply to a job with wage w if their θ falls in a given range. There is a threshold $\underline{\theta}(w)$ below which workers do not apply because their expected utility from working relative to the search cost is too low. There is an upper threshold $\bar{\theta}(w)$ beyond which workers with $\theta \geq \bar{\theta}(w)$ do not apply because the expected payoff is below their outside option. These thresholds can be seen in Panel (a) of Figure A.21 .

Impact of advertising a high wage. Higher wages increase $\bar{\theta}(w)$ by compensating high-ability workers for their higher outside option. But, as shown in Panel (b) of Figure A.21, they have two opposite impacts on $\underline{\theta}(w)$, making the net result ambiguous. First, they decrease $\underline{\theta}(w)$ by compensating low-ability workers for their lower enjoyment of work relative to the application cost, inducing workers with lower ability to apply. Second, they increase total demand for the job, reducing the likelihood that any individual worker gets hired, thereby discouraging applications (proof in Appendix Section B.3):

$$\frac{\partial\big[P(n(w),\theta)\big]}{\partial w} \;<\; 0,$$

But this effect is stronger for low-ability workers because higher ability workers are now selecting into the applicant pool, reducing the likelihood that low ability workers get hired (proof in Appendix Section B.4):

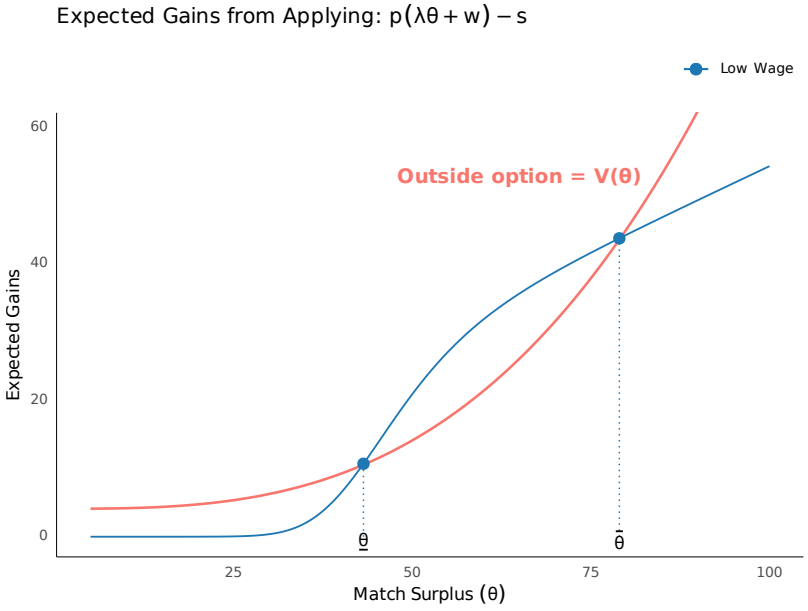
$$\frac{\partial^2\big[P(n(w),\theta\big]}{\partial w\,\partial\theta} \;>\; 0,$$

As a result, higher wages may increase $\underline{\theta}(w)$ because lower ability workers now experience tougher competition.

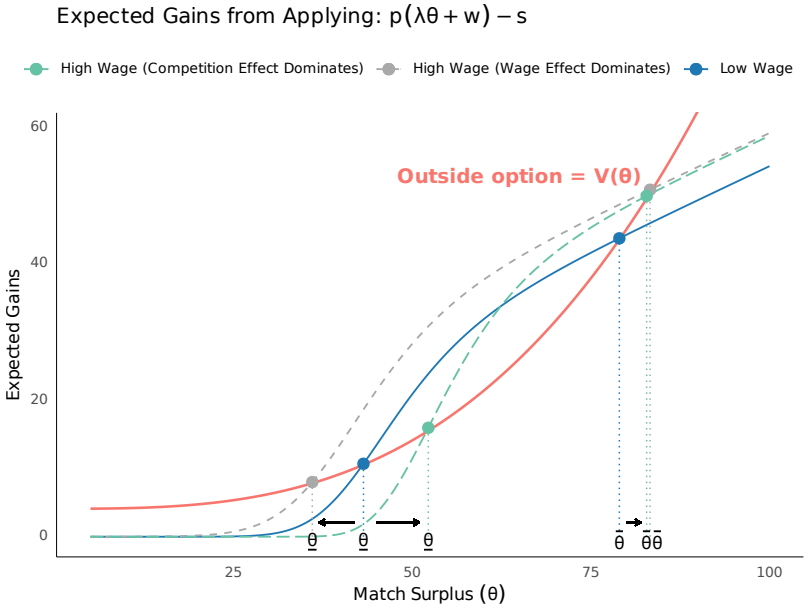
Result: Upper threshold is increasing in w (i.e., $\frac{\partial\bar{\theta}(w)}{\partial w} > 0$). But lower threshold may be increasing or decreasing in w. For details see Sections B.5 and B.6.

