Screening Women Out?
Pay Transparency in Job Postings

Amen Jalal*†

July 22, 2026

Job Market Paper

Frequently updated. Click **HERE** for latest version.

Abstract: Up to half of the gender pay gap stems from women's sorting into low-wage firms. Do women prefer amenities to wages, or face barriers to search? I tackle this question using data on 29 million job applications from Pakistan's largest job search platform, combined with firm and worker surveys and a field experiment mandating pay transparency. I document that large, high-paying firms are more likely to omit salaries in job ads and less likely to offer flexibility – an amenity women value slightly more than men. When pay is disclosed, men and women respond similarly to wages. But when undisclosed, behaviors diverge: men search randomly, while women sort negatively on pay. A theoretical framework shows that pay non-disclosure amplifies small gender differences in amenity preferences into large gender gaps in applications. To test whether transparency closes these gaps, I randomize mandatory versus optional pay disclosure in 20,088 jobs across 8,906 firms on the platform. The experiment leaves large-firm pay and amenities unchanged. Yet women's applications to these firms increase 95%, and men's 59%, reversing the gender gap in directed search. This implies women do not prefer flexibility to wages. Rather, they turn to flexibility when they cannot access wages. Meanwhile, large firms most exposed to mandated transparency become 30% more likely to disclose pay post-experiment, suggesting they overestimated the costs of transparency.

*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, Alan Manning and Paula Onuchic. I also benefited from comments by Emily Breza, Zoë Cullen, Matthias Doepke, Asim Khwaja, Camille Landais, Gabriel Leite Mariante, Maitreesh Ghatak, Marta Guasch-Rusinol, Guy Michaels, Canishk Naik, Steve Pischke, Nina Roussille, Pol Simpson, Kartik Srivastava, Daniel Sturm and seminar participants at LSE, UC Berkeley, IZA, Stanford, Warwick, HEC Paris, Royal Holloway, Exeter, LUMS, EAYE, the Wages, Employment and Inequality Conference, 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.

Introduction

Despite the global rise in female education and employment, women remain concentrated in small, low-wage, and low-productivity firms. This sorting into low-wage firms accounts for 15-50% of the gender pay gap (Card et al., 2016; Morchio and Moser, 2021; Bassier and Gautham, 2025). Pakistan, the setting of this study, offers a striking illustration: since 1990, female employment has doubled and college attainment has quadrupled. Despite these gains, gender gaps in sorting have widened – by 43% in salaried work and 114% in large-firm employment. Meanwhile, the gender pay gap has remained stagnant at around 30%, among the highest worldwide.

Women's sorting patterns may reflect preferences for low-wage firms, or barriers to accessing high-wage ones. A prevailing view emphasizes preferences, arguing that women may accept lower wages in exchange for family-friendly amenities (Goldin, 2014; Sorkin, 2017). However, the magnitude of the pay gap far exceeds most estimates of women's willingness to pay for these amenities (Mas and Pallais, 2017), suggesting that preferences alone cannot account for gendered sorting.

This paper identifies a critical but overlooked piece of the sorting puzzle: most workers search for jobs without knowing salaries. Across high- and low-income countries including the U.S., Chile, China, Germany and Ethiopia, over 80% of job ads omit salary information. Thus, when women "purchase" amenities with foregone wages, they may not know the price they are paying. Existing studies abstract from this possibility. Lab-based studies elicit preferences for amenities under full information, but in doing so, erase the role information frictions play outside the lab. Matched employer-employee data only capture realized matches, while disentangling preferences from frictions also requires observing unrealized opportunities – the jobs women choose not to pursue.

I partner with Pakistan's largest online job search platform to capture workers' application choices *and* consideration sets, together with the information environment that shapes these decisions. I also observe salaries even when workers do not, because the platform requires firms to report them internally.[1] About 55% of job ads omit salaries on the platform. In this setting, I document new facts around sorting, preferences, and pay disclosure, using 5 years of data on over 29 million job applications to 62,600 firms, supplemented by firm and worker surveys. The workers' surveys capture preferences with a discrete choice experiment, and expectations with incentivized belief elicitation. These data inform a search model that shows how wage uncertainty magnifies small gender differences in amenity tastes into large gender gaps in applications. The model yields predictions on the impacts of pay transparency, which I test in a field experiment with 20,088 jobs from 8,906 firms on the platform, randomizing mandatory or optional pay disclosure in job ads.

I find that pay non-disclosure steers women away from high-paying firms, while transparency reverses this pattern. The mechanism is threefold: high-paying jobs concentrate in large firms; these

[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 also filter applicants by whether the applicant's expected salary falls within the range the firm provided.

firms are more likely to omit salaries; and less likely to offer flexibility. Pay is therefore most often hidden precisely where it may compensate workers for inflexibility. Because women value flexibility slightly more than men, this friction disproportionately screens them out. Meanwhile, revealing pay nearly doubles women's applications to large firms and closes gender gaps in directed search, even though the wages and amenities at those firms do not change. This shows that gendered sorting arises not because women prefer flexibility to pay, but because they do not observe pay.

These findings are rooted in a sharp disconnect between real and perceived selection into pay disclosure. Large firms are 45% more likely to offer pay above the market average and 34% less likely to disclose it in job ads – making salaries more likely to be hidden when higher, for otherwise similar jobs.[2] Yet only 11% of surveyed job-seekers believe hidden salaries to be higher than posted ones, revealing a clear window for policy action.[3] Large firms also offer fewer family-friendly amenities: they are 27% less likely to provide flexible hours, 37% less likely to offer remote work, and 24% more likely to require overnight travel. The discrete-choice experiment shows 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. By contrast, their preferences for wages and wage transparency match men's.

Men and women respond similarly to pay information, but differently to its absence. Holding job type fixed, when pay is visible, a log-point increase in pay is associated with a 30% increase in male applications and 28% in female applications. But when pay is hidden in otherwise similar jobs, this pattern reverses: male applications become unresponsive to pay – reflecting the lack of information – while female applications decline by 16% with a log-point increase in hidden pay.

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 place equal weight on wages and wage uncertainty, but their flexibility preferences differ. When pay is hidden, wage uncertainty reduces the perceived value of undisclosed wages, discounting the appeal of large firms that hide pay. Two testable predictions follow. First, pay disclosure raises applications to large firms for both genders. Second, even without gender differences in risk aversion or information, transparency has gendered effects due to gendered amenity tastes. If women value flexibility more, more of them may sit on the margin where revealed wages just offset the amenity loss.

The field experiment tests these predictions: it removes pay uncertainty to test whether applications shift toward large firms, and whether the shift is gendered.[4] If revealing wages has no effect, observed gender gaps reflect gendered preferences over non-wage amenities or firm type. If, instead, transparency reallocates women's search toward large firms, both must be true: women learn new wage information, and are willing to choose pay over flexibility when the trade-off is transparent. In that case, information can reduce sorting gaps even in the presence of gendered amenity preferences.

[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 salary ranges, 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 large-firm applications is disproportionate for women. Treatment increases male applications to large firms by 59% and female by 95%.[5] This reverses the gender gap in sorting: in control, women's applications to large versus small firms trail men's by 20 pp. In treatment, this gap *flips* in favor of women to 43 pp, highlighting how incomplete information drives 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 study whether the relative appeal of key amenities changes with availability of salary information. By developing a natural language processing algorithm, I classify ads advertising remote work, flexible hours, transport support, safety, and inclusive language – features often stated in ads and known to matter more for women. Among these, remote work stands out. In control, remote jobs without posted salaries receive 167% more female applications and only 67% more male applications.

Salary disclosure does not change the distribution of amenities, but it sharply reduces their appeal. Female responsiveness to remote jobs falls by 82 pp. Other amenities also fade in importance as the appeal of large firms rises by 46 pp. 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 in the presence of gendered amenity tastes.

A limitation of this analysis is that amenity disclosure in job ads is not complete or 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 information about amenities with salary disclosure in the lab. Each job-seeker evaluates four pairs of stylized jobs in a discrete choice experiment. Ads vary randomly in whether they advertise salary, 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 pay. In fact, salary transparency

[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 and 120% of workers' expected wages.

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. Modest negative effects on average quality are also 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 high-ability women. 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 overestimated 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 pp (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 increases 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-the-envelope extrapolation: a representative woman is 18% more likely to be hired by a large firm in treatment than in control – not because treated large firms become more likely to hire women, but because women are 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 returns to tenure depress entry wages. Third, even if signals point in the right direction, inference is cognitively taxing. 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).

Other mechanisms could, in principle, drive gendered sorting, but none account for the results I observe. I rule out gender gaps in preferences over higher wages, ambiguity aversion, and the ability to infer hidden salaries, firm size or competition levels from the text of job ads that do not state salaries. Differences in perceived wage discrimination at large firms and in network size or composition also do not explain the results. Women are, however, 29% more likely than men to prefer fixed salary offers to bargaining, consistent with Leibbrandt and List (2015). Even so, women with and without bargaining aversion value salary transparency equally, plausibly because transparency allows bargaining-averse women to avoid negotiation, while it helps women who prefer bargaining to calibrate their wage demands. Incorporating bargaining aversion into the model only amplifies the gendered 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 a novel 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, avoid these firms unless they reveal the compensation for inflexibility.

The paper brings together various strands of research – usually studied in isolation – and shows how their interaction generates gendered sorting. It also advances each strand individually. First, hypothetical choice experiments find that women value flexibility more (Mas and Pallais, 2017; Wiswall and Zafar, 2018), but Mas and Pallais (2020) note that these survey valuations diverge from observed sorting – women often forego flexibility in practice – likely because flexibility is bundled with less desirable attributes outside the lab, e.g., lower wages. I test this explanation by replicating gendered preferences for flexibility in a discrete choice experiment, and then showing that behavior diverges in the field because many women are marginal in the face of constrained tradeoffs and incomplete information. Specifically, high-wage, high-flexibility jobs are rarely available in practice. When pay transparency reveals the price of flexibility to be high, workers shift toward inflexible jobs. This has design implications: to approximate field behavior in the lab, lab studies can vary wages and amenities jointly.

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 by pursuing changes to the supply of amenities: 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 in India. These approaches imply that for South Asian 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 longstanding literature shows that large firms pay 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. This provides another mechanism by which firms' recruitment and compensation policies can – often unintentionally – widen gender gaps, alongside flexible pay arrangements (Biasi and Sarsons, 2022), the elicitation of salary expectations and histories (Roussille, 2023; Hansen and McNichols, 2020), pay secrecy (Cullen and Pakzad-Hurson, 2023), gender-targeted job ads (Kuhn and Shen, 2013), and "potential"-based promotions (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 literature 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 low-income setting: as in prior studies, posted salaries in Pakistan rise by about 3% due to pay transparency, while ranges shrink, and applicant quality declines modestly. Third, I overcome a major challenge in prior work: non-random noncompliance with the mandate. Previously, undisclosed salaries were unobserved by researchers (e.g., in Colorado, 30% of postings ignored disclosure requirements). I induce high compliance and leverage my access to hidden salaries to study firm selection into disclosure. Last, I extend the focus to women's directed search to high-paying firms.

The paper is organized as follows: Section 1 describes the setting and Section 2 the data; Section 3 documents facts motivating this research; Section 4 provides a theoretical framework; Section 5 outlines the design of the field experiment; Section 6 lays out the empirical approach; Section 7 presents the results; Section 8 replicates the results in a lab setting; Section 9 discusses alternative supply-side mechanisms; Section 10 summarizes demand-side impacts; Section 11 provides back-of-the-envelope estimates of women's hiring likelihood, and finally, Section 12 concludes.

1 Context

This section describes three features of the paper's empirical environment: the country in which this research is based, the job platform that provides the data and experimental setting, and the global policy context motivating interest in pay transparency reforms.

1.1 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 stands at just 24.4%, far below countries at similar income levels, and less than one-third of the male rate of 81.0% (ILO, 2025). Even among college-educated women, participation rates barely exceed the national average (Bandiera et al., 2025). Among those who do enter the labor market, gender gaps remain stark: Pakistan records the world's highest adjusted gender pay gap at 36.3% – controlling for education, age, employment type, and sector – and the second-highest raw pay gap at 34%, trailing only India at 34.5% (ILO, 2018).

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 employment rate nearly doubled from 13.7% to 23% between 1991 and 2020, while the share of college-educated women nearly quadrupled from just 3.2% to 12.9%. Panel (b) shows that these dramatic gains have nearly halved gender gaps: the ratio of male-to-female employment rates fell from 6.0 to 3.4, and college attainment rates fell from 2.5 to 1.3.

However, these gains in women's education and employment have not translated to gains in where women work. Panel (c) of Appendix Figure A.1 shows that the gender gap in large-firm employment widened sharply: employed men are now 1.5 times more likely than employed women to work in large firms (up from 0.7 in 1991), and twice as likely to hold salaried jobs rather than casual-wage or self-employment (up from 1.4 in 1991).[9] Together these patterns highlight that despite dramatic growth in women's employment and human capital, they are increasingly underrepresented in high-return jobs and workplaces.

[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] In low-income countries, casual wage work and self-employment tend to offer lower returns and fewer formal protections relative to salaried jobs (Breza et al., 2021; World Bank, 2011).

1.2 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), the platform facilitated 29.5 million job applications.

Firms. There are 90,640 formal and informal 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.

Jobs. The platform receives over 100 new job postings daily, amounting to 326,145 job ads from January 2019 to April 2024. The jobs span 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 coverage of low-paying jobs, with about 27% of listings offering salaries at or below the monthly minimum wage.[10]

Job-seekers. About 3.1 million job-seekers are registered on the platform during 2019-2024. These job-seekers are 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.

1.3 Salary disclosure in job ads

Salary is a fundamental attribute of jobs, yet it is often missing in job ads globally. As shown in Figure 1, salary non-disclosure is widespread across countries, regardless of regional or economic context. In eight of the countries, over 80% of job ads omit salary information, including Chile (86.6%), the United States (83.5%), Ethiopia (83.4%) and China (83.0%). In Pakistan, the setting of this paper, 55% of ads do not disclose salaries.

[10] Minimum wage is not strongly enforced in Pakistan.

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 broader interest in pay transparency of different kinds; over the past two decades, 71% of OECD countries have adopted laws ranging from protecting workers who share pay information with colleagues from being penalized, to requiring firms to disclose internal pay data, including the reporting of gender pay gaps (Cullen, 2024). As such policies proliferate, assessing their effectiveness and broader economic consequences has become increasingly important.

2 Data and descriptive statistics

Platform 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).[11] 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 2.85/4 (SD: 0.80). Job-seekers submit around 8.84 applications (SD: 46.14). 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

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

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 those that men target. These differences underscore the importance of application behavior in driving downstream gender gaps in employment and pay.

Baseline firm surveys. I conduct baseline surveys with key decision-makers in over 300 firms to document how firms make salary disclosure decisions. These firms were randomly sampled from the platform's registry, conditional on 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 or executives), 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, and female share of applicants (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 in the field experiment. 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. Survey links were de-activated after the target sample size was reached. To incentivize

participation, respondents were offered a financial reward of PKR 5,000 (USD 17.69). The invitation list was randomly drawn from the approximately 310,000 job-seekers who were active in sending applications during the experiment. 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, job attributes and application decisions. It includes modules that assess the accuracy of respondents' beliefs about hidden salaries and associated application volumes (i.e., perceived competition for the job) and firm size. These beliefs are benchmarked against administrative data from the platform. The survey also measures preferences over bargaining, and ambiguity; the latter is elicited through a standard incentivized ambiguity aversion Ellsberg Paradox 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, respondents evaluate four job pairs (eight jobs total), where salary visibility, a large-firm label, remote work, and equal-opportunity language are independently randomized. Examples of job pairs presented to respondents are shown in Appendix Figure A.6.

3 Equilibrium facts

This section leverages platform and survey data to document a set of equilibrium facts about firms and workers that jointly explain why women may sort into low-paying firms.

Background: Why firms and firm-size matter. 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 critical 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 over 30% more than those in firms with fewer than 25 workers – a disparity as large as the gender wage gap itself. 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 high-return. Thus, understanding access to them is key to explaining labor market inequality.

Measurement: Firm size in this paper. Building on the literature, I use firm size as a proxy for access to high-return workplaces. This measure is particularly useful given my research design: while salaries may be directly affected by my pay transparency experiment (discussed later), firm size is a characteristic that remains fixed. As pre-registered, it provides a stable baseline along which

to test whether workers reallocate their search from low- to high-wage firms when salaries are revealed. I classify firms as large if they have over 50 employees. This threshold corresponds to the median firm size in my job-level data, and is frequently the threshold at which labor regulations bind. For instance, in most states in the U.S., pay transparency mandates apply to firms with more than 50 employees. However, results in the paper are robust to alternative definitions.

The following facts document how large firms' recruitment practices interact with workers' job search behavior to shape gender gaps in sorting across firms.

Fact 1: Large firms pay more for otherwise similar jobs. While prior work has documented that large firms pay more for otherwise similar workers, I find that this premium exists even at the job-posting stage. Jobs posted by large firms list substantially higher pay: the posted minimum is 23.1% higher and the posted maximum is 17.5% higher than at smaller firms. Large firms are also 44.5% (8.1 pp) more likely to post salary ranges in which both the minimum and maximum exceed the market average. To isolate size-specific pay differences from other job attributes that vary between small and large firms, Figure 3 Panel (a) reports salaries residualized on a comprehensive set of job-level controls (occupation, industry, city, required education and experience, career level, and work schedule). Even after this adjustment, the size premium persists: relative to firms with 1–10 employees, firms with over 301 employees pay at least 20.3% more. These patterns emerge before any worker selection or match-specific surplus generation, suggesting that small and large firms reward otherwise similar jobs differently from the outset.

Fact 2: 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 42.8% of job postings from the smallest firms (1-10 employees) hide salary, the rate climbs to 75.4% among firms with 301-600 employees and remains high at 73.7% for firms with over 600 employees. This pattern suggests that salary non-disclosure is most common among the largest and highest-paying firms. Moreover, leveraging a novel feature of the platform where firms must report salary ranges internally even if they choose to hide them from job-seekers,[12] I find that large firms' disclosure likelihood falls disproportionately with salary: one-log-point increase in the posted salary is associated with a 9.0 pp decrease in the probability of disclosure, whereas for small firms, the corresponding decrease is about half as large (4.6 pp). 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 thus, strongest impact on applications – among high-paying, large-firm jobs.

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

Fact 3: Salary disclosure is driven by firm – not job – characteristics. The two preceding facts yield a striking implication: jobs that offer the highest pay are also most likely to conceal it in job ads. One interpretation is that this pattern reflects firm behavior – larger firms both pay more and disclose less. Alternatively, certain types of jobs may be both higher-paying and less likely to disclose salaries, with these jobs concentrated in large firms. To distinguish between these explanations, I leverage novel access to hidden salaries to test whether salary levels rise with non-disclosure even within jobs of similar type. Appendix Figure A.8 reveals that they do: for otherwise similar jobs, salaries are more likely to be hidden when they are higher. A log-point increase in the minimum, midpoint, or maximum salary is associated with a 2.6%, 8.6%, and 9.4% increase in the probability of non-disclosure, respectively. This suggests that job characteristics alone cannot fully explain disclosure patterns. While job-level traits such as career level, education, and experience requirements do predict salary 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%. Appendix Figure A.4 confirms this finding using the Bayesian information criterion. Together, these results point to the firm, not the job, as the dominant driving force behind salary non-disclosure.