Figure A.21: Theoretical framework: Application thresholds

(a) Selection into Application



(b) Increase in advertised wage



Notes: This figure illustrates the relationship between worker ability θ, the expected gains from applying to a job (solid blue line), and the outside option $V(\theta)$ (solid red line). The intersections of these two curves define the lower and upper application thresholds, denoted as $\underline{\theta}$ and $\bar{\theta}$, respectively. The expected gains from applying are determined by the hiring probability and the expected wage, adjusted for search costs. Specifically, it is: $P(n, \theta) \cdot (\lambda\theta + w) - s$. Panel (b) shows two counterfactuals when a higher wage (w) is advertised. In one case, the competition effect dominates, and in the other, the wage effect dominates. [Return to page 104]

How salary non-disclosure can self-screen. If firms do not post wages, workers form expectations over wages: $\tilde{w} = w + \epsilon$ where $\epsilon \sim N(0, \sigma^2)$. Accessing wage information when it is hidden is more costly for workers with lower θ because they lack the experience or skills in the relevant role that would enable them to extract wage signals from job posts, or are less well-connected to relevant networks. Alternatively, they have to spend more time researching a firm and its specific job. As a result, application cost is heterogeneous under this regime, with $s_\theta = s + \frac{1}{\theta_i}$ where the cost rises as the fit with the job declines. Workers apply if:

$$F(\theta)^{n-1}[1 - F(\theta_{min})]\left(\lambda\theta_i + \tilde{w}\right) - V_\theta - s - \frac{1}{\theta_i} \leq 0.$$

Wage hiding is, thus, a rational response by firms who believe that as θ falls, the cost of application rises prohibitively, deterring lower ability workers from applying and raising $\underline{\theta}$. This means that the process of learning wages when they are hidden is a screening ordeal that disproportionately affects workers who are worse fits, leading them to self-screen out of the applicant pool while the right type of worker optimally chooses to apply. By contrast, a law that mandates salary range disclosure in job ads eliminates $\frac{1}{\theta}$, all workers of all types equal access to the job.

Prediction. As $\theta \to 0$, the cost of application grows, deterring lower-ability workers, and raising the threshold $\underline{\theta}(w)$. Thus, hiding wages may reduce the total number of applicants, while the quality of the applicant pool increases.

B.1 Impact of n and $\bar{\theta}$ on firm's hiring probability

The impact of n on the hiring probability can be seen by differentiating it with respect to n:

$$\frac{d}{dn}\left(1 - [F(\theta_{\min})]^n\right) = -\ln F(\theta_{\min}) \cdot [F(\theta_{\min})]^n.$$

Since $0 \leq F(\theta_{\min}) < 1$, it follows that $\ln F(\theta_{\min}) < 0$, ensuring:

$$\frac{d}{dn}Q(n, \bar{\theta}) > 0$$

The second effect comes from the fact that if the expected maximum ability $\bar{\theta}$ increases, the probability of exceeding any fixed threshold $\theta_{\min}$ also rises. Thus:

$$\frac{d}{d\bar{\theta}}Q(n, \bar{\theta}) > 0$$

B.2 Impact of n and θ on workers' hiring probability

Workers' hiring probability is:

$$P(n, \theta) \equiv [F(\theta)]^{n-1} \cdot (1 - F(\theta_{min}))$$

Differentiating with respect to n:

$$\frac{\partial}{\partial n} P(n, \theta) = (1 - F(\theta_{\min})) \cdot [F(\theta)]^{n-1} \ln F(\theta).$$

Since $0 \leq F(\theta) \leq 1$, we have $\ln F(\theta) \leq 0$, which implies:

$$\frac{\partial}{\partial n} [F(\theta)]^{n-1} \leq 0.$$

Thus, as n increases, $[F(\theta)]^{n-1}$ decreases, and since $(1 - F(\theta_{\min}))$ is independent of n, we conclude:

$$\frac{\partial}{\partial n} P(n, \theta) < 0.$$

Differentiating with respect to θ:

$$\frac{\partial}{\partial \theta} P(n, \theta) = (n-1)[F(\theta)]^{n-2}(1 - F(\theta_{min}))f(\theta),$$

which is positive for $n > 1$.