Fact 4: Large firms offer fewer 'family-friendly' amenities. In many low-income countries, micro and home-based enterprises are pervasive, where formal human resource departments tend to be absent and work arrangements are relational or informal: workers bring children to work, shift hours, take unplanned leave, or work from home (World Bank, 2011). As firms expand and professionalize, employees may lose the ability to negotiate flexible work arrangements, and encounter more demanding and competitive environments that reward long hours, in-person visibility, limited time off, and networking. While such requirements are often compensated with higher pay, they are less appealing to women, even in high-income countries (Blau and Kahn, 2017; Goldin, 2014; Cullen and Perez-Truglia, 2023; Niederle and Vesterlund, 2007; Niederle et al., 2014). Matched employer-employee data confirm that women's employers are not only lower-paying but also more family-friendly: they are located closer to home, offer more part-time and stable schedules, pay that is less elastic to hours worked, and employ a disproportionate share of women with young children (Pertold-Gebicka et al., 2016; Sorkin, 2017; Hotz et al., 2017; Kleven et al., 2019; Vattuone, 2023; Tô et al., 2025). I find similar patterns in my context. For a limited time prior to my experiment, my partner job platform mandated reporting of family-friendly amenities. Table B.4 uses these baseline data to examine the relationship between such amenities and firm size. Panel A presents firm-level results. Large firms are 10 pp (17.9%; p<0.01) less likely to offer flexible work hours (Column 1) and 6 pp (23.1%; p<0.01) less likely to offer remote work (Column 3), but 7 pp (116.7%; p<0.01)

[13] Jobs requiring more education and experience, or situated at higher career levels, are more likely to conceal salary information. See Appendix Figure A.7. These patterns align with prior findings from the U.S. and Chile (Batra et al., 2023; Banfi and Villena-Roldán, 2019).

more likely to require vehicle ownership (Column 5) and 5 pp (20.8%; p<0.01) more likely to require overnight travel (Column 7). These patterns persist after controlling for industry fixed effects. Panel B examines job-level patterns, and the associations strengthen: large-firm jobs are 12 pp (26.7%; p<0.01) less likely to offer flexible hours (Column 1); 7 pp (36.8%; p<0.01) less likely to offer remote work (Column 3); 3 pp (42.9%; p<0.01) more likely to require vehicle ownership (Column 5) and 5 pp (23.8%; p<0.01) more likely to require overnight travel (Column 7). All correlations except overnight travel remain significant and large despite controlling for industry and occupation fixed effects, education and experience requirements, and career level of the job (even-numbered columns). Finally, while large firms are 3 pp (12.5%) more likely to provide transport support at the firm level (Panel A, Column 9), this association becomes statistically insignificant with industry controls and turns negative at the job-level, particularly after including controls (Panel B, Columns 9-10; p<0.05). Large firms are also 1 pp (11.1%; p<0.01) less likely to have jobs with female supervisors (Column 10) but this coefficient attenuates and becomes insignificant after controlling for job characteristics (Column 11). Overall, these patterns confirm that large firms systematically offer fewer family-friendly amenities.[14]

Fact 5: Women have a modest preference for flexibility. To measure preferences for job attributes, 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:

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

Here, $\mathbf{X}_{wij}$ is a vector of these randomized job attributes, 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 estimated coefficients by gender. It shows that women place significantly greater value on flexible work arrangements: they are 6.51 pp more likely than men to choose remote jobs (p=0.032), consistent with Mas and Pallais (2017). By contrast, women are 6.24 pp less likely to prefer jobs at large firms (p=0.043). These are the only two attributes that exhibit statistically significant gender gaps. In contrast preferences for salary visibility, inclusive language, and transport subsidies show no meaningful differences between men and women.

[14] In high-income countries, large firms often leverage economies of scale to provide on-site childcare facilities at higher rates than small firms. However, formal childcare arrangements are exceedingly rare in Pakistan, where families typically rely on informal care by relatives. Consequently, employer-provided childcare is not a commonly offered amenity by firms of any size in this context.

Fact 6: Both genders lack salary information, and value transparency similarly. Job-seekers of both genders seem unaware that salaries are higher when hidden for otherwise similar jobs: only 11% of respondents (12.8% of men and 10.1% of women) in the job-seeker survey believe that hidden salaries are higher than disclosed ones. This creates a clear opportunity for pay transparency to correct beliefs and redirect search. Moreover, both genders share a similar inability to infer wages from non-wage characteristics of a job. I establish this by presenting job-seekers with two job ads from their occupation that had been recently posted on the platform without a salary range, one by 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 salary reported by the firm to the platform. I find that an overwhelming majority of male and female respondents fail to guess salaries accurately, and there is no detectable gender gap in guessing accuracy (for details on the results, see Section 9). Unsurprisingly then, I also find that both men and women value pay transparency highly, and to similar degrees. The DCE results reported in Table 1 reveal that salary visibility more than doubles the odds of job-choice for men and women, and the coefficients are statistically indistinguishable across genders (odds-ratios of 2.14 and 2.65, respectively with p=0.152).

Fact 7: Women sort on salaries as much as men, but only when salaries are revealed. I examine application behavior and show that women's responsiveness to pay mirrors men's when salaries are disclosed but diverges sharply when salaries are hidden. Figure 4 plots the relationship between job applications and salary by gender, separately for postings with and without disclosed pay. Application counts are normalized by gender-specific means to account for the fact that women make up only 27% of job-seekers on the platform, so raw application counts are mechanically lower for women than men. All specifications include industry, occupation, and city fixed effects, as well as detailed job characteristics – experience and education requirements, career level, number of vacancies, required skills, and job schedule. The results show that men and women respond similarly to wage information when it is available: both increase their applications to higher-paying jobs, with a 30% and 28% rise in applications, respectively, per log-point increase in salary. When salaries are hidden, however, the gender parity breaks down. Women's application rate flips sign, declining by 16% per log-point increase in salary, while men's becomes flat. Thus, in the absence of wage information, men appear to apply broadly across the pay distribution – consistent with random search – while women actively avoid higher-paying jobs. The difference in slopes is statistically significant (p<0.01).

4 Theoretical framework

This section formalizes how salary non-disclosure amplifies amenity preferences, driving gendered sorting across firms. The key insight is that pay transparency can generate gendered sorting even without gender differences in preferences over wages, information about wages, or aversion to wage

uncertainty. The primitive forces underlying this model come directly from the empirical facts in Section 3.

4.1 Workers

A continuum of workers chooses where to apply based on job postings. Each worker is characterized by gender $g \in \{f, m\}$ where f is a female worker and m is a male worker. Expected utility from a job j offering wage w and amenity level a is

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

where $\gamma \geq 0$ denotes the worker's amenity preference drawn from a gender-specific distribution P_g with density p_g, and $\theta \in (0, 1)$ captures risk aversion over wages.

When the salary is visible ($d = 1$), a job-seeker observes the exact wage w and evaluates utility directly at $w^\theta a^\gamma$. When the salary is hidden ($d = 0$), the worker does not observe the realized wage ex ante and must evaluate a distribution of possible wages, $\mathbb{E}[w^\theta]$.[15]

For $\theta \in (0, 1)$, the function $u(w) = w^\theta$ is concave. Hence, Jensen's inequality implies $\mathbb{E}\left[w^\theta\right] < \left(\mathbb{E}[w]\right)^\theta$, whenever w is non-degenerate. Intuitively, evaluating an uncertain wage through a concave utility function yields lower expected utility than receiving the mean wage $\mathbb{E}[w]$ for sure. Hidden pay thus enters expected utility with a discount relative to a posted wage equal to $\mathbb{E}[w]$.

In addition, wage non-disclosure leads to a pooling of jobs. When wages are hidden, applicants assign any posting the pooled value $\mathbb{E}[w^\theta]$ rather than its realized utility w^θ. A job that in fact pays $\tilde{w}$ incurs a perceived utility loss of $\tilde{w}^\theta - \mathbb{E}[w^\theta]$, which is increasing in $\tilde{w}$. Consequently, wage hiding imposes a larger penalty on genuinely high-paying jobs than on jobs with wages closer to the mean, disproportionately dampening applications to positions that would otherwise stand out as high-paying. The model extension in Appendix Section C explores this implication in more detail.

Finally, when pay is hidden, wages contribute less to perceived utility, and perceived wage differences across hidden-pay postings are compressed. As a result, within the hidden-pay set, application choices are driven more by amenities. So holding true preferences fixed at γ_g, workers behave as if they are more amenity-loving when wages are unobserved because jobs with better amenities look relatively more attractive.

[15] The model assumes a common prior across genders because empirically I do not find gender differences in beliefs about hidden salaries (see Section 9). Extensions allowing women to have systematically worse wage information or biased beliefs would amplify transparency's benefits for female workers without altering the core mechanism.

4.2 Jobs

The worker-side formulation highlights two general forces. First, wage non-disclosure imposes an uncertainty penalty on jobs even when workers have rational expectations. This penalty is increasing in the job's true wage. Second, by discounting and compressing perceived wage differences across postings, hidden pay magnifies the role of observable amenities in application decisions.

To translate these forces into testable predictions about sorting, I now introduce structure on the job side. Specifically, I model the set of jobs workers face in the market and how wages, amenities, and disclosure policies systematically vary across them using the empirical patterns documented in my data (Section 3). This allows me to generate testable predictions about how transparency can induce reallocation of search in my setting.

I allow for 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 family-friendly 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 in job ads 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 cutoff 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 similar comparative-statics. Additionally, the model extension in Appendix C allows for wages to be stochastic rather than determinate, and accounts for empirical differences in wage dispersion across small and large firms, but preserves the core insights of the simpler framework presented below.

4.3 Parameters and assumptions

I discipline the preference parameters (γ, θ) using my lab-in-field discrete-choice experiment (DCE) where workers evaluate four pairs of stylized job ads that vary along five randomized attributes: salary visibility, a large-firm label, remote work, transport subsidy, and inclusive language ("equal opportunity employer"). Preferences are then estimated using a conditional logit model.

Risk aversion over wage uncertainty. The DCE results confirm that workers dislike pay uncertainty. The coefficient on salary visibility (Table 1) measures the gain in latent utility when wages are shown rather than hidden – essentially, the reduced-form counterpart to removing wage uncertainty in the model. It is large and positive for both genders: the implied odds ratios are around two, meaning that non-disclosure reduces job-choice odds by about 0.4. The estimated effect is also statistically indistinguishable across genders ($p = 0.152$), supporting the use of a common risk-aversion parameter θ for men and women.

Gender-neutral wage preferences. A common θ also implies a shared valuation of wage levels themselves, consistent with two pieces of evidence. First, in observed application behavior, men and women sort on salaries equally when pay is revealed (see Fact 7), indicating similar responsiveness to wage information. Second, in the DCE, both genders display nearly identical preferences for higher-paying jobs once wages are visible (discussed in more detail in Section 9.2). Together, these findings support common preferences over wages. Allowing women to value higher wages more would mechanically generate larger effects of transparency for women but the data provide no evidence for such differences.

Gendered amenity preferences. The DCE shows that women choose flexible jobs (remote work) 6.51 pp more often than men ($p = 0.032$). Motivated by this, I allow amenity preferences for the two genders to come from different distributions with CDF $P_g(\gamma)$ and density $p_g(\gamma)$. The common wage risk parameter (θ) combined with a gender-specific parameter for amenities (γ_g) means that any gendered response to transparency will operate through the amenity channel.

Limited inference of wages from non-wage characteristics. The model implicitly assumes that workers evaluate jobs only on directly observed attributes, and do not systematically infer an unobserved wage w from observed amenities a, firm size j, or disclosure d.

(i) Amenities. Abstracting from wage inference based on amenities is a conservative choice. In the job-level data of this and other studies (e.g., Mas and Pallais (2017)), cross-sectional wage–amenity correlations are weak and often positive because they conflate two forces with opposite signs: within a job type, amenities carry a negative compensating differential (they reduce the wage required to attract applicants to a given job), but across jobs of different types, higher-wage employees can afford to "buy" more amenities (Lavetti, 2023). Thus, amenity-rich jobs pay less than they would absent the amenity, yet still pay more than amenity-poor jobs because of their higher unobserved productivity. The positive component often dominates, making the cross-sectional correlation upward-biased. If workers nevertheless extrapolate from this weak but positive correlation, they would interpret low-amenity postings to be a signal of low wages, making large-firm jobs look more unattractive. Revealing large-firm wages through pay transparency would then generate even larger reallocations

toward these firms. Thus, ignoring this channel biases the model against finding large transparency effects.

(ii) Firm size and disclosure. Meanwhile, the assumption that workers do not infer hidden wages from firm size and disclosure is motivated by survey evidence. Empirically, large firms both pay more and disclose pay less, implying that hidden salaries should exceed posted ones in expectation. Yet, only $\sim$11% of surveyed workers believe that hidden salaries are higher than posted ones – despite correctly identifying large firms from the text of job ads $\sim$73% of the time. This pattern is consistent with workers either lacking information about the correlation between firm size and pay or neglecting to apply it when evaluating jobs with hidden salaries. The former is plausible because while large firms offer higher pay on average, they also have higher wage dispersion, particularly offering lower wages at the point of entry (Song et al., 2019). Moreover, when they do disclose, it's typically for lower-paying jobs, creating misleading signals. In either case, workers do not reliably infer missing wage information from observable firm size. This bias is accounted for in the comparative statics described below.

(iii) Extension: Rational expectations. Appendix C extends the baseline model presented in this section, allowing wages to be draws from firm-size–specific distributions, and workers to form rational expectations about hidden salaries based on these distributions. This extension shows that under the expected utility framework, wage non-disclosure compresses the continuum of large-firm wages into a single risk-discounted certainty equivalent which – for sufficiently high risk aversion – can make small firms more appealing to workers even when large firms pay more and workers know this fact. Pay transparency reverses this by allowing workers to sort on realized wages, increasing applications to large firms by leveraging their greater mass of high-wage realizations.

4.4 Sorting under different information regimes

Consider a setting in which each worker applies to a single job and chooses the option that maximizes expected utility. The model then predicts how salary transparency affects the sorting of male and female workers between large and small 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. Workers prefer the large-firm job if the expected utility from its (hidden) wage and lower amenities exceeds that of the small-firm job:

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

Rearranging and taking logs:

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

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:

$$\gamma \;\leq\; \gamma^C \;\equiv\; \frac{\theta \Delta w - \kappa}{\Delta a}. \tag{2}$$

This defines a cutoff 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 L_C line represents expected utility of the large-firm job that hides wages, while S represents the small-firm job's utility. L_C has a flatter slope than S because as γ rises, utility of the small-firm job rises faster as it offers more flexibility ($a_S > a_L$). L_C, however, has a higher intercept than S. This is because the figure depicts the parameter region in which a nonzero share of workers applies to large firms under opacity ($\mathbb{E}[w^\theta] > w_S^\theta$). In the limiting case of very high risk aversion (θ close to 0), or a negligible large-firm wage premium, the curves may not intersect and no worker would apply to L_C. This edge case is theoretically possible but less relevant empirically.

The intersection of γ^C with the utility functions shows where workers are indifferent between large and small firms under the control regime. Those with low amenity demand ($\gamma \leq \gamma^C$) choose the large-firm job despite its hidden wage, while those with higher amenity demand ($\gamma > \gamma^C$) prefer the more flexible small-firm option. The cutoff γ^C 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 risk aversion captured in the uncertainty penalty (κ).

When the wage gap between large and small firms becomes arbitrarily large ($\Delta w \to \infty$), the cutoff 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 cutoff collapses to zero. In this case, only the worker with $\gamma = 0$ 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 $w_L^\theta a_L^\gamma$ versus $w_S^\theta a_S^\gamma$. 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 cutoff is therefore:

$$\gamma \;\leq\; \gamma^T \;=\; \frac{\theta\,\Delta w}{\Delta a} \;=\; \gamma^C + \frac{\kappa}{\Delta a} \;>\; \gamma^C. \tag{3}$$

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 new cutoff inherits the same comparative statics as before, except that it no longer falls with κ because $\kappa = 0$ under transparency.

Switcher band. 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. 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 switcher band is highlighted by the gray shading in Figure 5.

Amplification of relative amenity weight. Workers trade off wages and amenities according to the ratio γ/θ: a higher γ/θ implies a stronger preference for amenities relative to pay. Under transparency, a worker chooses the large firm (L) if the wage premium compensates for the amenity loss:

$$\theta\,\Delta w \;\geq\; \gamma\,\Delta a \quad\Longleftrightarrow\quad \frac{\Delta w}{\Delta a} \;\geq\; \frac{\gamma}{\theta}.$$

Thus, γ/θ defines the required wage premium per unit of lost amenity. Uncertainty about pay effectively reduces the perceived value of the wage premium by κ:

$$\theta\,\Delta w - \kappa \;\geq\; \gamma\,\Delta a \quad\Longleftrightarrow\quad \frac{\Delta w}{\Delta a} \;\geq\; \frac{\gamma}{\tilde{\theta}}, \qquad \tilde{\theta} \equiv \theta - \frac{\kappa}{\Delta w}.$$

Because $\tilde{\theta} < \theta$, the effective ratio $\gamma/\tilde{\theta}$ is larger than γ/θ:

$$\frac{\gamma}{\tilde{\theta}} \;=\; \frac{1}{1 - \kappa/(\theta\,\Delta w)} \cdot \frac{\gamma}{\theta} \;>\; \frac{\gamma}{\theta}$$

Wage uncertainty therefore amplifies the relative weight on amenities: workers with a given γ preference appear more amenity-loving and behave as though amenities matter more compared to wages. When $\tilde{\theta} \leq 0$ (i.e., $\kappa \geq \theta\Delta w$), uncertainty is so high that the large firm is strictly dominated, and the trade-off is irrelevant.

Gender-specific response. The mass of gender g workers who switch from small to large firms is:

$$\text{Switchers}_g = \int_{\gamma^C}^{\gamma^T} p_g(\gamma)\,d\gamma = P_g(\gamma^T) - P_g(\gamma^C)$$

This switching mass depends entirely on the gender-specific density $p_g(\gamma)$ in the switcher band $(\gamma^C, \gamma^T]$. Since amenity preferences differ by gender, transparency generates heterogeneous responses by gender.

One possible scenario is illustrated in the lower panel of Figure 5. The blue curve shows the CDF $P_m(\gamma)$ for men, and the red curve shows $P_f(\gamma)$ for women, with the female CDF lying to the right. The switching masses is illustrated by vertical arrows Δ_m (blue) and Δ_f (red). They reflect each gender's density in the switcher band. If women's distribution places more mass than men's in this interval – consistent with the DCE finding that women value flexibility moderately more – then transparency will induce proportionally more women to switch to large firms, illustrated as a longer red arrow in the figure. The field experiment is designed to test whether this scenario holds.

Alternatively, the differential effect by gender could emerge from other factors typically invoked in the literature, such as gender differences in risk preferences, information asymmetries, biased beliefs, or discrimination. All of these are inconsistent with my empirical evidence, discussed at length in Section 9. While these channels would amplify transparency's effects if present, the model's insight is that they are not necessary. Even with gender-neutral risk aversion, and identical wage preferences and beliefs, a differential response can emerge from the interaction of wage uncertainty with differences in amenity preferences. Moreover, even modest differences in flexibility valuations can generate large gendered responses when preferences cluster near the wage-amenity trade-off margin.

4.5 Testable predictions

The framework above delivers two testable predictions:

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.* If women's amenity preference distribution (P_f) places more mass than men's (P_m) in the interval $(\gamma^C, \gamma^T]$, then pay transparency will induce a larger increase in women's application rates to large firms relative to men's: $\big[\Pr_f(L \mid T) - \Pr_f(L \mid C)\big] > \big[\Pr_m(L \mid T) - \Pr_m(L \mid C)\big]$.

Takeaways. The framework embeds two potential forces underlying gendered sorting: preferences for amenities, and information frictions. Posting wages does not erase gendered amenity tastes, but it shifts a disproportionate share of more amenity-loving workers to L by removing the information friction. Thus, pay non-disclosure has gendered impacts even when men and women share identical risk preferences, valuation of wages, access to pay information, and beliefs about hidden pay.

4.6 Implications for research design

The literature uses two main approaches to infer how gender differences in amenity preferences shape sorting: (i) observational analyses with matched employer-employee data and (ii) hypothetical discrete-choice vignettes in surveys. My theoretical framework shows that when pay is undisclosed, these designs can overstate the role of preferences.

Observational studies. Studies such as DeLeire and Levy (2004); Hotz et al. (2017); Vattuone (2023) infer gender differences from observed job moves or allocations. This approach interprets choices as full-information trade-offs between wages and amenities. Under pay non-disclosure, however, risk-averse workers evaluate jobs at a certainty-equivalent wage below the expected value. Assuming full information for observed choices made under opacity will misattribute the uncertainty discount to stronger amenity preferences, inflating estimates of gender differences in tastes.

To see this, consider matched employer-employee data revealing large-firm employment shares: $P_f(\gamma^C)$ for women and $P_m(\gamma^C)$ for men. A researcher observes sorting at threshold $\gamma^C = \frac{\theta\Delta w - \kappa}{\Delta a}$ but assumes workers have full information and thus should sort at $\gamma^T = \frac{\theta\Delta w}{\Delta a}$. To rationalize the observed sorting under this incorrect assumption, the researcher must infer a distorted preference distribution $\tilde{P}_g$ such that:

$$\tilde{P}_g(\gamma^T) \;=\; P_g(\gamma^C).$$

Since $\gamma^T = \gamma^C + \kappa/\Delta a$, this implies

$$\tilde{P}_g(\gamma) \;=\; P_g\Big(\gamma - \frac{\kappa}{\Delta a}\Big),$$

i.e., the inferred distribution is shifted to the right by $\kappa/\Delta a$: preferences appear stronger than they truly are because the uncertainty discount is absorbed into γ. Critically, this misattribution is not gender-neutral. The observational study would conclude the gender gap in preferences is:

$$\text{Inferred gap} \;=\; \tilde{P}_m(\gamma^T) - \tilde{P}_f(\gamma^T) \;=\; P_m(\gamma^C) - P_f(\gamma^C),$$

whereas the true gap under full information is

$$\text{True gap} \;=\; P_m(\gamma^T) - P_f(\gamma^T).$$

Hence, the overstatement is

$$\text{Bias} \;=\; [P_m(\gamma^C) - P_f(\gamma^C)] - [P_m(\gamma^T) - P_f(\gamma^T)] \;=\; \Delta_f - \Delta_m$$

which equals exactly the differential response to transparency that the experiment identifies. Thus, observational studies that ignore pay non-disclosure mechanically over-attribute sorting to preferences by exactly the amount that the experiment recovers when pay is posted.

This will be positive when women's amenity preference distribution (P_f) places more mass than men's (P_m) in the switcher band. Appendix Figure A.9 illustrates this bias visually: the gray arrow shows the inflated preference gap that observational studies would infer at γ^C under hidden pay, i.e., $P_m(\gamma^C) - P_f(\gamma^C)$. The right coral arrow shows the smaller, true preference gap at γ^T under full information, i.e., $P_m(\gamma^T) - P_f(\gamma^T)$. The difference between them is exactly $\Delta_f - \Delta_m$, the bias derived above.

Hypothetical discrete-choice vignettes. Vignette designs present respondents with full information on wages and amenities, randomize the two independently, and use the induced choice responses to recover willingness to pay for amenities (e.g., Wiswall and Zafar (2018); Mas and Pallais (2017)). This approach isolates preferences under full transparency (γ^T) but abstracts from two features that shape realized choices outside the lab. First, wages and non-wage amenities often co-vary in the market, and choices differ depending on whether they are made along a constrained or unconstrained frontier. This is shown in Appendix Figure A.10 where vignettes present choices from the full wage-amenity space (including purple dots), while actual job choices are limited to the market frontier (gray band) where wages and amenities must be traded off. Second, vignettes always reveal wages but actual sorting often occurs under pay non-disclosure at γ^C. This overlooks that many women with moderate flexibility preferences are marginal (sit in Figure 5's gray shaded "switcher" region) where changes in wage information can shift choices while holding tastes fixed. Thus, vignette estimates have strong internal validity for measuring the marginal rate of substitution between wages and amenities under full transparency. However, they may not predict observed sorting well, as noted in Mas and Pallais (2020), because it occurs along a constrained frontier and under incomplete information. Vignette designs that jointly vary amenities and wages, reflecting their true covariance can overcome some of these challenges.

Field experiment. Observational studies can overstate the role of preferences by inferring them from choices made under incomplete information. By contrast, lab experiments recover true preferences by removing frictions, but miss how frictions distort behavior at the margin. A field experiments that randomizes pay transparency can address both limitations. It removes information frictions (as in the lab), while preserving realistic trade-offs (as in the field) to capture marginal responses to the interaction of information with preferences.

5 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 to randomly assign 20,088 job ads posted by 8,906 firms to one of two conditions. In the treatment group, firms are required to display salaries in their job ads; in the control group, they retain the option to hide them.

Job-level randomization. Treatment is assigned automatically at the time of job posting. When a firm representative initiates a new post, a random number drawn by the platform's backend system determines assignment to the treatment or control group. Appendix Figure A.13 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. Clicking the "Learn More" link opens a pop-up window (Figure A.11, 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 firms assigned to treatment are automatically enrolled in mandatory salary disclosure. However, those that object to participation can request to opt out. Opting out is intentionally designed to impose a small but salient administrative cost on firms – filing a request to withdraw by phone, and a two-to-three day delay in posting the job due to manual processing by platform staff. When firms contact the platform to request withdrawal, staff follow a standardized phone script explaining the study's purpose, the random assignment of treatment, and the rationale for transparency. Firms that still wish to opt out after this explanation formally confirm their decision, after which their job post is reclassified to allow salary non-disclosure. Platform staff are instructed to approve all such requests, regardless of how strongly firms make the case for non-disclosure. Importantly, even firms that opt out of public disclosure must still privately report a salary range to the platform. This ensures that salary information is recorded for all job ads, including those that do not comply with public disclosure. This feature is a central strength of the experimental environment as it allows direct analysis of selection into non-compliance. By contrast, while studies of wage transparency laws often document substantial non-compliance, they cannot

examine its determinants since undisclosed salaries are unseen by both job-seekers and researchers. For example, despite a legal mandate in Colorado, nearly 30% of job ads continue to omit salary information (Arnold et al., 2022).

Managing inattention. Although the treatment interface replaces the option to hide salary with new text explaining the salary disclosure requirement, some inattentive firm representatives may overlook this change and inadvertently disclose salaries. To guard against this, the platform displays an additional confirmation pop-up immediately before the job post is finalized. This message reminds the firm that the salary information they have entered will be publicly visible and provides an opportunity to review or revise it before submission. The confirmation screen, shown in Figure A.11, Panel (b), ensures that treatment participation reflects an informed decision rather than inattention to the modified interface.

Managing strategic manipulation. Firms cannot evade treatment assignment by refreshing the webpage or re-posting the same job under an identical title. The platform links treatment status to the initial attempt to create a given job post, ensuring that it remains fixed even if the post is left incomplete and resumed later. To further guard against evasion, a dedicated sales team reviews all ads before publication. When a firm attempts to bypass treatment by making minor edits 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 reinstates the original treatment status. This enforcement process preserves fidelity to the experimental design and minimizes scope for strategic non-compliance. Finally, to address any residual inconsistencies due to platform oversight, the analysis aggregates all job titles posted by the same firm into a single observation, assigning it the treatment status from the job's first appearance on the platform during the experiment.

Saturation design. Randomized controlled trials rely on the assumption that a given unit's treatment assignment does not affect the outcomes of its neighbors (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.12. 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 enables identification of both the direct effects of salary disclosure on treated jobs and the indirect spillover effects operating through competitors' disclosure behavior. Comparing outcomes across markets with different saturation levels provides an explicit empirical test of potential interference between units.

6 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. I also report OLS estimates to assess whether these percentage changes translate to economically meaningful differences in application volumes.

7 Field experiment: Results

This section presents the empirical results from the field experiment, starting with compliance and first-stage, and then discussing impacts on wages and job applications. After this, I use the cluster-saturation design to test for spillover effects and explore how transparency reshapes the relative importance of non-wage amenities such as flexibility and remote work.

7.1 First stage and compliance

Table B.5 shows that compliance with the randomized assignment is high. The table reports first-stage results, where the dependent variable is an indicator for whether the job ad displays a visible salary. In the control group, 56.2% of job posts include salary information. Assignment to treatment raises this share by 40.1 pp, implying a non-compliance rate of only 3.7% among treated jobs.

Column 2 disaggregates the results by firm size and shows that most non-compliance arises among large firms. At baseline, salary disclosure is 19.2 pp more common among small than large firms, with 64% of small firms and 44.8% of large firms voluntarily disclosing salaries. In response to treatment, disclosure rises by 35.8 pp among small firms, yielding near-perfect compliance (0.3%). Meanwhile, it rises by 47.0 pp among large firms, corresponding to a moderate non-compliance rate of 8.3%.

Figure A.15 further shows that the number of job postings and active firms remains stable after the intervention, indicating that the disclosure requirement did not reduce platform participation. The combination of negligible non-compliance and stable participation likely reflects the platform's substantial market power. Most firms appear to view adherence to platform rules as a prerequisite for accessing the country's largest pool of job-seekers. Large firms may hold somewhat greater leverage, as their potential exit could reduce the platform's appeal to job-seekers. This leverage likely makes them more comfortable requesting exceptions to platform rules. However, because such firms constitute only a small share of all participants, overall compliance remains very high.

7.2 Advertised salaries

Salary disclosure requirements can raise or lower posted wages. Salaries may increase as firms compete for talent, especially since firms that withhold pay offer 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 treatment on 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 (p<0.05). Treatment also has a negative and statistically significant effect of -0.03 (p<0.01) on the width of the range (defined as maximum minus minimum divided by the midpoint) indicating that firms tend to post narrower salary ranges in response to the disclosure requirement.

That salary ranges become more compact runs counter to the notion that firms may widen posted ranges to obscure pay information. In the firm survey, employers report that widening the range would attract more applicants. This response suggests that firms expect job-seekers to anchor on the upper end of the range when forming wage expectations. As a result, firms may refrain from widening ranges if they anticipate that doing so would attract excess applications or raise applicants' expected pay and complicate subsequent wage negotiations.

The results documented here 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.

7.3 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 sort in response to salaries. To test whether providing 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.16, include salary range information when available, and thus, an applicant choosing to view further details of the job may be responding to this 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.[16] 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).

Why do these effects concentrate among large firms? First, these firms are significantly more likely to withhold salary information: Column 2 of 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 (Column 1), and minimum salaries are 32% higher (Column 2). Moreover, treatment does not affect salary levels at large firms. While small-firm minimum salaries increase by 4%, 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 signal for these jobs. Together, these results suggest that workers respond strongly when new information reveals higher pay.

Women. The model predicts that salary transparency increases applications for both men and women, but the effects may be gendered if preferences over non-wage amenities are gendered. Figure 7 tests heterogeneity in treatment effects by gender. In the full sample (on the left), treatment increases male applications by 46.1% from a control mean of 65.7 applications, and female applications by 60.2% (control mean = 15.9). This amounts to a gender difference of 14.1 pp (p-value<0.1) in favor of women. In the large-firm subsample (right panel), where wages were less accessible prior to treatment, the disclosure requirement makes higher salaries visible to job-seekers. Accordingly, both the treatment effects and gender differences are more pronounced among large-firm jobs. Male applications to large firms increase by 58.8% (from a mean of 89.9), while female applications nearly double (95.2% from a control mean of 20.6), yielding a 36.4 pp differential effect in favor of women.

I next consider whether salary disclosure changes the gender gap in directed search toward large firms. Figure 8 plots the ratio of applications that large firms receive relative to small firms, separately by gender and treatment status. In the control condition (left panel), for every application a small firm gets, large firms receive 1.82 male applications and 1.62 female applications – a gender gap of 20.4 pp in favor of men. Under treatment (right panel), large firms receive 2.31 male applications and 2.74 female applications for every small-firm application. Thus, mandating pay transparency not only raises application volumes; it also reallocates applications from small to large

[16] 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.

firms. This reallocation is stronger for women than for men by 42.6 pp, reversing the gender gap observed in the control group.

The same results are reproduced in tabular form in Appendix Table B.6, while a corresponding graph overlaying OLS and Poisson treatment effects can be found in Appendix Figure A.17.

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. Second, 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.14 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 considerable 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. I test for three potential channels: (i) application spillovers to control jobs arising from a higher intensity of treated competitors; (ii) firm disclosure responses to competitors' treatment status; and (iii) heterogeneity in treatment effects by local treatment intensity.

The first test considers whether application volumes to control jobs vary with the treatment intensity of the labor market. If transparency generates additional traffic to the market, control jobs in high saturation clusters may receive more applications. Alternatively, if transparency draws applicants away from untreated jobs, control jobs in high saturation clusters may receive fewer applications. I find that control jobs receive similar volumes of male and female applications regardless of

the treatment intensity of competing jobs. In Table B.8, Columns 1–4, the coefficients on the high-saturation cluster indicator are small and statistically insignificant across all models. For example, the estimated effect on female applications is 1.31 (SE = 3.04) across all clusters in Column 1 and –0.93 (SE = 2.90) dropping the 3 highest traffic clusters in Column 3, ruling out meaningful crowding in or out of job-seekers.

The second test examines whether firms in the control group adjust their own disclosure behavior in response to competitors' treatment status. If transparency were contagious – for instance, if firms feared losing applicants to competitors that disclose salaries – control jobs in high-saturation clusters would be more likely to reveal pay information voluntarily. Column 5 of the same table tests this hypothesis by regressing a salary-visibility indicator on high-saturation. The coefficient is precisely estimated at zero (–0.00, SE = 0.02), showing that control firms are no more likely to disclose salaries when surrounded by treated competitors.

Finally, I test whether the magnitude of treatment effects varies with the treatment density of neighboring jobs. If transparency works partly through market-level mechanisms – for example, if job-seekers become more responsive to disclosed salaries when many postings in the market are treated – then the direct treatment effects would attenuate once I control for whether a given job is located in a high-saturation cluster. Table B.9 shows that direct treatment effects on applications remain large and statistically significant, particularly among large firms. For instance, without any controls, the interaction between "Treated Job" and "Large firm" is 17.67 ($p < 0.01$) for female applications in Column 1 and becomes 17.85 ($p < 0.01$) after accounting for treatment intensity in the market in Column 2. Similar patterns persist for male applications, indicating the observed treatment effects on applications reflect direct behavioral responses to salary disclosure, rather than indirect effects operating through local market exposure to transparency.

7.4 Family-friendly amenities

The average treatment effects show that salary disclosure attracts more women to large firms. What remains unclear is where in the amenity distribution those additional applications come from. Do women simply add large-firm jobs to an existing portfolio of flexible, family-friendly applications, or do they re-optimize their search? To answer this, I examine how the relative appeal of key non-wage amenities changes once wages are made visible. Specifically, within each experimental arm, I estimate how applications respond to core amenities and firm size in a Poisson specification, conditioning on other amenities and job characteristics.

Classifying amenities. I classify relevant 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 selected because they are both frequently mentioned in job ads, and have been

highlighted in prior research as attributes that women may value differentially, and that therefore may contribute to observed gender gaps in job search and employment.

Supply of amenities. Table 6 examines how the prevalence of these amenities varies by firm size, separately for control and treatment jobs. Consistent with the baseline facts reported in Section 3, 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 (45.5%) less likely than small firms to offer remote work and 2 pp (50%) less likely to offer flexible work hours (Columns 1 and 3). These gaps persist under treatment (Panel B). In contrast, transport assistance is less frequently offered by small and large firms, and 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). Thus, transport support, a safe work environment and equal-opportunity language are unlikely candidates for explaining women's lower application rates to large firms.

Demand for amenities: Specification. 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 the amenities described above, along with other job characteristics such as education and experience requirements, career level, and industry and occupation fixed effects.

Demand for amenities: Results. Table 7 reports the results from this exercise, estimating how specific job-level amenities and firm size predict the number of applications submitted by women and men, separately for control jobs with hidden salaries (Columns 1 to 2) and treatment jobs with disclosed salaries (Columns 3 to 4). Each coefficient reflects the incident rate ratio (IRR), where values above 1 indicate increases and below 1 indicate decreases in application volume. The final four columns formally test for differences in the coefficients. Columns 5 and 6 capture gender gaps within each salary regime, while Columns 7 and 8 report differences by salary disclosure separately for each gender.

When salaries are hidden, women's applications are strongly driven by amenities rather than firm size. Remote work (Row 2) is the dominant force: it increases applications by 167% for women and only 67% for men (Columns 1 and 2), generating a large and statistically significant gender gap of 99.5 pp (Column 5). Equal opportunity language (Row 6) also raises applications for both genders,

but with a 32.8 pp larger effect on female than male applications (Column 5; $p < 0.10$). By contrast, transport support and flexible work hours (Rows 3 and 4) do not significantly influence application volumes for either gender when salaries are concealed. Taken together, these patterns indicate that in the absence of wage information, women disproportionately sort into jobs that signal flexibility and inclusiveness.

Salary disclosure substantially weakens the role of these amenities in directing women's applications. Under transparency, the appeal of remote work among women declines sharply: its effect falls by 82.4 pp (Column 7; $p < 0.01$), while the corresponding decline for men is smaller and statistically insignificant (19.2 pp; Column 8). Other amenities similarly lose explanatory power. Transport support becomes negatively associated with applications for both genders, and flexible work hours remain insignificant while becoming more negative in magnitude.

As the influence of flexibility and inclusiveness weakens under transparency, firm size becomes more salient. Absent salary information, large firms have a null association with female applications (Row 1, Column 1) but attract 28% more male applications ($p < 0.01$; Row 1, Column 2). Once pay transparency is mandated, however, large firms attract 45.5 pp more female applications (Row 1, Column 7; $p < 0.05$), while the corresponding increase for men is smaller and statistically insignificant (13.6 pp; Row 1, Column 8).

Overall, these patterns point to a re-weighting of job attributes when wages become visible, with pronounced gendered effects. When salaries are hidden, women rely heavily on observable amenities – particularly remote work – to guide search. Salary disclosure deflates the appeal of these amenities and shifts applications toward large firms, even though those firms are not more amenity-rich. The resulting increase in female applications to large firms therefore reflects a decline in the relative importance of amenities under transparency, rather than a uniform expansion of applications across job types.

8 Lab experiment: Cross-randomizing amenities and salaries

A limitation of the preceding analysis is that amenity disclosure is not randomly assigned across jobs. Even with rich controls, unobserved job attributes may jointly influence both the likelihood that an amenity is advertised and applicants' responsiveness to that amenity, limiting a causal interpretation of amenity coefficients. To isolate preferences over non-wage amenities, I turn to a complementary discrete-choice lab experiment (DCE), in which salary information is cross-randomized with amenity disclosure within stylized job vignettes. In the DCE, 448 respondents evaluate 4 pairs of hypothetical job postings that vary independently in advertised salary ranges, remote-work availability, transport support, and equal-opportunity language, while holding constant core job characteristics such as its title, location, education requirements, and schedule. Salary ranges, when shown, are

randomized relative to each respondent's expected wage.[17] Figure A.6 presents illustrative examples of the vignettes.

The lab experiment has been discussed more broadly in Section 3, where it showed that while transport support and equal-opportunity language do not drive significant gender differences, women are significantly more likely than men to choose remote jobs and less likely to choose large-firm jobs (Table 1). The latter potentially reflects women's awareness that large firms offer fewer family-friendly amenities, as documented in Table B.4. Taken together with the results in the preceding section – where salary disclosure appears to weaken the appeal of remote work and shift women's applications toward large firms – the DCE provides a natural setting to test whether the appeal of remote work attenuates for female job-seekers once large firms reveal their salaries, and whether higher disclosed wage offers make large firms more attractive to women. This exercise also speaks directly to a central policy question of the paper: whether increasing women's applications to large firms requires them to expand the set of family-friendly amenities they offer, or whether salary transparency is sufficient.

Estimation. I test whether women pivot from flexibility to higher pay when large firms reveal salaries – and when the revealed salary levels 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 the advertised salary levels $s \in \{none, low, medium, high\}$; 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:

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

Standard errors are clustered at the respondent level in both specifications.

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

[17] The displayed range is generated by randomly assigning a minimum equal to {0.7, 0.75, 0.8} times, and a maximum equal to {1.1, 1.15, 1.2} times the respondent's previously reported expected salary.

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 large-firm salaries are hidden, women are 13.2 pp more likely to choose remote over on-site jobs. However, this preference attenuates once these salaries are disclosed, and it no longer significantly influences job choice at any wage level. In the highest salary tier, the appeal of remote work falls to 6.9 pp and is statistically insignificant. Conditional logit estimates in Appendix Table B.7 confirm this pattern: under pay non-disclosure, 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. By contrast, consistent with results from the field experiment, men's valuation of remote work is largely unchanged by salary disclosure. Across wage tiers and in both specifications, the probability of choosing a remote job over an on-site job is never statistically different from zero; if anything, at the highest salary tier men exhibit a modest (and imprecise) preference for on-site jobs.

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.

9 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 2) to test each of these alternative explanations.

9.1 Gender gaps in information or inference

In Pakistan, gender differences in labor market exposure are large. As a result, when salaries are hidden, men and women may rely on different information sets or form different inferences from the same non-wage information. Women, for instance, may systematically infer lower pay or tougher competition from the text of job ads, or may misjudge firm size based on firm names and job descriptions. Such differences in information or inference can translate to gender gaps in applications, even absent underlying differences in amenity preferences.