B.3 Hiring probability is decreasing in w: $\frac{\partial P}{\partial w} < 0$

$$P(n(w), \theta) = [F(\theta)]^{n(w)-1} \cdot (1 - F(\theta_{\min})).$$

Since the number of applicants $n(w)$ is increasing in w, we differentiate $P(n(w), \theta)$ with respect to w:

$$\frac{\partial}{\partial w} P(n(w), \theta) = \frac{\partial}{\partial n} P(n, \theta) \cdot \frac{dn(w)}{dw}$$

where

$$\frac{\partial}{\partial n} P(n, \theta) = (1 - F(\theta_{\min})) \cdot [F(\theta)]^{n-1} \ln F(\theta),$$

and $\ln F(\theta) \leq 0$ since $0 \leq F(\theta) \leq 1$. This implies:

$$\frac{\partial}{\partial n} P(n, \theta) \leq 0.$$

Since $\frac{dn(w)}{dw} > 0$ (i.e., the number of applicants increases with w), we conclude that:

$$\frac{\partial}{\partial w} P(n(w), \theta) < 0.$$

Thus, as w increases, competition intensifies, reducing the probability that any individual worker matches.

B.4 Impact is smaller for workers with lower θ: $\frac{\partial^2 P}{\partial w \partial \theta} > 0$

$$\frac{\partial^2 P}{\partial w \partial \theta} = \frac{\partial}{\partial \theta} \left([F(\theta)]^{n(w)-1} \ln F(\theta) \cdot \frac{dn(w)}{dw} \right)$$

$$\frac{\partial^2 P}{\partial w \partial \theta} = \frac{dn(w)}{dw} \cdot \left(\frac{\partial}{\partial \theta} \left[[F(\theta)]^{n(w)-1} \ln F(\theta) \right] \right)$$

Now, differentiating the term inside the parentheses:

$$\frac{\partial}{\partial \theta} [F(\theta)]^{n(w)-1} \ln F(\theta) + [F(\theta)]^{n(w)-1} \frac{1}{F(\theta)} f(\theta)$$

where:

$$\frac{\partial}{\partial \theta} [F(\theta)]^{n(w)-1} = (n(w) - 1)[F(\theta)]^{n(w)-2} f(\theta),$$

we get:

$$(n(w) - 1)[F(\theta)]^{n(w)-2} f(\theta) \ln F(\theta) + [F(\theta)]^{n(w)-1} \frac{f(\theta)}{F(\theta)}.$$

Factoring out $f(\theta)$ and $[F(\theta)]^{n(w)-2}$:

$$f(\theta)[F(\theta)]^{n(w)-2} \left[(n(w) - 1) \ln F(\theta) + 1 \right]$$

Thus:

$$\frac{\partial^2 P}{\partial w \partial \theta} = f(\theta)[F(\theta)]^{n(w)-2} \left[(n(w) - 1) \ln F(\theta) + 1 \right] \frac{dn(w)}{dw}.$$

where:

- $f(\theta) > 0$, since it is a probability density function and must be non-negative for all θ.

- $\frac{dn(w)}{dw} > 0$, because an increase in wages attracts more applicants.
- For the expression to be positive, we need $(n(w) - 1) \ln F(\theta) + 1 > 0$.
  - Rearranging:

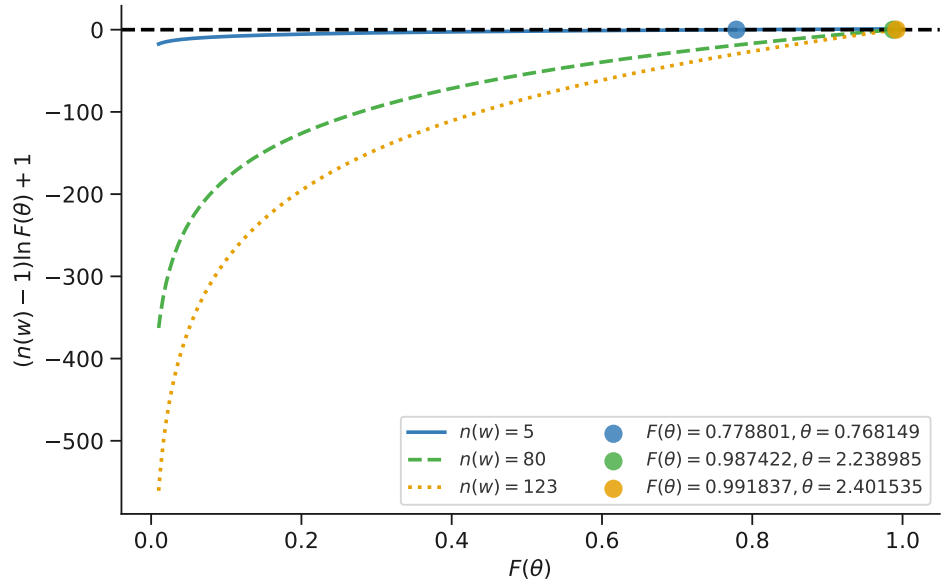
$$(n(w) - 1) \ln F(\theta) > -1$$

$$\ln F(\theta) > -\frac{1}{n(w) - 1}$$

$$F(\theta) > e^{-1/(n(w)-1)}$$

  - This condition is met for finite sufficiently large θ. For larger applicant pools, even moderately high θ values will satisfy this condition as shown in the graph below.

Figure A.22: Threshold values of $F(\theta)$ for which $(n(w) - 1) \ln F(\theta) + 1 > 0$.



Notes: This figure illustrates the relationship between the cumulative distribution function $F(\theta)$ and the condition $(n(w) - 1) \ln F(\theta) + 1$. Each line represents a different value of $n(w)$, the number of applications. Circular markers indicate the threshold values where the expression transitions from negative to non-negative. The corresponding θ values, derived from the inverse normal CDF. The legend is structured into two columns: the first listing $n(w)$ values corresponding to each curve, and the second listing threshold values and their associated θ.

Thus, as θ grows we have:

$$\frac{\partial^2 P}{\partial w \partial \theta} > 0.$$

B.5 Upper Threshold: $\bar{\theta}(w)$ increases with w.

Define the function

$$F\big(\theta, w\big) \;=\; P\big(n(w), \theta\big)\,\big[\lambda\,\theta + w\big] \;-\; V(\theta) \;-\; s.$$

The threshold $\bar{\theta}(w)$ is implicitly defined by the condition

$$F\big(\bar{\theta}(w),\, w\big) = 0.$$

Applying the Implicit Function Theorem,

$$\frac{\mathrm{d}}{\mathrm{d}w}\,F\big(\bar{\theta}(w),\, w\big) \;=\; F_\theta\big(\bar{\theta}(w), w\big)\,\bar{\theta}' \;+\; F_w\big(\bar{\theta}(w), w\big) \;=\; 0 \quad\Longrightarrow\quad \bar{\theta}(w)' \;=\; -\,\frac{F_w}{F_\theta}.$$

Where:

$$F_\theta = \underbrace{P\big(n(w), \bar{\theta}(w)\big)\,\lambda}_{\text{(i)}} \;-\; \underbrace{\frac{\partial V_{\bar{\theta}}}{\partial \theta}}_{\text{(ii)}},$$

$$F_w = \underbrace{\frac{\partial P}{\partial w}\Big[\lambda\,\bar{\theta}(w) + w\Big]}_{\text{(negative since } \frac{\partial P}{\partial w}<0)} + \underbrace{P\big(n(w), \bar{\theta}(w)\big)}_{(>0 \text{ because it's a probability})} - \underbrace{\frac{\partial V_{\bar{\theta}}}{\partial w}}_{(=0 \text{ since } V(\theta) \text{ depends only on } \theta)}$$

Sign conditions:

- $\boldsymbol{F_w > 0}$: A shown in Section B.4, while $\frac{\partial P}{\partial w} < 0$ since more competition lowers hiring probability, its magnitude falls as θ rises, thus the first term becomes smaller. Meanwhile, the hiring probability rises in θ as shown below:

$$P(n, \theta) = (n-1)[F(\theta)]^{n-1}(1 - F(\theta_{min}))$$
$$\frac{\partial P(n, \theta)}{\partial \theta} = (n-1)[F(\theta)]^{n-2}(1 - F(\theta_{min}))f(\theta)$$

  Thus, as θ rises, the magnitude of $P(.)$ dominates, making $F_w > 0$.