To measure potential gaps in information and inference, I present job-seekers with two recent job ads from the platform, and ask them to guess wages, competition, and firm size based on the text of the ad. Their responses are incentivized with financial rewards for accuracy, and benchmarked

against administrative data from the platform. 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 differences in knowledge about the broader labor market. I find that 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. Appendix Table B.10 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: 0.028 for small and large firms, compared to male means of 0.068 and 0.088, 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 are 0.063 and 0.137, respectively). Importantly, men and women fail to accurately guess salaries for both small and large firms. Overall, these results confirm that the disproportionate impact of pay transparency on women is not driven by a gendered lack of information about hidden salaries.

Competition. A second possibility is that men and women differ in their knowledge of competition for jobs, and thus, in their perceived likelihood of getting the job. 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 ask 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. Appendix Table B.11 reports the results. Columns (1)–(2) show the probability of guessing the number of female applications within 10% of the true value, while Columns (3)–(4) repeat the exercise for male applications. Across all specifications, there is no evidence of gender differences in beliefs about competition. For female applications, women's accuracy rates (coefficients = –0.015 for small firms and –0.005 for large firms) are statistically indistinguishable from male means of 0.05 and 0.04, respectively. Similarly, for male applications, women's accuracy rates (coefficients = 0.008 for small firms and –0.009 for large firms) are statistically similar to male means of 0.035 and 0.04, respectively. 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 small or 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 at 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). Appendix Table B.12 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 0.037 is statistically indistinguishable from a male mean of 0.347. For large firms (Column 2), the female coefficient is –0.017, 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.

9.2 Gender gaps in preferences for higher salaries

Pay transparency will have larger effects on women if they place greater value on salary information than men. Since large firms both offer higher pay and are more likely to hide it, transparency disproportionately changes perceptions of those jobs, narrowing the information gap that previously discouraged women from applying. I test this hypothesis using the discrete choice experiment embedded in my job-seekers' survey.

In the experiment, respondents make four binary choices between stylized job postings that vary randomly across five dimensions: salary visibility, large-firm label, remote-work option, transport subsidy, and diversity language. When salary is visible, the displayed range is generated by randomly assigning a minimum equal to {0.7, 0.75, 0.8} times, and a maximum equal to {1.1, 1.15, 1.2} times the respondent's previously reported expected salary. To identify gender differences in preferences over salary levels, I restrict the sample to choice sets with visible salaries and estimate gender-specific conditional logit models exploiting the randomized variation in the displayed salary range.

Let $Sal(j) \in \{low, medium, high\}$ denote the salary level of job j, with low as the omitted reference category. The probability that worker w selects job j from choice set i is given by

$$\Pr(\text{choice}_{wij} = 1 \mid \text{salary visible}) = \frac{\exp\left(\alpha_{Sal(j)} + \boldsymbol{\beta}'\mathbf{Z}_{wij}\right)}{\sum_{k \in C_i} \exp\left(\alpha_{Sal(k)} + \boldsymbol{\beta}'\mathbf{Z}_{wik}\right)},$$

where $\alpha_{Sal(j)}$ represents the utility weight associated with salary level $Sal(j)$ (with $\alpha_{low} = 0$), and $\mathbf{Z}_{wij}$ is the vector of other randomized job attributes (remote work, diversity language, transport

subsidy, and firm size) with corresponding coefficients β. Salary levels are defined by the midpoint of the displayed range: *low* (90–92.5% of expected salary), *medium* (95–97.5% of expected salary), and *high* (100% of expected salary).[18]

Appendix Table B.13 presents the results. The first two columns show odds ratios – the factor by which the odds of choosing a job increase when moving from low to medium, or low to high salary levels. Both men and women show a strong preference for higher pay, and this preference does not differ significantly by gender. Women are 2.17 times more likely to choose a medium-salary job over a low-salary alternative, and men are 2.48 times more likely. For high salaries, these odds ratios increase to 3.41 for women and 2.89 for men. The third column tests for gender differences in these utility weights by comparing the coefficients underlying Columns 1 and 2. To aid interpretation, the gender gap is reported in terms of marginal effects – the difference in choice probability (in percentage points) when salary levels rise. The fourth column reports p-values from tests of equality of the coefficients underlying Column 1 vs. Column 2. Column 4 shows that there are no statistically significant differences between men and women in their responsiveness to salary levels. For medium salaries, the difference in marginal effects is -1.67 percentage points ($p = 0.670$), and for high salaries it is 3.17 percentage points ($p = 0.756$), neither remotely approaching statistical significance. These results provide no support for the hypothesis that women place greater value on high salaries than men. The larger effects of pay transparency on female applications must therefore be explained by mechanisms other than differential preferences for salary levels.

9.3 Gender gaps in network size and quality

Section 9.1 established that when presented with job ads and asked to infer hidden salaries, competition levels, and firm size, men and women demonstrated similar inaccuracy. This suggests they brought comparable information sets into the survey. However, these incentivized guessing tasks measure the information job-seekers already possessed, not their ability to acquire new information when making higher-stakes application decisions outside the lab. If 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 have larger or higher quality networks, they may be better positioned to gather information about particular jobs, even if their baseline knowledge is similar to women's.

To test this, I examine gender differences in self-reported network characteristics. Appendix Table B.14 reports the results. 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

[18]For example, low salary arises from min=0.7× and max=1.1× or 1.15× expected salary, or min=0.75× and max=1.1×. Medium arises from min=0.75× and max=1.15× or 1.2×, or min=0.8× and max=1.1× or 1.15×. High corresponds to min=0.8× and max=1.2×.

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.

If this difference in network composition affects access to information, job-seekers with fewer employed contacts should rely more on publicly disclosed salaries and therefore apply disproportionately to treated jobs. Column (5) directly tests this hypothesis by matching survey data to platform application records and regressing the share of treated applications on the respondent's share of employed contacts, separately by gender. The results show no such relationship for either men or women. Thus, although men report 11.2% more employed contacts, this difference does not translate into differential responses to pay transparency. This finding is consistent with the earlier discussion of Figure 4, which shows that when salaries are hidden, men's application behavior is not responsive to actual wages – their application slope is flat rather than upward-sloping with respect to pay, suggesting they lack information about hidden salaries despite having a larger network of employed contacts. Together, these results indicate that gender differences in network composition cannot explain the disproportionate impacts of pay transparency on women.

9.4 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 a specific salary range. The workers' survey provides evidence against the prevalence of such perceptions. After viewing a large-firm job ad in their occupation that did not reveal the salary for the job, respondents were first asked to guess the minimum and maximum salary for the job. They were then shown two otherwise identical candidate profiles, one male and one female, that match the education, age, city, and occupation profile of the respondent (see Figure A.18 for an example). Respondents were asked how much each candidate would be offered for this job. Column (1) of Table B.15 is coded as 1 if their answers implied they expect a positive gender gap between male and female salary offers. The results reveal that a minority of both male and female job-seekers expect discrimination against women at large firms: 22.4% of men and an additional 5.4 pp of women believe a male candidate would receive a higher salary offer than an equally qualified female candidate. Column (2) examines whether women underestimate their own earnings potential relative to the perceived salary of the job. The ratio of respondents' expected personal offer to the maximum salary they guessed for the job is 0.908 for men and 0.994 for women, a difference of 8.6 pp that is statistically insignificant. Thus, if anything, women expect to be offered salaries closer to the perceived maximum of the job compared to men. Together, these results suggest that women's lower application rates to large firms with hidden salaries are not driven by substantial fears of discrimination.

9.5 Intra-household bargaining in the absence of salaries

Women may face intra-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 to working are ambiguous, making it harder to justify such jobs alongside household responsibilities. Appendix Table B.16 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. Among women, younger respondents seek permission from fathers, but as they get older they continue needing to seek permission – now from their from spouses. As a result, for women, the constraint that they require permission remains binding across their life-cycle, while for men, it relaxes with time. 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.17 examines whether this household constraint translates into differential responses to treatment in the field experiment. Column (1) re-states the gender gap in permission-seeking for reference. 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. To the contrary, it is highly likely that the gendered division of labor within households is the reason women prefer flexibility. This section merely suggests that permission-seeking or household-bargaining at the time of application do not interact with salary disclosure alone to drive gender differences in sorting. Rather, opinions of household members may be fully internalized by women when they trade-off wages, amenities and domestic responsibilities.

9.6 Gender differences in ambiguity and bargaining aversion

Do gender gaps in ambiguity or bargaining aversion explain gendered responses to salary transparency? Ambiguity aversion is a preference for risks with well-specified odds over risks with unknown odds. In job search, hidden pay can create ambiguity by making it difficult to assess how likely a given posting is to pay well. Disclosure of a salary or range supplies an informational anchor that reduces this uncertainty. Similarly, when wages are not posted, hiring may involve bargaining. Workers who find bargaining stressful, costly, or disadvantageous may avoid these jobs. If women differ from men in either preference, this could drive gendered responses to pay disclosure policies.

Measuring ambiguity aversion. I measure ambiguity aversion using a standard incentivized Ellsberg Paradox game following Ashraf et al. (2009). Respondents choose between two urns. 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 ambiguity aversion. The exact question presented to respondents is shown in Appendix Figure A.19.

Measuring bargaining aversion. 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. Respondents can only respond to one firm. Choosing the non-negotiable offer indicates a preference for avoiding bargaining.

Results. Appendix Table B.18 reports the results. Columns (1) and (2) show gender gaps in preferences. Women are no more likely than men to exhibit ambiguity aversion (Column 1). By contrast, they are substantially more likely to be bargaining averse (Column 2): the female coefficient of 18 pp – about 30% relative to the male mean – is large and statistically significant. If women's stronger response to pay transparency stems from their higher bargaining aversion, then bargaining-averse women should value disclosed salaries significantly more than women who are comfortable with negotiation. To test this, I use the discrete choice experiment from my job-seeker survey. Each respondent evaluates four pairs of job postings that vary randomly along five attributes: salary visibility, firm size, remote work, transport subsidy, and inclusive language. I estimate conditional logit models separately for bargaining-averse women and those who prefer bargaining to test whether their preference for salary visibility differs:

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

where $\mathbf{X}_{wij}$ is a vector of job attributes and $\boldsymbol{\alpha}$ captures preference weights. Columns (3) and (4) of Appendix Table B.18 present the results. Column (3) focuses on the salary-visibility attribute, comparing how strongly bargaining-averse women and those who prefer bargaining value transparent pay. Both groups exhibit a pronounced preference for jobs that disclose salaries: the odds of choosing a job with visible pay are 2.73 times higher for bargaining-averse women and 2.83 times higher for women who prefer bargaining. Column (4) translates these to marginal effects in percentage points: salary disclosure increases the probability of choosing a job by 18.1 pp for bargaining-averse women and 18.8 pp for women who prefer bargaining. Importantly, this difference of 0.7 pp is small and statistically insignificant (p=0.888). Thus, even women comfortable with negotiation place similar value on transparency, likely because disclosed salaries provide valuable information that informs their salary demands. This suggests that bargaining aversion cannot explain women's stronger preference for salary disclosure.

Implications. While ambiguity and bargaining aversion do not appear to drive the differential responses to transparency, the model can readily accommodate such preference differences – and doing so would amplify the disproportionate effects of disclosure policies on women. In the baseline model, expected utility from a job without wage disclosure is

$$V(w, a) = \mathbb{E}[w^{\theta}]\, a^{\gamma},$$

where θ governs preferences over wages and γ captures the valuation of non-wage amenities (a). Gender differences in bargaining aversion, risk aversion, or attitudes toward ambiguity can be incorporated by introducing a gender-specific parameter $\beta_f < 1$ that effectively increases concavity in the wage dimension under non-disclosure for women:

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

Lower values of β_f magnify the uncertainty penalty attached to hidden wages, making undisclosed pay less attractive to women. This mechanism would raise the gains from transparency for women relative to men. Whether driven by bargaining or ambiguity aversion, or other factors governing preferences over wage risk, such gender differences would strengthen the role of wage disclosure in reshaping women's sorting patterns toward high-paying, low-flexibility firms.

10 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 prefer to 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.

10.1 Why do large firms hide attractive salaries?

Advertising high salaries plausibly helps attract top talent. Yet large firms often withhold this information, implying that disclosure carries potential costs. To explore these costs, I field a baseline survey with roughly 300 employers. Among firms that do not disclose pay, 71% report that transparency could constrain wage-setting by anchoring candidates' expectations to posted salaries and limiting firms' flexibility to tailor pay to applicant ability.

In addition, 63% of firms expect transparency to increase the number of applicants, while 72% anticipate that average applicant quality will weakly decline (and 40% expect a strict decline). These responses reveal two underlying concerns among employers. First, that posting high salaries may attract more unsuitable candidates. Second, that screening applicants is costly – if it were not, firms would maximize application volumes and filter later. These concerns are not unfounded. In my setting, the average job receives 82 applications, and 70–80% of candidates meet the basic education and experience requirements (on paper). This reflects a broader challenge: as education levels rise globally, more 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 characteristics are weak predictors of ability (Bertrand et al., 2024; Bandiera et al., 2024), often forcing firms to rely on costlier processes, such as interviews and referrals to identify suitable candidates (Jayachandran, 2021; Chandrasekhar et al., 2020). 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.

A simple model presented in Appendix Section B formalizes these screening concerns. It shows how advertising high salaries can increase the number of applicants but reduce their expected quality, making wage non-disclosure a rational self-screening strategy for firms. The model also highlights a key trade-off: while transparency may lower average applicant quality, it unambiguously attracts higher-ability candidates. A pay transparency mandate could make the gains among top applicants more salient to firms, prompting them to adjust their disclosure behavior.

10.2 Are firms' screening concerns justified?

Firms expect transparency to attract more applicants. However, they appear to believe that lower average quality will raise screening costs that outweigh the gains potential match quality. This section assesses whether the data supports these beliefs.

Impacts on average applicant quality. Table 9 assesses how average applicant quality responds to the intervention. Panel A examines the share of applicants who meet experience requirements. The results are null for large firms, but for small firms the share declines by 2 pp among women and 1 pp among men. Given control-group means of 64–74%, this corresponds to a relative decline of 1.4–3.1% in experience match rates. Panel B reports effects on education match. Treatment leads to a 1 pp decline for both men and women, across small and large firms; relative to control means of 88–92%, this represents a reduction of about 1.1%. Panel C shows more variation in the skills match. Among applicants to large firms, treatment reduces the share meeting all skills 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 4.1% for women.

Taken together, these results point to modest declines in observed applicant quality, concentrated in skill-based mismatches at large firms. Similar patterns have been documented in other settings following pay transparency mandates (Skoda, 2022). Additionally, the consistently high control-group means for education and experience highlight the underlying screening challenge: most applicants appear formally qualified on paper, making it harder for firms to identify true match quality without costlier evaluation. The small observable declines may reflect broader reductions in unobserved applicant quality, lending some empirical support to firms' concerns that transparency raises the cost of screening.

Selection into the applicant pool by ability. Average applicant quality is informative about screening costs, since a lower average implies that firms must review more applications to identify suitable candidates. However, it reveals little about match quality, as firms ultimately hire the best rather than the average applicant. If transparency attracts stronger applicants, it could improve the quality of hires even if overall applicant quality declines. Although I do not observe the characteristics of hired workers, I test this idea by examining whether salary range disclosure changes the likelihood that high- or low-ability job-seekers apply to a given job. Table 10 reports results for all jobs in Panel A, while Panels B and C show results separately for large- and small-firm jobs.

I classify workers by their position in the market-wide distribution of résumé characteristics, defining high-ability applicants as those in the top quintile and low-ability applicants as those in the bottom quintile of an Anderson (2008) index summarizing four traits: education level, GPA, years of experience, and management experience. The outcome variable is the probability that at least one applicant of each type applies to a given job.

Panel A shows that large-firm jobs are already highly likely to attract at least one high-ability male applicant even without mandated salary disclosure. In the control group, this probability is 89.7% (Column 1). Transparency therefore has limited scope to further improve male selection: the treatment effect is only 1.6 pp. High-ability women, by contrast, present a starkly different picture. Only 56.8% of control-group jobs attract at least one high-ability female applicant (Column 2), which is 33 pp lower than the likelihood of attracting high-ability men. This gap leaves substantial room for improvement, and transparency delivers: the probability of attracting a high-ability woman rises by 6.5 pp – an 11.4% increase relative to control jobs, and more than four times the male effect. The gender difference in treatment effects is statistically significant ($p < 0.001$), indicating that transparency disproportionately draws high-ability women into the applicant pool.

A trade-off is that transparency also increases applications from low-ability workers, though these effects are smaller. For men, the treatment effect on low-ability applications is 1.7 pp (Column 3). For women, it is 3.3 pp (Column 4) – larger than for men but only half the 6.5 pp effect observed for high-ability women. This asymmetry suggests a favorable compositional shift: among women, transparency attracts high-ability applicants at roughly twice the rate of low-ability applicants.

Panels B and C confirm that these patterns hold for both large and small firms, though effects are stronger at large firms. For large firms (Panel B), transparency increases the probability of attracting at least one high-ability woman by 7.2 pp (11.5%) compared with 1.7 pp (1.8%) for men. For small firms (Panel C), the effects are 5.7 pp (10.8%) for women and 1.4 pp (1.6%) for men. In both subsamples, gender differences are highly significant for high-ability applicants ($p < 0.001$), and the compositional shift favoring high-ability over low-ability women remains evident.

Together, these results help explain why firms – especially large ones – may prefer to hide salaries. The modal applicant is male, and even without disclosure, firms already have a high likelihood of receiving high-ability male applications. Transparency therefore offers little upside on this margin while, in their view, it risks a surge of lower-ability candidates. The data, however, show that this concern about negative selection is largely misplaced: the composition of male applicants – high or low ability – changes little, while there are considerable gains in applications from high-ability women. Firms are correct to anticipate that transparency also brings in some lower-ability women, but these increases are modest relative to the increase in stronger female candidates.

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 on average 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 job-seeker characteristics: 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. The outcome variables are 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.

10.3 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 that large firms hide salaries by 2.5 pp. Thus, firms who had all their job posts treated are 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 women. 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.

11 Women's hiring probability: Back-of-the-envelope

While this paper primarily examines how pay transparency reshapes job search and sorting behavior, a natural downstream question is whether transparency affect women's likelihood of being hired into better-paid jobs. Because treatment is randomized at the *job* level while this outcome of interest operates at the *worker* level, I develop a back-of-the-envelope calculation that maps job-level experimental variation into worker-level hiring probabilities. To do so, I collect job-level hiring data through firm surveys, recording the number of male and female hires for each position.

Framework. Consider a labor market with N_g job-seekers of gender $g \in \{f, m\}$ and a set of job positions randomly assigned to treatment (salary disclosed) or control (salary disclosure is optional). For a representative worker of gender g, the probability of being hired by a representative large firm job can be decomposed into two components: the probability of applying to the job, and the probability of being hired conditional on applying.

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

Let A_g denote the number of applicants of gender g the job gets. Then the probability a representative worker applied is:

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

The representative worker's hiring probability conditional on applying is:

$$\Pr\big(\text{hired}_g \mid \text{applied}_g\big) = \frac{H_g}{A}.$$

where H_g is the number of applicants of gender g hired by the job, and A is the total number of applicants. Combining both terms, the unconditional probability that a representative man or woman is hired becomes:

$$\Pr(\text{hired}_g) = \frac{A_g}{N_g} \cdot \frac{H_g}{A}.$$

The treatment effect on hiring probability can be expressed as the ratio:

$$R_g = \frac{\Pr(\text{hired}_g \mid \text{treatment})}{\Pr(\text{hired}_g \mid \text{control})} = \frac{(A_g^T/N_g) \times (H_g^T/A^T)}{(A_g^C/N_g) \times (H_g^C/A^C)} \tag{6}$$

where T denotes the treatment group, and C the control group. Since the labor pool size N_g remains constant across treatment arms, this simplifies to:

$$R_g = \frac{(A_g^T/A^T) \times H_g^T}{(A_g^C/A^C) \times H_g^C} \tag{7}$$

Appendix Table B.19 reveals a crucial empirical finding: gender-specific hiring remains approximately constant across treatment and control. That is, firms hire roughly the same number of women per job whether or not salaries are disclosed ($H_f^T \approx H_f^C$), and similarly for men ($H_m^T \approx H_m^C$). This finding – that firms don't adjust their gender-specific hiring in response to changed applicant pool composition – simplifies the expression to:

$$R_g \approx \frac{A_g^T/A^T}{A_g^C/A^C} \tag{8}$$

This ratio captures how salary disclosure changes the rate at which gender-g workers appear in applicant pools relative to total competition.