- $\boldsymbol{F_\theta < 0}$: As ability θ increases, the outside option $V(\theta)$ grows at rate $\frac{\partial V(\theta)}{\partial \theta}$. Since $P(\cdot)$ is capped at 1 and $\lambda < 1$, we typically have

$$\frac{\partial V(\theta)}{\partial \theta} > P(n(w), \bar{\theta}) \cdot \lambda,$$

  ensuring that $F_\theta < 0$.

Since $F_w > 0$ and $F_\theta < 0$, we conclude:

$$-\frac{F_w}{F_\theta} > 0 \quad \Longrightarrow \quad \bar{\theta}'(w) > 0.$$

Hence, the upper threshold $\bar{\theta}(w)$ increases with the wage.

B.6 Lower Threshold $\underline{\theta}(w)$ has an Ambiguous Relationship w.r.t. w

Similarly, define

$$G\big(\theta, w\big) = P\big(n(w), \underline{\theta}(w)\big)\big[\lambda\,\theta + w\big] - V(\theta) - s.$$

The lower threshold $\underline{\theta}(w)$ is implicitly defined by

$$G\big(\underline{\theta}(w), w\big) = 0.$$

Applying the Implicit Function Theorem,

$$\frac{\mathrm{d}}{\mathrm{d}w} G\big(\underline{\theta}(w),\, w\big) = G_\theta\big(\underline{\theta}(w), w\big)\,\underline{\theta}'(w) + G_w\big(\underline{\theta}(w), w\big) = 0 \quad \Longrightarrow \quad \underline{\theta}'(w) = -\frac{G_w}{G_\theta}.$$

Where:

$$G_\theta = \underbrace{P\big(n(w), \underline{\theta}(w)\big)\,\lambda}_{\text{(i)}} - \underbrace{\frac{\partial V(\theta)}{\partial \theta}}_{\text{(ii)}},$$

$$G_w = \underbrace{\frac{\partial P}{\partial w}\Big[\lambda\,\underline{\theta}(w) + w\Big]}_{\text{(negative if competition effect dominates)}} + \underbrace{P\big(n(w), \underline{\theta}(w)\big)}_{\text{(positive)}}$$

Sign Conditions:

- $G_\theta < 0$: As θ increases, $V(\theta)$ grows at rate $\frac{\partial V(\theta)}{\partial \theta}$. Since $P(\cdot)$ is capped at 1 and $\lambda < 1$,

  $$\frac{\partial V(\theta)}{\partial \theta} > P(n(w), \underline{\theta}(w)) \cdot \lambda$$

  ensures $G_\theta < 0$.

- G_w is ambiguous: The direct wage effect increases $\big[\lambda\,\theta + w\big]$, while competition lowers hiring probability ($\frac{\partial P}{\partial w} < 0$). If $\underline{\theta}(w)$ is low, crowding-out is strong, making $G_w < 0$. If the direct wage effect dominates, then $G_w > 0$.

Hence, we have that:

$$\underline{\theta}'(w) = -\frac{G_w}{G_\theta}$$

- If $G_w > 0$, then $\underline{\theta}'(w) < 0$: More low-$\theta$ workers apply.
- If $G_w < 0$, then $\underline{\theta}'(w) > 0$: The lower threshold rises.

Thus, the effect of w on $\underline{\theta}(w)$ is ambiguous. High wage increases the best applicant's expected ability. Finally, note that $\underline{\theta}(w)$ is defined as the expected maximum (or best) ability in the applicant pool. As w rises, the upper threshold $\bar{\theta}(w)$ strictly increases, expanding the supply of high-ability applicants. Even if $\underline{\theta}(w)$ moves, the distribution's upper support becomes strictly larger. Hence, under standard single-peaked or monotonic hazard assumptions on θ's distribution, the *max* of the drawn pool shifts right in expectation, implying $\underline{\theta}'(w) > 0$.

Conclusion. Higher wages unambiguously *raise* the upper threshold $\bar{\theta}(w)$ and thus *increase* the expected top ability $\underline{\theta}(w)$. They have *opposing* influences on the lower threshold $\underline{\theta}(w)$: A positive wage effect that encourages low-θ types to apply. A negative crowding-out effect that discourages them by lowering $P(n(w), \underline{\theta}(w))$ when $\underline{\theta}(w)$ is large. Thus, a large crowding-out effect can push $\underline{\theta}(w)$ *upward*, reducing the participation of marginal workers, while a small crowding-out effect can drive $\underline{\theta}(w)$ *downward*.