Results. Using job-level data, I compute this ratio for both women and men. For women, the average number of female applicants per job is $A_f^C = 20.61$ in the control group and $A_f^T = 40.24$ in the treatment group, while total applicants per job increase from $A^C = 110.45$ to $A^T = 182.91$. This yields:

$$\frac{A_f^C}{A^C} = 0.1867, \quad \frac{A_f^T}{A^T} = 0.2200$$

Therefore: $R_f = 0.2200/0.1867 = 1.178$. A representative woman is 17.8% more likely to be hired by a large firm in treatment than control.

For men, the average number of male applicants per job is $A_m^C = 89.79$ in control and $A_m^T = 142.58$ in treatment, yielding:

$$\frac{A_m^C}{A^C} = 0.8131, \quad \frac{A_m^T}{A^T} = 0.7795$$

Therefore: $R_m = 0.7795/0.8131 = 0.959$. This implies that a representative man is 4.1% less likely to be hired under salary disclosure.

Mechanisms and interpretation. These results emerge entirely from compositional changes in applicant pools rather than changed firm hiring behavior. When salaries are disclosed, women increase their application rates to well-paid jobs by 95% (from 20.61 to 40.24 per job), while men increase theirs by 59% (from 89.79 to 142.58 per job). This differential response shifts the gender composition of applicants: women comprise 18.7% of applicants in control jobs but 22.0% in treatment jobs. Despite increased overall competition – total applications rise by 66% (from 110 to 183 per job). Thus, women's hiring probability increases because their presence in applicant pools grows more than men's. This back-of-the-envelope calculation suggests that salary transparency can improve women's hiring outcomes at large firms through the search channel. By correcting informational asymmetries, disclosure induces women to apply more frequently to well-paid positions, increasing their representation among applicants, and consequently their likelihood of being hired.

12 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 co-vary, 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.

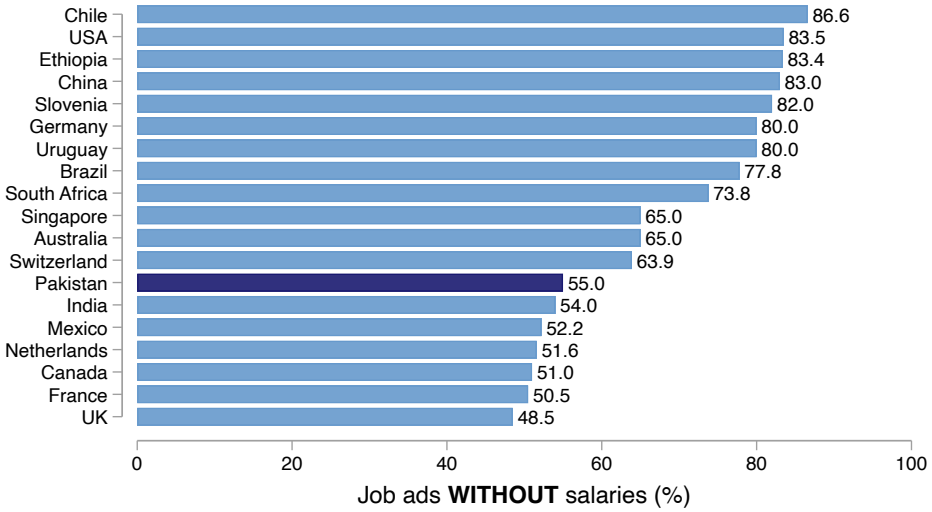
The findings carry clear policy implications. One approach to narrowing gender gaps in sorting is to alter the amenity mix offered by high-paying firms – for example, expanding flexible or remote work options. A cheaper and more immediately effective lever, however, is to improve pay information, which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

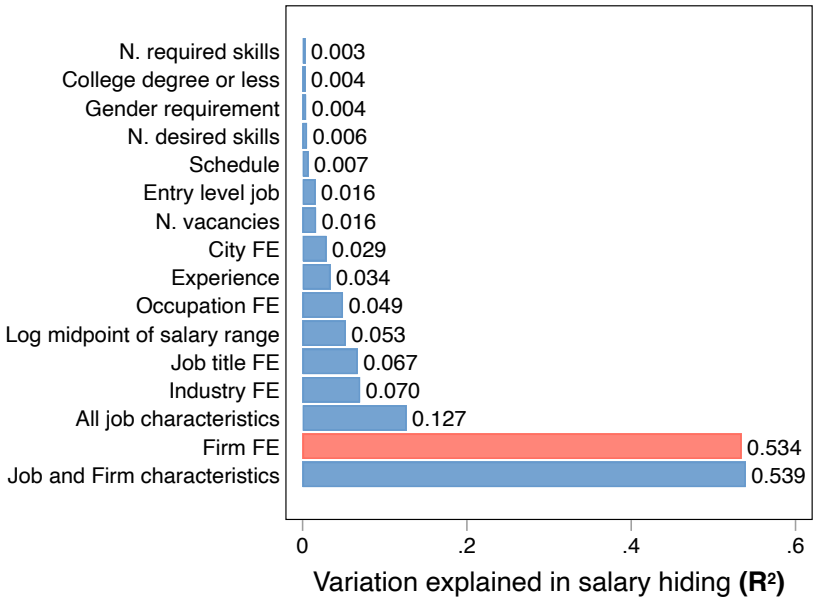
13 Figures

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); Escudero et al. (2024). [Return to page 8]

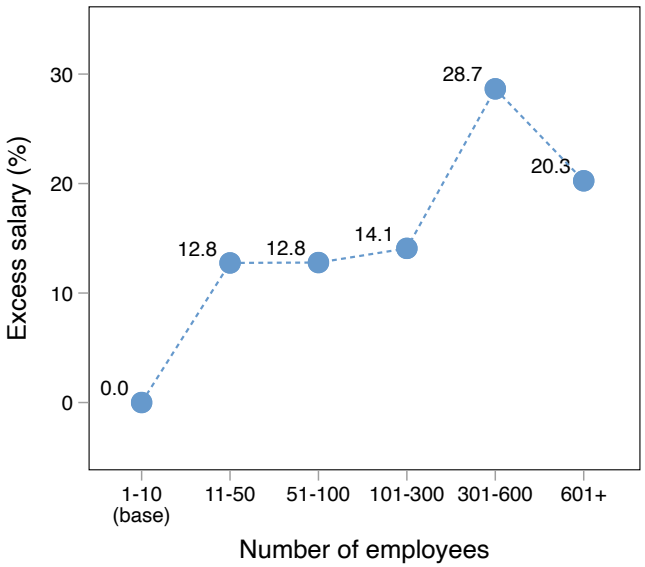
Figure 2: Explaining variation in salary non-disclosure



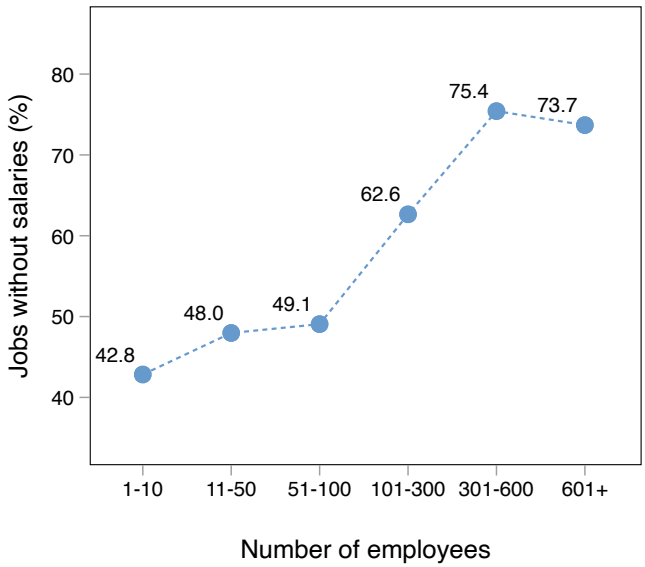
Notes: The figure shows R^2 from a job-level regression of an indicator for whether the salary of the job is hidden on each characteristic of the job, one at a time, until the third-last row in which all job characteristics are entered together. In the second to last row, firm fixed effects are controlled for by themselves, and the final row includes firm effects together with job characteristics. The sample consists of all baseline jobs. [Return to page 13]

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

(a) Firm-size pay premiums

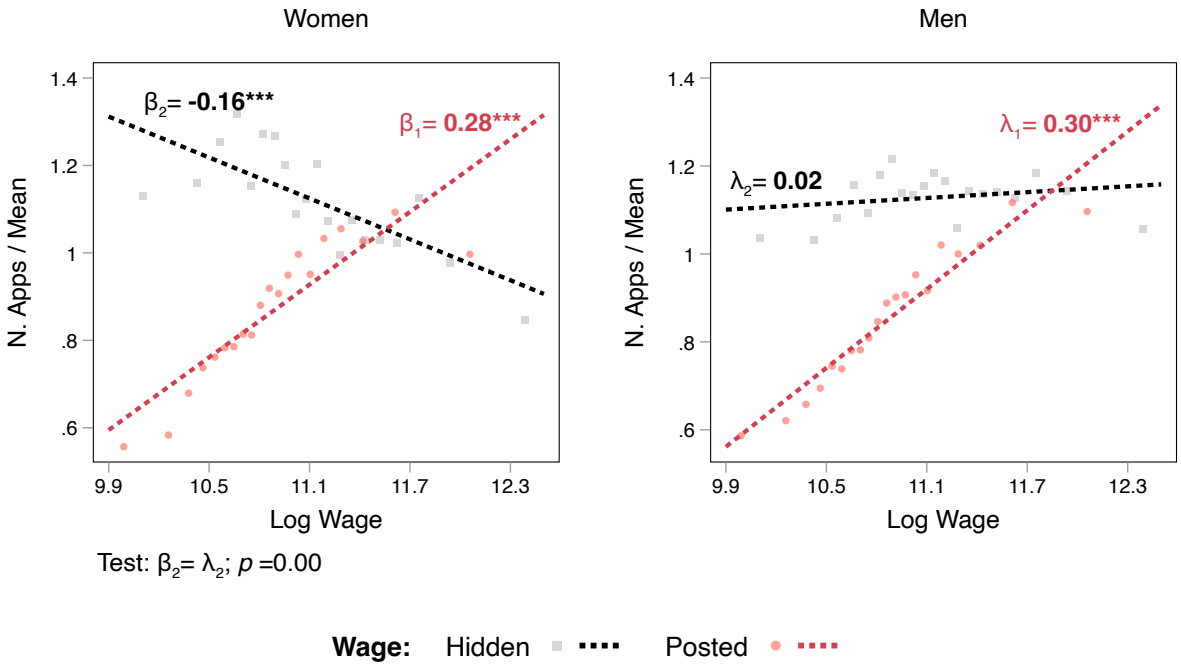


(b) Salary non-disclosure by firm size



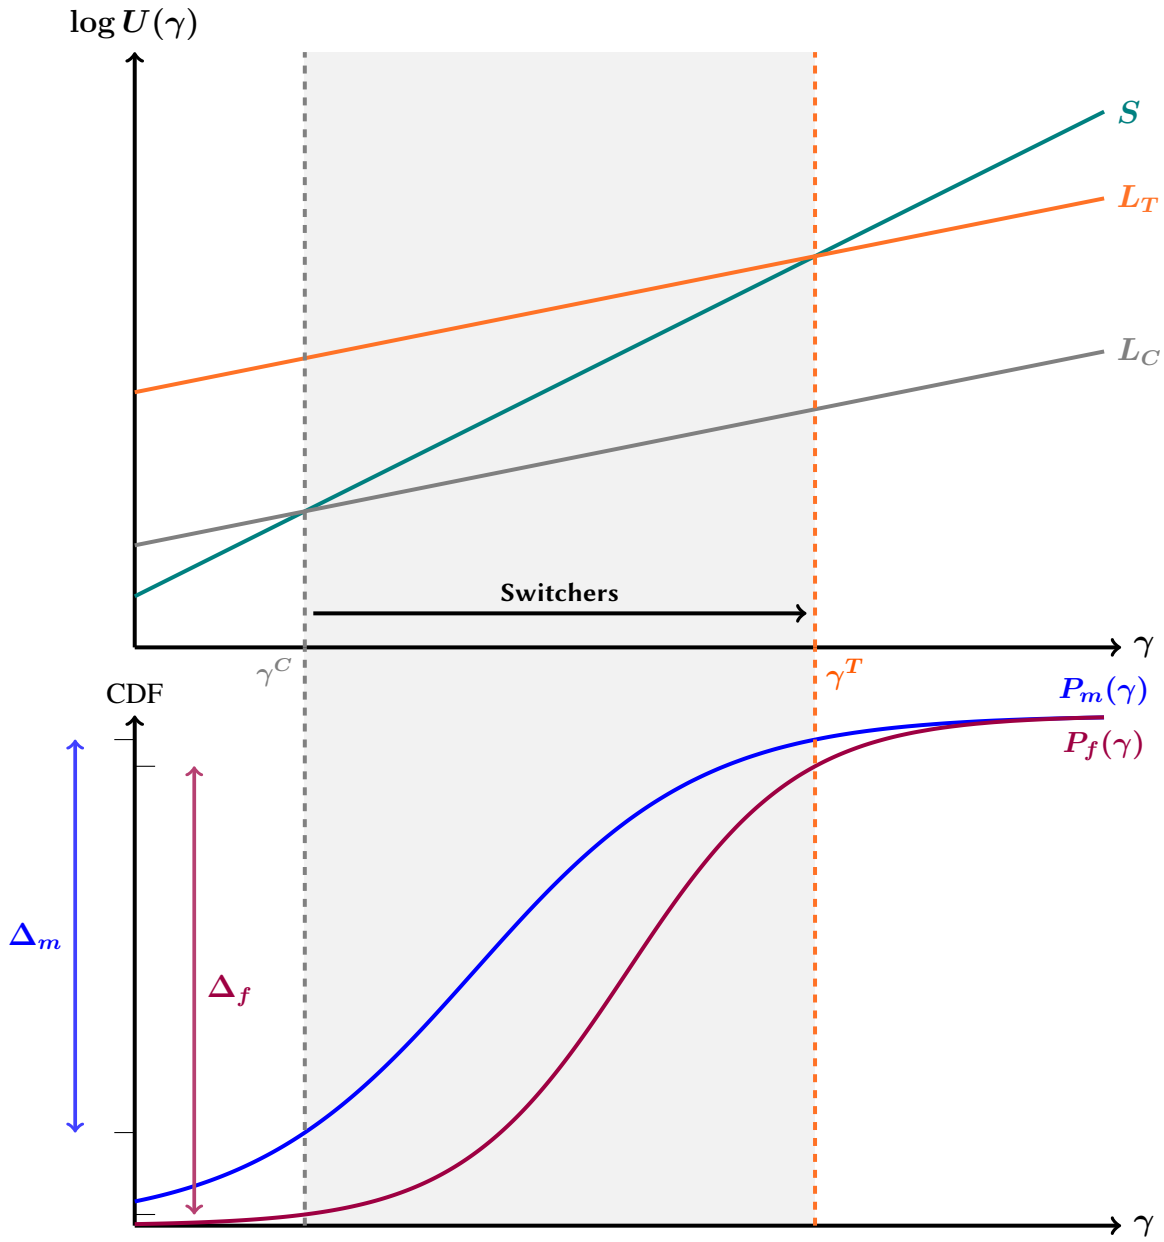
Notes: Panel A plots the relationship between firm size and salary premiums. Each point represents the coefficient from a regression of log maximum posted salary on firm size indicators and detailed job-level controls (SOC occupation codes, industry, city, experience and education requirements, career level, gender preference, year and month of posting, number of required and nice-to-have skills, job shift, and job type). The y-axis shows the percent salary premium relative to the smallest firm size category (1-10 employees), with robust standard errors. To hold sample size constant, missing values in continuous control variables are imputed with -9,999,999 and in categorical control variables with 0 and missing value indicators are included for all control variables. Panel B shows the percentage of job postings that hide salary information by firm size, estimated from a linear regression of a salary hiding indicator (multiplied by 100) on firm size indicators. Both panels use consolidated firm size categories: 1-10 (base), 11-50, 51-100, 101-300, 301-600, and 601+ employees. Point labels show the exact coefficient values. Sample includes job postings from 2019-2024, prior to the experiment. [Return to page 12]

Figure 4: Gendered application response by salary disclosure status



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 (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. 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). To hold sample size constant, missing values in continuous control variables are imputed with -9,999,999 and in categorical control variables with 0, and missing value indicators are included for all control variables. Sample includes job postings from 2019-2024, prior to the experiment. [Return to page 15 or 40.]

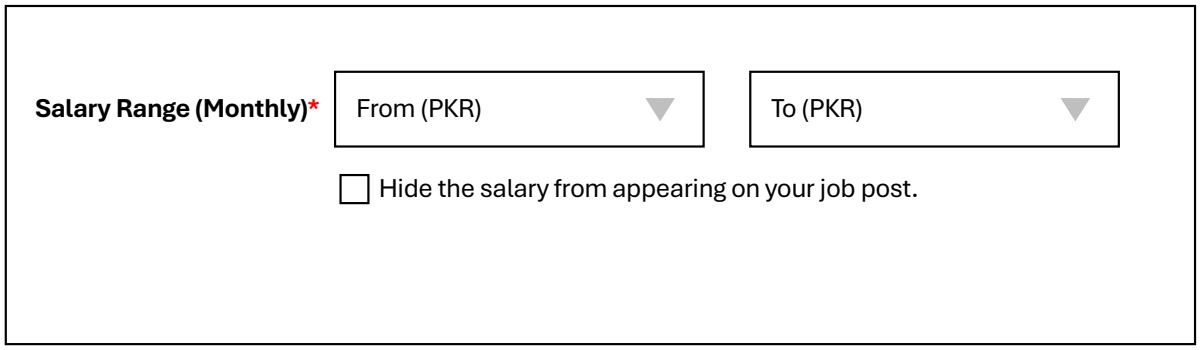
Figure 5: Theoretical framework



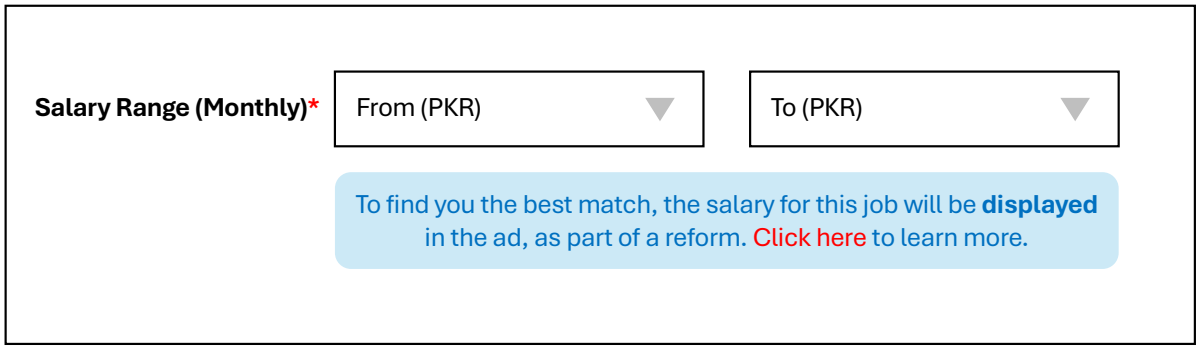
Notes: **Upper panel:** Utility functions by amenity preference γ. The blue line S represents small-firm utility. The gray line L_C shows large-firm utility under control (wage hidden), while the red line L_T shows large-firm utility under transparency (wage revealed). Transparency shifts the large-firm 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 **??** or 24]

Figure 6: Experiment Design

(a) Control Ads

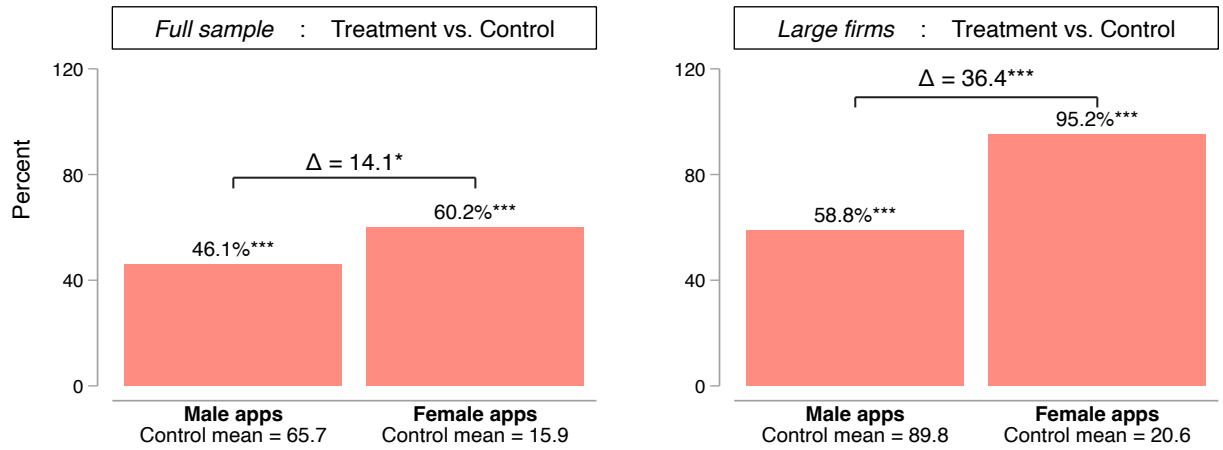


(b) Treatment Ads



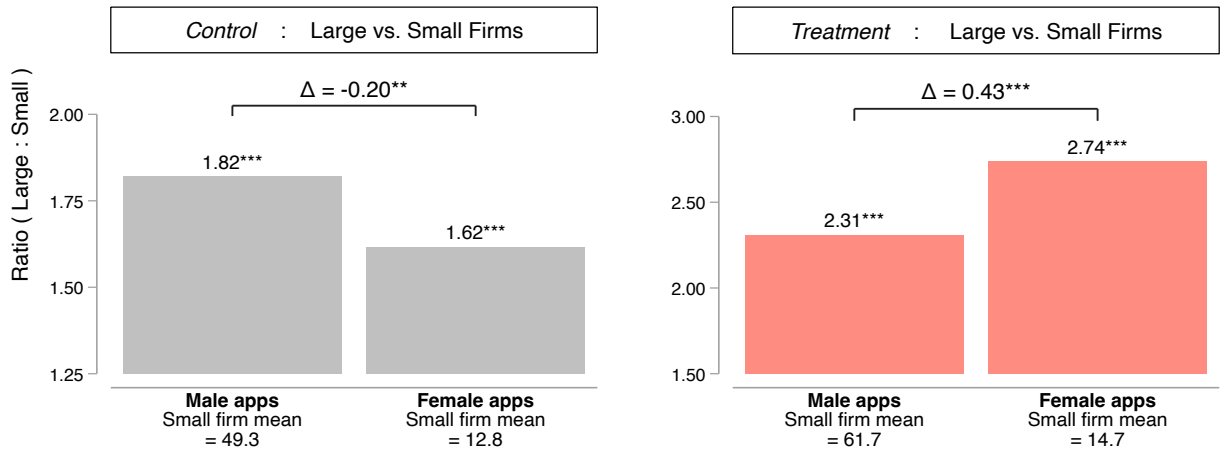
Notes: Panel (a) Control shows the standard job posting interface for specifying salary ranges. Employers can optionally check a box to hide the salary range from appearing in the job post. Panel (b) Treatment shows the modified interface for the treatment group, where salary ranges are displayed by 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 25]

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 30]

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 30]

14 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 15, 18 or 35]

Table 2: Treatment effects on salaries

	(1) **Log Max Salary**	(2) **Log Min Salary**	(3) **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 29]

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 Equation 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.20. For similar graphs on gender specific results in Columns 2-3 and Columns 5-6, see Figure A.21. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. [Return to page 29]

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.22. Robust standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. [Return to page 30]

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 30]

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	0.11	0.02	0.04	0.02	0.02
Observations	8,742	8,742	8,742	8,742	8,742
Panel B: Treatment 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***	0.02***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Small firm	0.11	0.01	0.04	0.02	0.02
Observations	10,972	10,972	10,972	10,972	10,972

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 33]

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

	Control: Salary Hidden (IRR)		Treatment (IRR)		Diff: Female - Male (pp)		Diff: Treatment - Control (pp)	
	Female apps (1)	Male apps (2)	Female apps (3)	Male apps (4)	Control (1)-(2) (5)	Treatment (3)-(4) (6)	Female (3)-(1) (7)	Male (4)-(2) (8)
Large firm	1.13	1.28***	1.57***	1.40***	-14.80	16.21**	43.54**	12.53
	(0.11)	(0.06)	(0.11)	(0.08)				
Remote work	2.57***	1.65***	1.84***	1.47***	92.21***	36.28***	-73.53**	-17.60
	(0.11)	(0.11)	(0.12)	(0.09)				
Transport	1.15	1.00	0.55**	0.73**	15.51	-17.84**	-60.05**	-26.70
	(0.21)	(0.21)	(0.23)	(0.15)				
Flex work	0.98	0.88	0.59*	0.73*	10.71	-14.05	-38.96	-14.21
	(0.18)	(0.14)	(0.28)	(0.16)				
Safe environment	1.39	1.41*	1.03	1.01	-2.10	1.37	-36.10	-39.57
	(0.23)	(0.19)	(0.17)	(0.12)				
Equal opp. employer	1.80**	1.39**	1.27	1.24	40.55*	2.88	-52.71	-15.04
	(0.23)	(0.16)	(0.24)	(0.16)				
Dependent var. mean	18.31	72.07	25.52	95.98				
Observations	3,851	3,851	10,972	10,972				

Notes: The table reports Poisson regression estimates of how firm size and job-level amenities (remote work, transport support, flexible work hours, references to safe work environment, and equal opportunity language) 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 where salary transparency is mandated. Coefficients in columns 1-4 are incident rate ratios (IRR). For reference, an estimate of 1.13 implies a 13% increase while 0.93 implies a 7% decrease. The difference columns report percentage point (pp) differences in IRR: columns (1)-(2) and (3)-(4) show Female minus Male differences within control and treatment groups, respectively; columns (3)-(1) and (4)-(2) show Treatment minus Control differences for female and male applicants, respectively. In addition to the rows displayed in the table, all specifications control for industry, occupation, education requirements, gender preferences, career level, experience requirements, and work shift. Stars on coefficients in columns 1-4 indicate statistical significance of the amenity effect; stars in difference columns indicate statistical significance of the difference between the two coefficients being compared. $*p < 0.1, **p < 0.05, ***p < 0.01$. [Return to page 33]

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***
	(0.028)	(0.038)	
Low salary	0.512***	0.526***	0.014
	(0.054)	(0.064)	
Medium salary	0.647***	0.694***	0.048
	(0.045)	(0.062)	
High salary	0.636***	0.705***	0.069
	(0.050)	(0.053)	
Observations			1042

Panel B: Men

	On-site	Remote	Difference
No salary	0.491***	0.473***	-0.018
	(0.036)	(0.041)	
Low salary	0.524***	0.532***	0.008
	(0.059)	(0.070)	
Medium salary	0.525***	0.625***	0.099
	(0.056)	(0.061)	
High salary	0.794***	0.672***	-0.122
	(0.045)	(0.068)	
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 35]

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 44]

Table 10: Selection into application pools

Panel A: All firms

	High Ability		Low Ability	
	(1) **Men**	(2) **Women**	(3) **Men**	(4) **Women**
Treated Job	0.016***	0.065***	0.017***	0.033***
	(0.004)	(0.007)	(0.006)	(0.007)
Control mean	0.897	0.568	0.721	0.497
Observations	20,088	20,088	20,088	20,088
H_0: $\beta_f = \beta_m$; p-value		0.000		0.051

Panel B: Large firms

	High Ability		Low Ability	
	(1) **Men**	(2) **Women**	(3) **Men**	(4) **Women**
Treated Job	0.017***	0.072***	0.008	0.028**
	(0.006)	(0.010)	(0.010)	(0.011)
Control mean	0.924	0.627	0.677	0.476
Observations	8,341	8,341	8,341	8,341
H_0: $\beta_f = \beta_m$; p-value		0.000		0.120

Panel C: Small firms

	High Ability		Low Ability	
	(1) **Men**	(2) **Women**	(3) **Men**	(4) **Women**
Treated Job	0.014**	0.057***	0.026***	0.038***
	(0.006)	(0.009)	(0.008)	(0.009)
Control mean	0.879	0.528	0.751	0.512
Observations	11,747	11,747	11,747	11,747
H_0: $\beta_f = \beta_m$; p-value		0.000		0.265

Notes: The table shows the effect of treatment on the probability that at least one high or low ability applicant applies to a job. I classify job-seekers in the market as high or low ability based on their position in the market-wide distribution of a composite résumé index. High-ability applicants are those in the top quantile of this distribution; low-ability applicants are those in the bottom quantile. The composite index is an Anderson index across four traits: education level, GPA, years of experience, and management experience. Panel A reports results for all firms; Panel B reports results for large firms; Panel C reports results for small firms. The Control mean row shows the probability in control jobs. H_0: $\beta_f = \beta_m$ tests whether the treatment effect differs between men and women. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 45]

Table 11: Characteristics of the top applicants

Panel A: All firms						
	(1) **Expected Salary (SDs)**	(2) **Education Level (SDs)**	(3) **Experience (SDs)**	(4) **GPA (SDs)**	(5) **Team Size Managed (SDs)**	(6) **At least 1 Female**
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 46]

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 46]

References

Adzuna (2022). Adzuna salary transparency white paper. Accessed: 2024-12-23. [Return to page **51**]

AfriWorket (2024). Afriworket: Empowering africa's workforce. Accessed: 2024-12-23. [Return to page **51**]

Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American statistical Association*, 103(484):1481–1495. [Return to page **45**]

Arnold, D., Quach, S., and Taska, B. (2022). The impact of pay transparency in job postings on the labor market. *SSRN Electronic Journal*. [Referenced on pages **6, 26, and 29**]

Ashraf, N., Bandiera, O., Davenport, E., and Lee, S. S. (2020). Losing prosociality in the quest for talent? sorting, selection, and productivity in the delivery of public services. *American Economic Review*, 110(5):1355–1394. [Return to page **110**]

Ashraf, N., Giné, X., and Karlan, D. (2009). Finding missing markets (and a disturbing epilogue): Evidence from an export crop adoption and marketing intervention in kenya. *American Journal of Agricultural Economics*, 91(4):973–990. [Referenced on pages **11, 41, and 108**]

Bandiera, O., Jalal, A., and Roussille, N. (2025). The illusion of time: Gender gaps in job search and employment. Working Paper 34051, National Bureau of Economic Research. [Return to page **7**]

Bandiera, O., Kotia, A., Lindenlaub, I., Moser, C., and Prat, A. (2024). Meritocracy across countries. 82 pages. Posted: 18 Apr 2024. Multiple versions available. [Return to page **44**]

Banfi, S. and Villena-Roldán, B. (2019). Do high-wage jobs attract more applicants? directed search evidence from the online labor market. *Journal of Labor Economics*, 37(3):715–746. [Referenced on pages **6, 13, and 51**]

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

Batra, H., Michaud, A., and Mongey, S. (2023). Online job posts contain very little wage information. Technical Report 31984, National Bureau of Economic Research. [Referenced on pages **6, 13, and 51**]

Benson, A., Li, D., and Shue, K. (2021). Potential" and the gender promotion gap. *Unpublished working paper*. [Return to page **6**]

Bertrand, M., Crépon, B., and Caria, S. (2024). Improving job seekers' employment and earnings through credible skills signals. Accessed: 2024-12-23. [Return to page **44**]

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

Blau, F. D. and Kahn, L. M. (2017). The gender wage gap: Extent, trends, and sources. *Journal of Economic Literature*, 55(3):789–865. [Referenced on pages **6 and 13**]

Brenčič, V. (2012). Wage posting: evidence from job ads. *Canadian Journal of Economics/Revue canadienne d'économique*, 45(4):1529–1559. [Return to page **51**]

Breza, E., Kaur, S., and Shamdasani, Y. (2021). Labor rationing. *American Economic Review*, 111(10):3184–3224. [Return to page **7**]

Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(1):S13–S70. [Return to page **11**]

Card, D., Cardoso, A. R., and Kline, P. (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly Journal of Economics*, 131(2):633–686. [Referenced on pages **1 and 5**]

Card, D., Heining, J., and Kline, P. (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics*, 128(3):967–1015. [Return to page **11**]

Card, D., Mas, A., Moretti, E., and Saez, E. (2012). Inequality at work: The effect of peer salaries on job satisfaction. *American Economic Review*, 102(6):2981–3003. [Return to page **5**]

Chandrasekhar, A. G., Morten, M., and Peter, A. (2020). Network-based hiring: Local benefits; global costs. Technical report, National Bureau of Economic Research. [Return to page **44**]

Colonnelli, E., Tåg, J., Webb, M., and Wolter, S. (2018). A cross-country comparison of dynamics in the large firm wage premium. *AEA Papers and Proceedings*, 108:323–327. [Return to page **11**]

Cullen, Z. and Perez-Truglia, R. (2022). How much does your boss make? the effects of salary comparisons. *Journal of Political Economy*, 130(3):766–822. [Return to page 5]

Cullen, Z. and Perez-Truglia, R. (2023). The old boys' club: Schmoozing and the gender gap. *American Economic Review*, 113(7):1703–1740. [Return to page 13]

Cullen, Z. B. (2024). Is pay transparency good? *Journal of Economic Perspectives*, 38(1):153–180. [Return to page 9]

Cullen, Z. B. and Pakzad-Hurson, B. (2023). Equilibrium effects of pay transparency. *Econometrica*, 91(3):765–802. [Return to page 6]

DeLeire, T. and Levy, H. (2004). Worker sorting and the risk of death on the job. *Journal of Labor Economics*, 22(4):925–954. [Return to page 23]

Escudero, V., Liepmann, H., and Vergara, D. (2024). Directed search, wages, and non-wage amenities: Evidence from an online job board. Technical Report IZA Discussion Paper No. 17211, IZA — Institute of Labor Economics. Also issued as ILO Working Paper No. 136 (2025). [Return to page 51]

Frimmel, W., Schmidpeter, B., Wiesinger, R., and Winter-Ebmer, R. (2022). Mandatory wage posting, bargaining and the gender wage gap. Working Paper 2202, Linz. [Return to page 6]

Garlick, R., Field, E., and Vyborny, K. (2025). Women's mobility and labor supply: Experimental evidence from pakistan. [Return to page 6]

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119. [Referenced on pages 1, 6, and 13]

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

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

Hansen, B. and McNichols, D. (2020). Information and the persistence of the gender wage gap: early evidence from california's salary history ban. Technical report, National Bureau of Economic Research. [Return to page 6]

Ho, L., Jalota, S., and Karandikar, A. (2024). Bringing work home: Flexible work arrangements as gateway jobs for women in west bengal. Technical report, Technical Report. [Return to page **6**]

Hotz, V. J., Johansson, P., and Karimi, A. (2017). Parenthood, family friendly workplaces, and the gender gaps in early work careers. NBER Working Paper 24173, National Bureau of Economic Research, Cambridge, MA. Revised June 2018; previously circulated as "Parenthood, Family Friendly Firms, and the Gender Gaps in Early Work Careers". [Referenced on pages **13 and 23**]

HRD Asia (2023). Pay transparency: More singapore job postings including salary offers. Accessed: 2024-04-15. [Return to page **51**]

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

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

Indeed Hiring Lab (2023a). Pay transparency: More singapore job postings including salary offers. Accessed: 2024-04-15. [Return to page **51**]

Indeed Hiring Lab (2023b). Salaires dans les offres d'emploi : l'importance croissante de la transparence en europe. Accessed: 2024-04-15. [Return to page **51**]

Indeed Hiring Lab (2024). Salary transparency increasingly the norm in canadian job postings. Accessed: 2024-04-15. [Return to page **51**]

Jäger, S., Roth, C., Roussille, N., and Schoefer, B. (2024). Worker beliefs about outside options. *The Quarterly Journal of Economics*, 139(3):1505–1556. [Return to page **5**]

Jayachandran, S. (2021). Microentrepreneurship in developing countries. *Handbook of labor, human resources and population economics*, pages 1–31. [Return to page **44**]

Kleven, H., Landais, C., and Søgaard, J. E. (2019). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4):181–209. [Return to page **13**]

Kuhn, P. and Shen, K. (2013). Gender discrimination in job ads: Evidence from china. *The Quarterly Journal of Economics*, 128(1):287–336. [Referenced on pages **6 and 51**]

Lavetti, K. (2023). Compensating wage differentials in labor markets: Empirical challenges and applications. *Journal of Economic Perspectives*, 37(3):189–212. [Referenced on pages **5 and 18**]

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

Lucas, R. E. (1978). On the size distribution of business firms. *Bell Journal of Economics*, 9(2):508–523. [Referenced on pages **6 and 11**]

Marinescu, I. and Wolthoff, R. (2020). Opening the black box of the matching function: The power of words. *Journal of Labor Economics*, 38(2):535–568. [Return to page **78**]

Mas, A. and Pallais, A. (2017). Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–3759. [Referenced on pages **1, 5, 14, 18, and 24**]

Mas, A. and Pallais, A. (2020). Alternative work arrangements. *Annual Review of Economics*, 12(1):631–658. [Referenced on pages **5 and 24**]

Michelacci, C. and Suarez, J. (2006). Incomplete wage posting. *Journal of Political Economy*, 114(6):1098–1123. [Return to page **44**]

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

Morchio, I. and Moser, C. (2021). The gender pay gap: Micro sources and macro consequences. CEPR Discussion Paper DP16383, Centre for Economic Policy Research. [Return to page **1**]

Muralidharan, K. and Singh, A. (2023). India's new national education policy: Evidence and challenges. Accessed via RISE. [Return to page **44**]

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

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

Oi, W. Y. and Idson, T. L. (1999). Firm size and wages. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 3, chapter 33, pages 2165–2214. Elsevier, Amsterdam. [Return to page **11**]

People Matters (2023). Pay transparency on the rise in india, 50salary details. Accessed: 2024-04-15. [Return to page **51**]

Pertold-Gebicka, B., Pertold, F., and Gupta, N. D. (2016). Employment adjustments around childbirth. IZA Discussion Paper 9685, IZA Institute of Labor Economics. [Return to page **13**]

Roussille, N. (2023). The role of the ask gap in gender pay inequality. *Working Paper*. [Return to page **6**]

Skoda, S. (2022). Directing job search in practice: Mandating pay information in job ads. *Job Market Paper*. [Referenced on pages **4, 6, 29, 45, and 51**]

Song, J., Price, D. J., Guvenen, F., Bloom, N., and Von Wachter, T. (2019). Firming up inequality. *The Quarterly journal of economics*, 134(1):1–50. [Referenced on pages **5 and 19**]

Sorkin, I. (2017). The role of firms in gender earnings inequality: Evidence from the united states. *American Economic Review: Papers & Proceedings*, 107(5):384–387. [Referenced on pages **1 and 13**]

Tô, M., Coudin, É., Maillard, S., and Tréguier, J. (2025). Sorting into lower-paying firms: The firm premium child penalty in the french labour market. *Journal of Human Resources*. Revise & Resubmit. [Return to page **13**]

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

Vattuone, G. (2023). Worker sorting and the gender wage gap. Job Market Paper. [Referenced on pages **13 and 23**]

Wiswall, M. and Zafar, B. (2018). Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics*, 133(1):457–507. [Referenced on pages **5 and 24**]

World Bank (2011). *World Development Report 2012: Gender Equality and Development*. World Bank, Washington, DC. Public Disclosure Authorized. [Referenced on pages **7 and 13**]

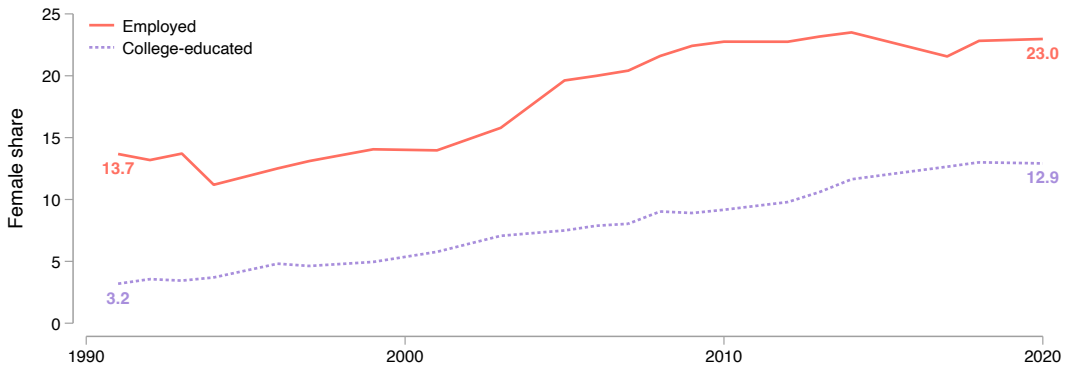
World Economic Forum (2025). Global gender gap report 2025. Technical report, World Economic Forum. 19th annual Global Gender Gap Index benchmarking 148 economies. [Return to page **7**]

A Appendix

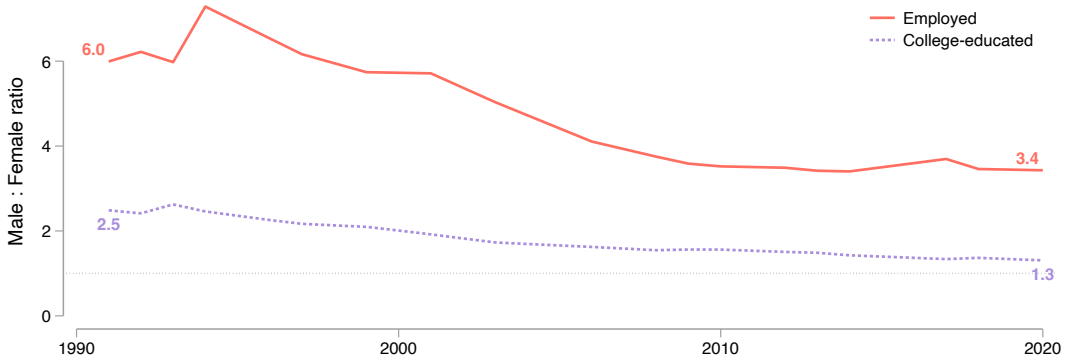
A.1 Appendix Figures

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

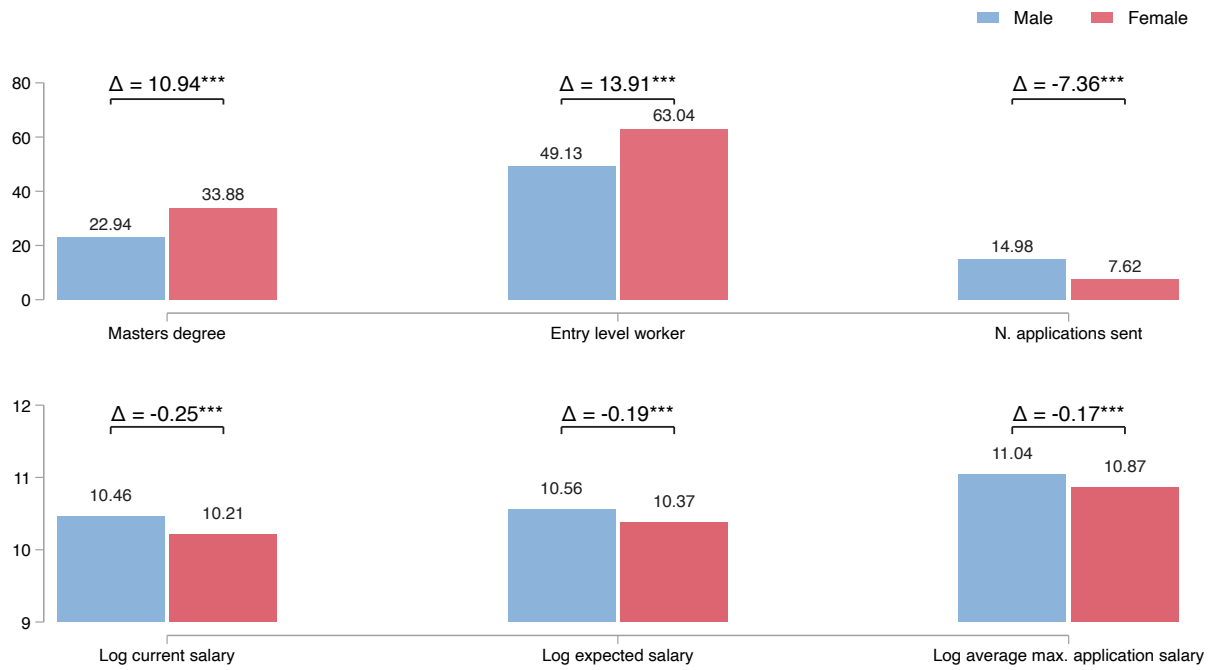


(c) Gender gaps in large-firm and salaried employment



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 10]

Figure A.3: First Word of Job Title

(a) Jobs with Hidden Salaries

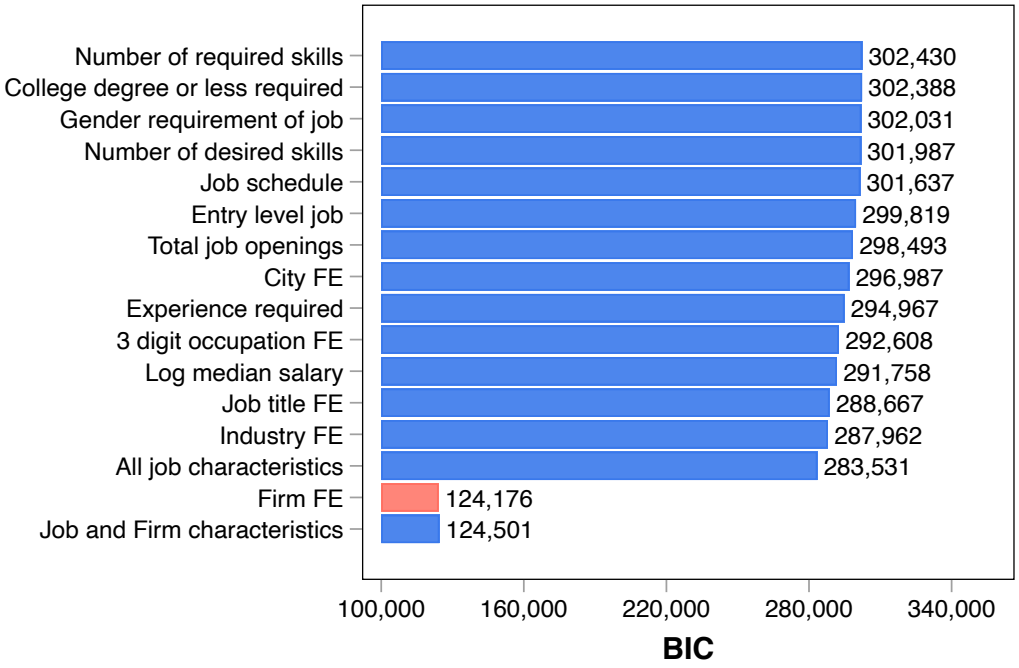


(b) Jobs with Visible Salaries



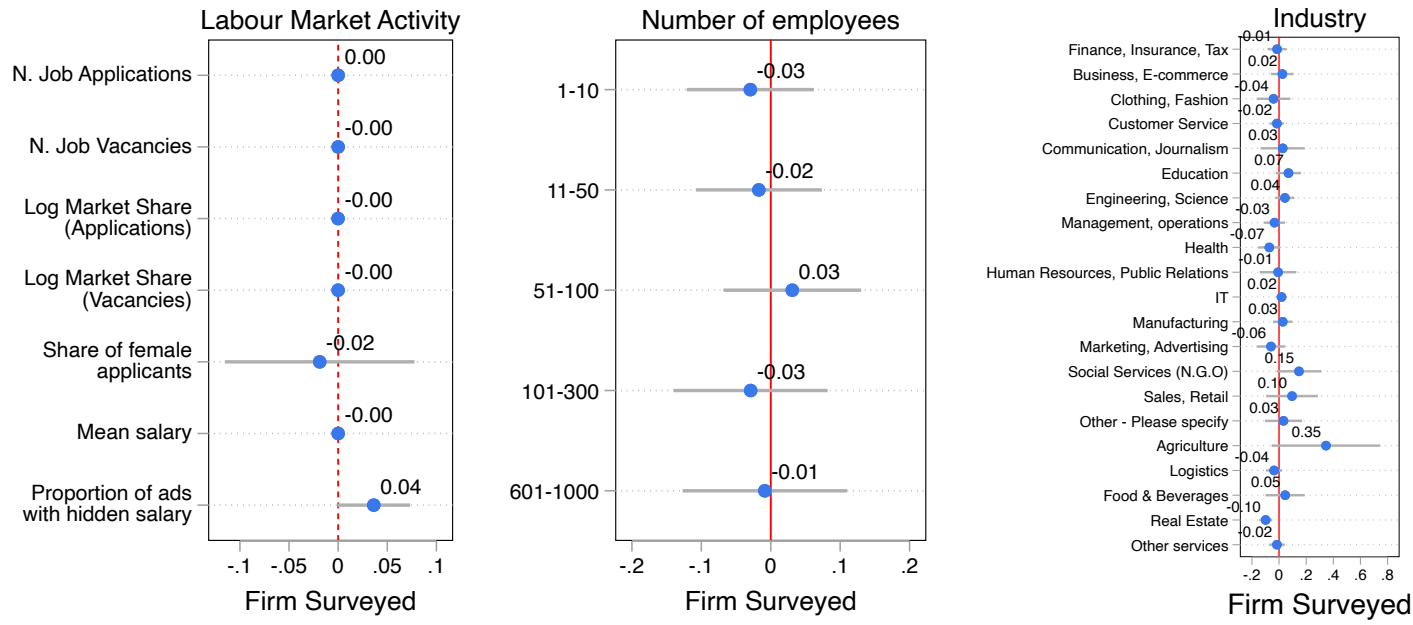
Notes: The figure shows the forst word of job titles separately for jobs with hidden salaries (panel a) and jobs with visible salaries (panel b). The job titles are cleaned following the approach in Marinescu and Wolthoff (2020). [Return to page 8]

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 13]

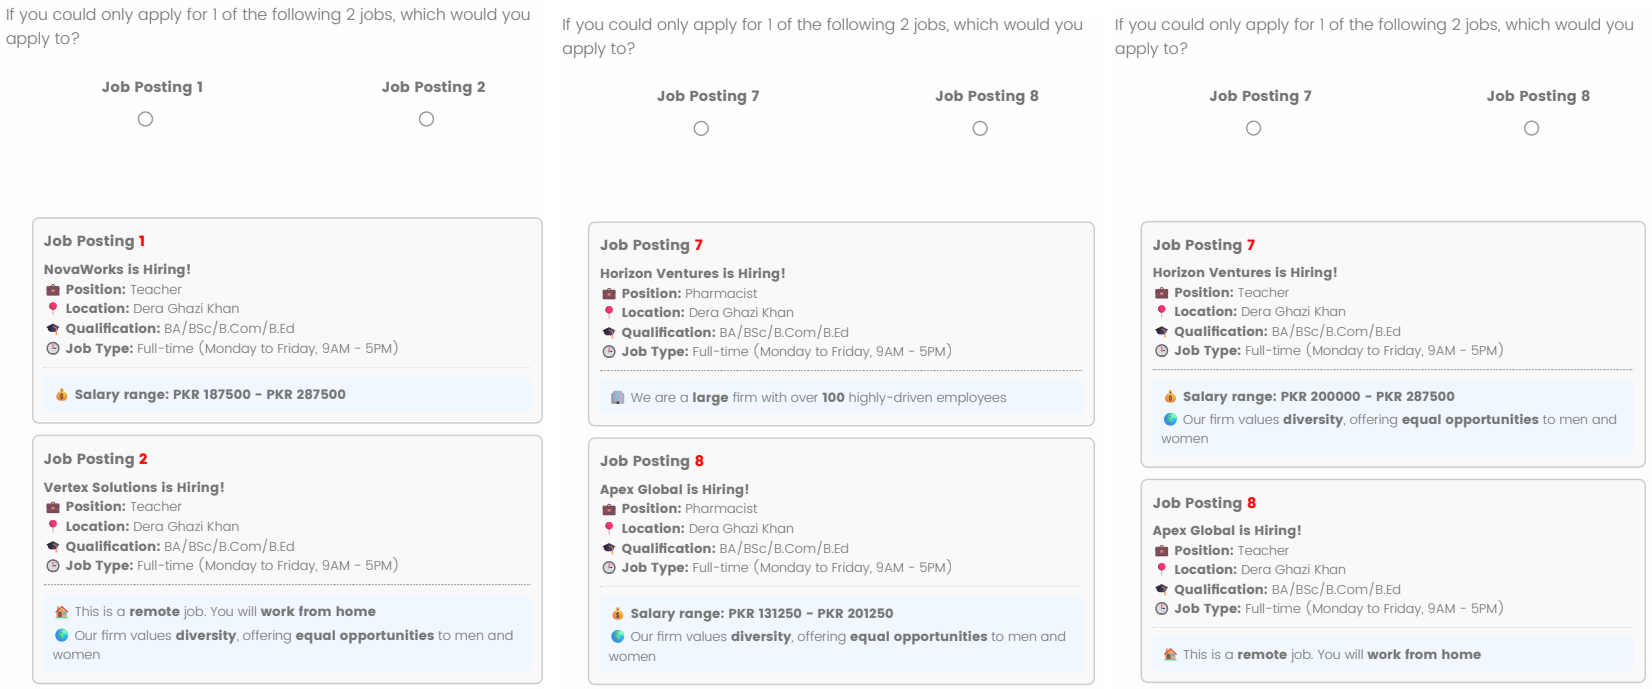
Figure A.5: Selection into firm survey



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

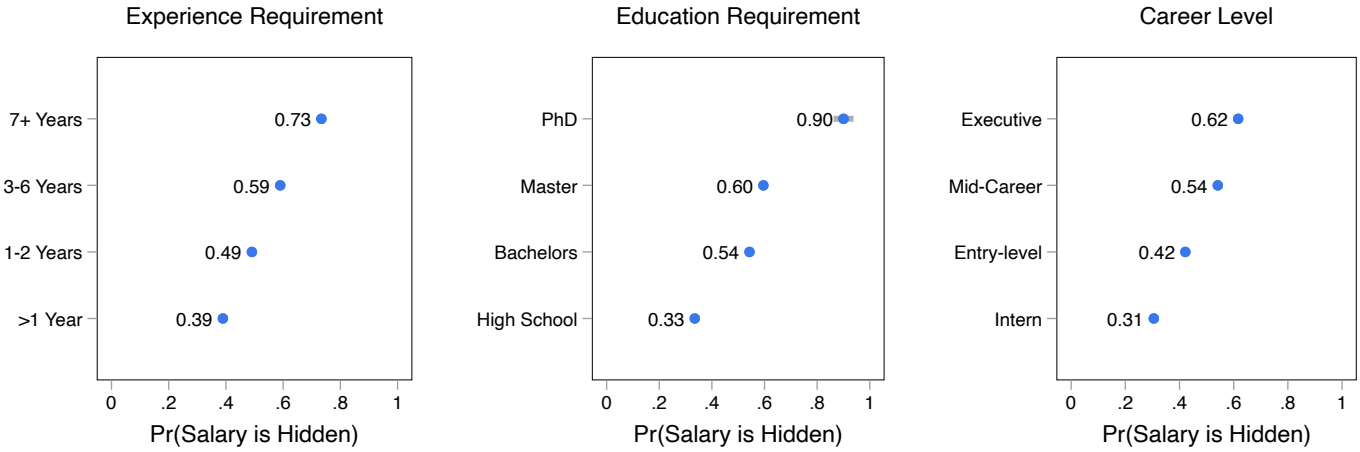
Notes: The figure tests whether there was selection in the firms who consented to participate in the survey. The first subfigure tests whether there is differential selection into the survey by the number of job applications received by the firm, the number of job vacancies, the log market share of applications 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: Examples of vignettes from discrete choice experiment



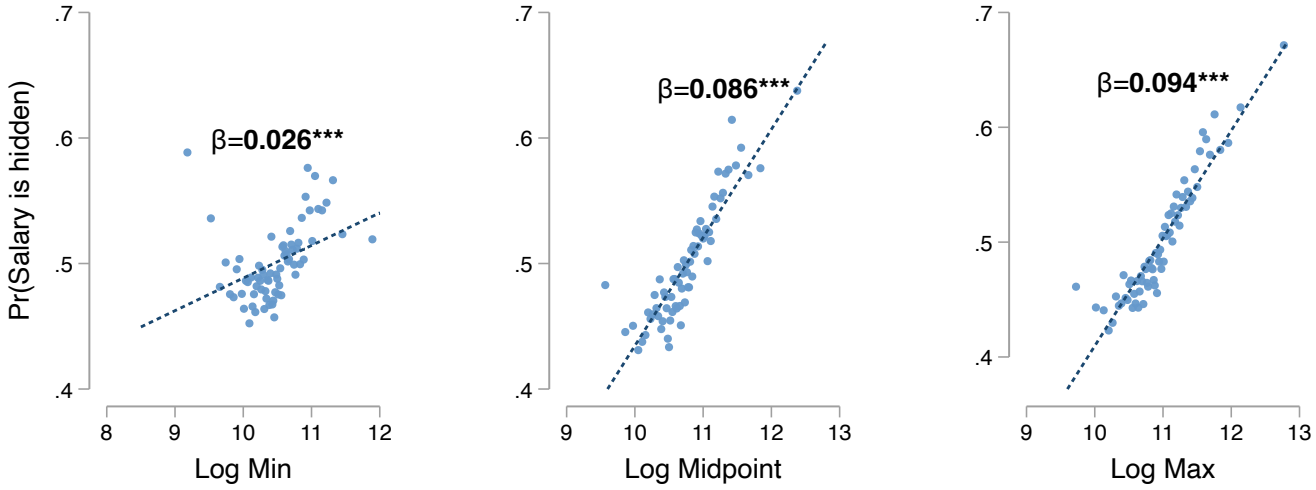
Notes: Each panel presents an example vignette from the discrete choice experiment shown to job-seekers, who were asked to choose between two job postings. Vignettes randomly vary salary visibility, non-wage amenities (remote work, transport support, and equal opportunity language), and a large firm 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 gray 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 pages 11 or **??**]

Figure A.7: Job complexity and salary non-disclosure



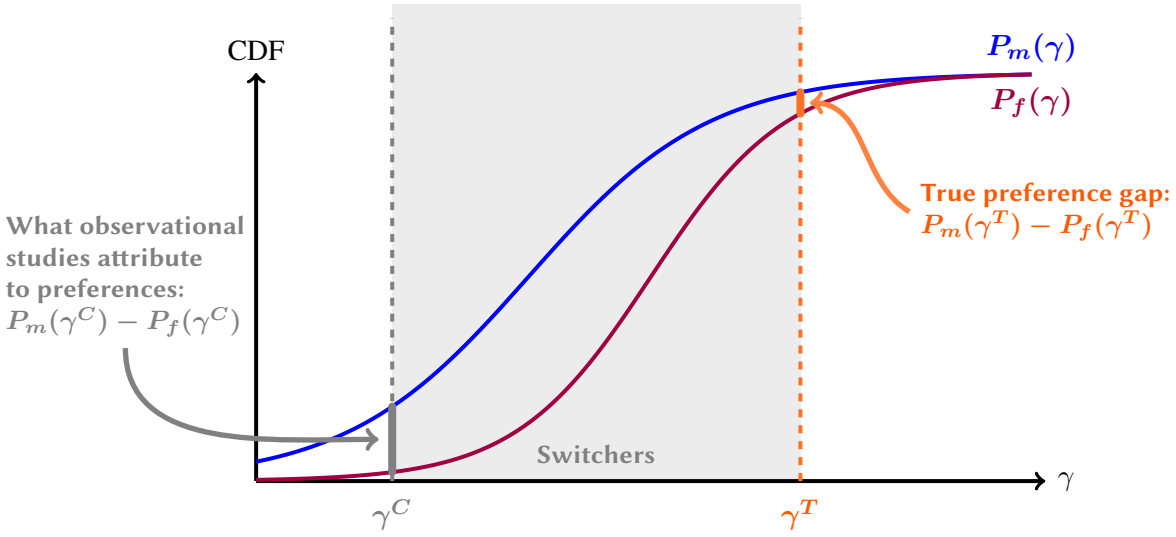
Notes: The figure shows correlations at the job level between an indicator for whether the salary is hidden, and various job characteristics. The sample consists of all baseline jobs. [Return to page 13]

Figure A.8: Salary levels and non-disclosure



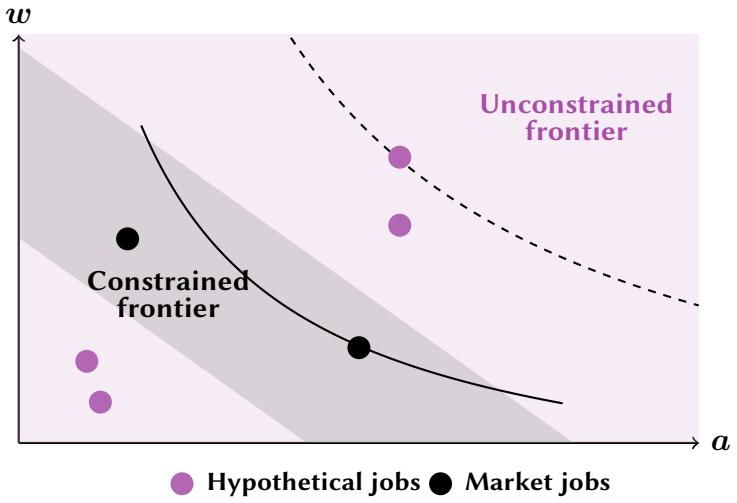
Notes: The figure shows the job-level relationship between an indicator for whether the salary of the job is hidden, and the log minimum, median and maximum salary of the job, residualized on industry, occupation and city fixed effects, and job characteristics (experience requirement, education requirement, 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 13]

Figure A.9: Gender gaps in observation studies that assume no frictions



Notes: The figure plots the cumulative distribution functions (CDFs) of the amenity index γ for men (P_m, blue) and women (P_f, red). The dashed lines mark the application cutoffs without transparency γ^C and with transparency γ^T; the gray shading between them is the "switcher band" $\gamma \in (\gamma^C, \gamma^T)$, where more women than men switch to large firms ($\Delta_f > \Delta_m$). The left gray arrow labels what observational studies typically attribute to preferences when they ignore frictions: the gap evaluated at γ^C, i.e., $P_m(\gamma^C) - P_f(\gamma^C)$. The right coral arrow shows that the true preference gap once frictions are removed is smaller: the gap at γ^T, i.e., $P_m(\gamma^T) - P_f(\gamma^T)$. [Return to page 24]

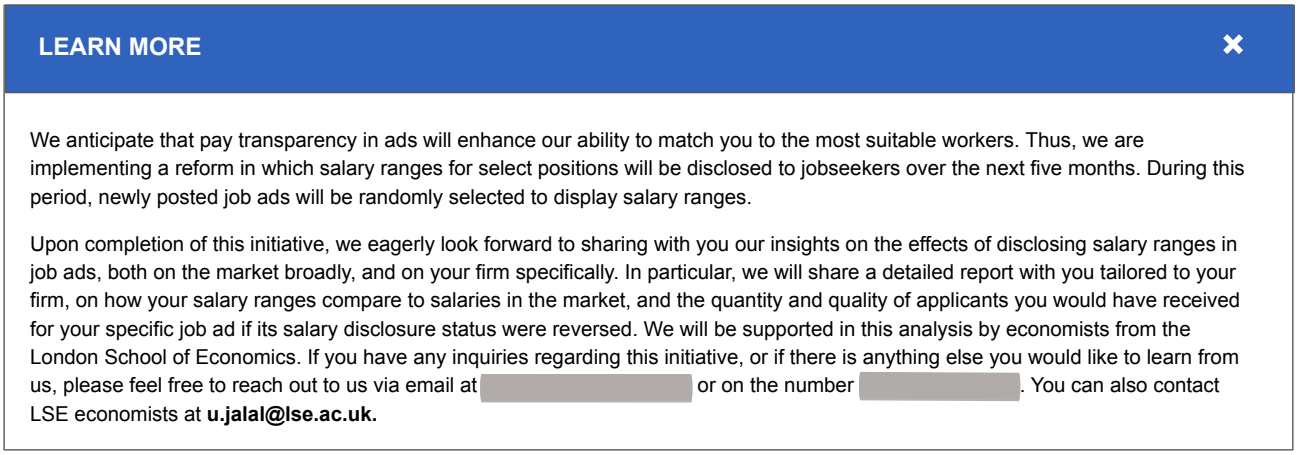
Figure A.10: Choices under constrained trade-offs



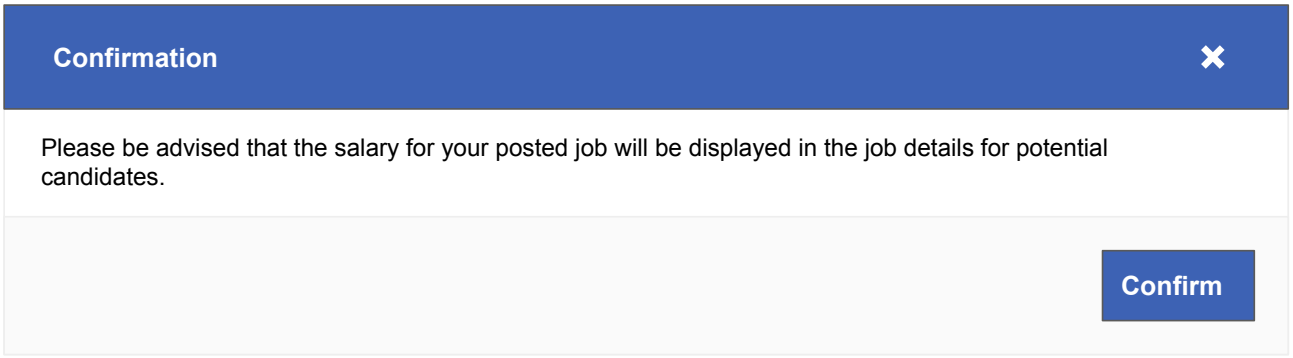
Notes: The graph presents the wage (w on y-axis) and amenity (a on x-axis) space. Gray area shows the constrained frontier of actual market jobs (black dots), where wages and amenities trade off negatively. Purple area represents the expanded choice set in vignette experiments, including hypothetical jobs (purple dots) with unrealistic wage-amenity combinations. The solid curve is the actual market frontier; the dashed curve shows preferences estimated from unconstrained vignette choices. This difference illustrates why vignette-based preference estimates may poorly predict real sorting behavior. [Return to page 24]

Figure A.11: Further information and inattention

(a) Further information for treated ads

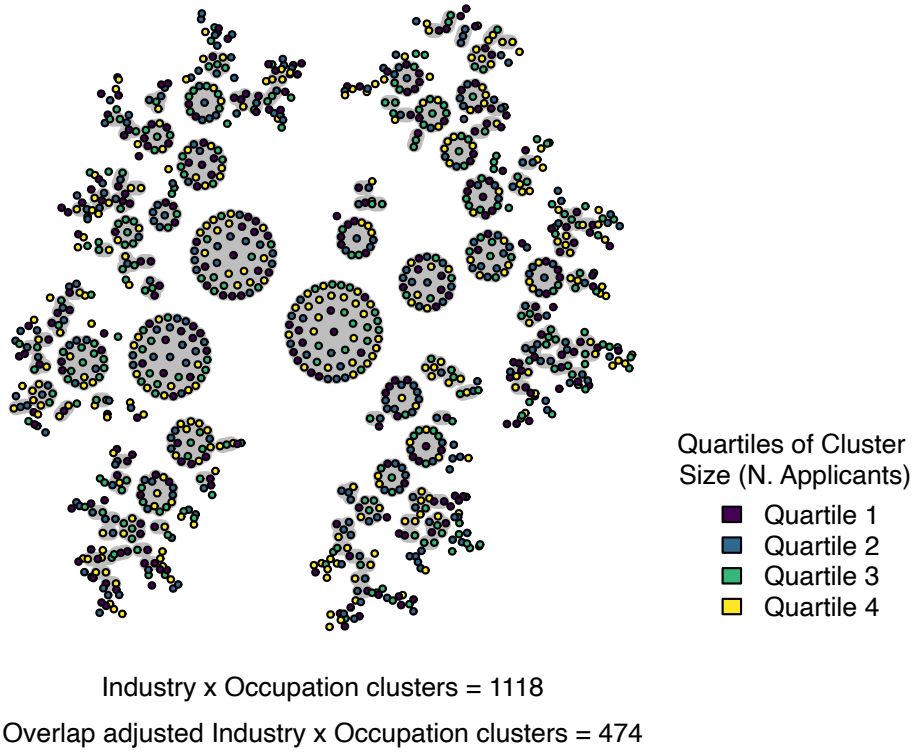


(b) Inattention



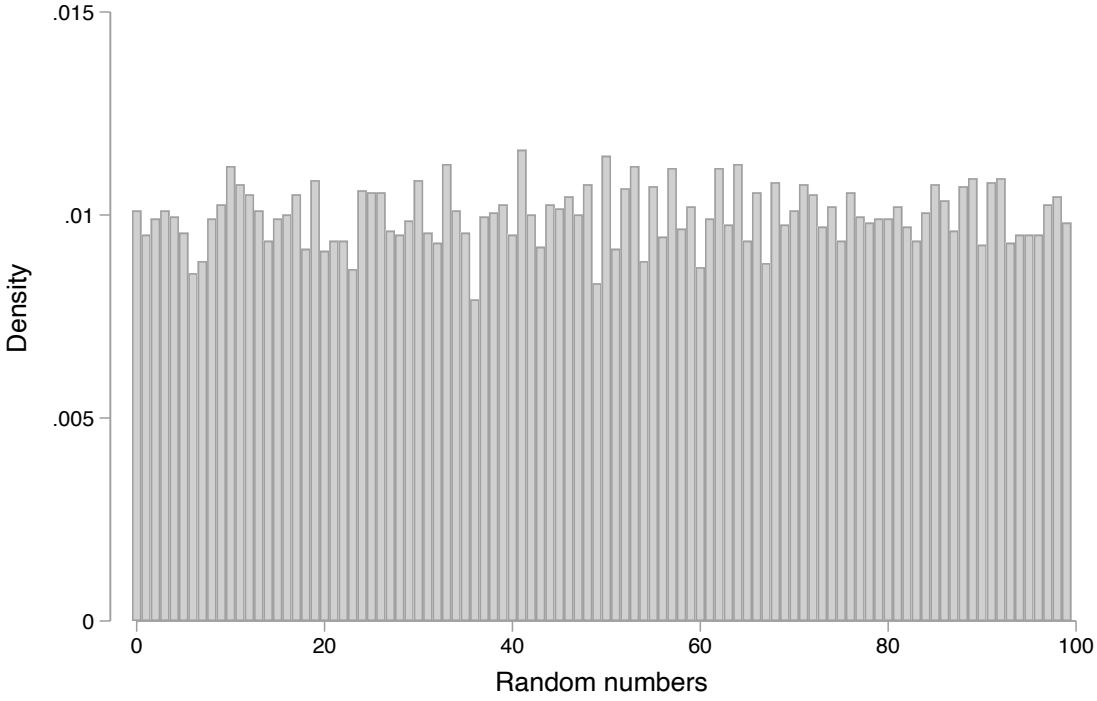
Notes: Panel (a) shows the "Learn More" pop-up provided to employers posting treated ads. It details the reform's purpose, emphasizing improved candidate matching and transparency. The pop-up invites employers to reach out with inquiries and mentions the involvement of researchers from the London School 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 25]

Figure A.12: Merged industry-occupation clusters based on historical overlap of applicants



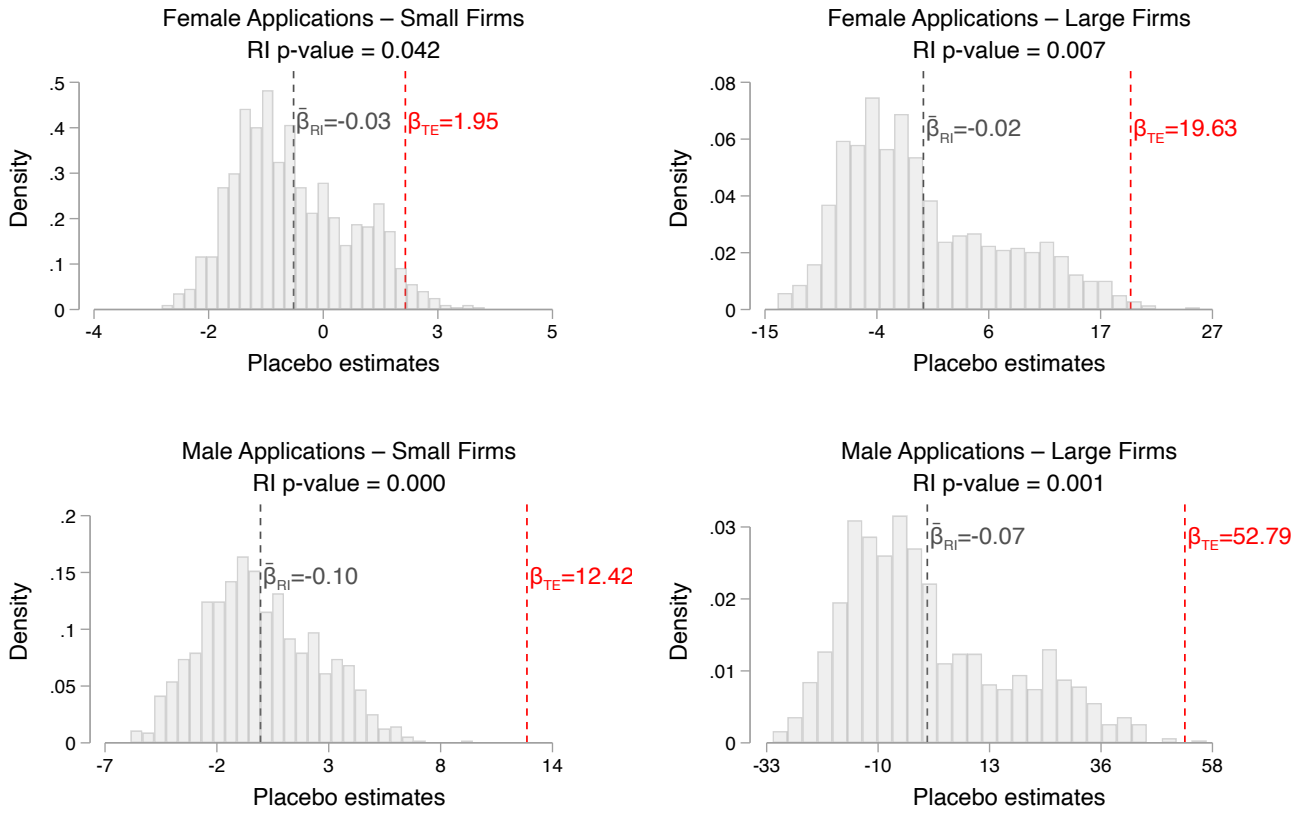
Notes: This figure visualizes the 474 labor market clusters used in the saturation design, formed by merging 1,118 initial industry–occupation cells based on historical application data. Industry-occupation cells with more than 10% mutual overlap in their applicant pools were merged to ensure that clusters are relatively self-contained in terms of job-seeker flows. Each node represents a unique industry-occupation cluster. The gray 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 26]

Figure A.13: Uniformity of Random Numbers



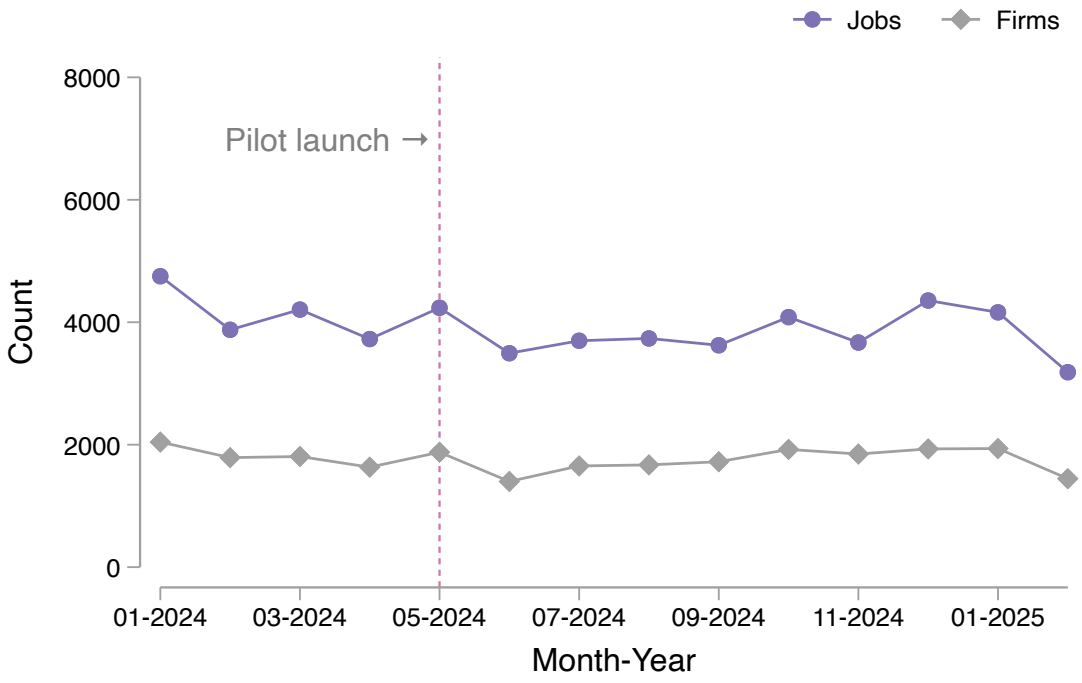
Notes: Distribution of the random numbers used to assign treatment across all job postings. Jobs were assigned to treatment if their random number—drawn from a continuous uniform distribution over $[0, 100]$—was less than or equal to a pre-specified threshold: 75 in high-saturation clusters and 25 in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

Figure A.14: 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 31]

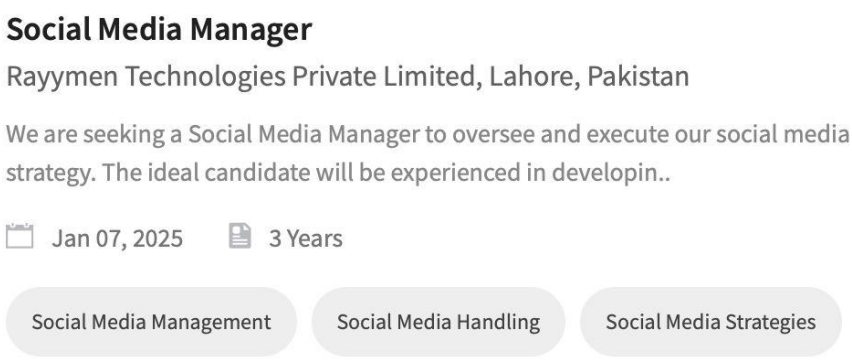
Figure A.15: No evidence of job posting or firm exits following experiment launch



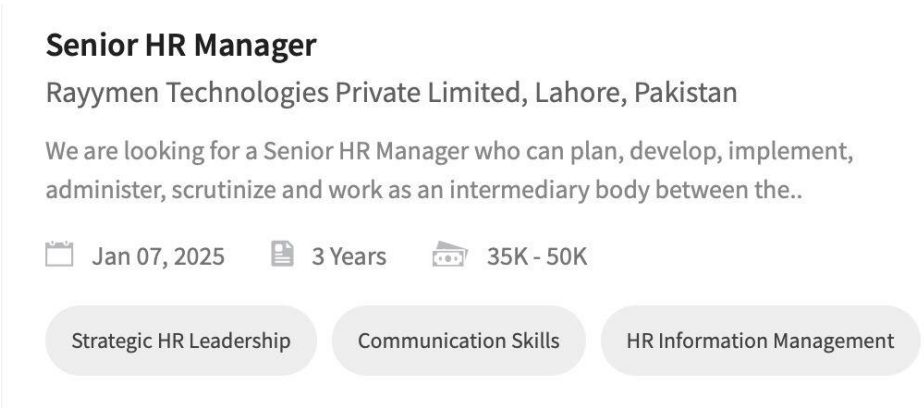
Notes: The figure shows the monthly counts of job postings (purple line) and active firms (orange line) on the platform from January 2024 to November 2024. The vertical dashed line marks the launch of the pilot intervention in May 2024. There is no evidence of a significant decline in the number of job postings or firms following the intervention, indicating no discernible exit of participants from the platform. [Return to page 28]

Figure A.16: Short previews of jobs

(a) Job with hidden salaries



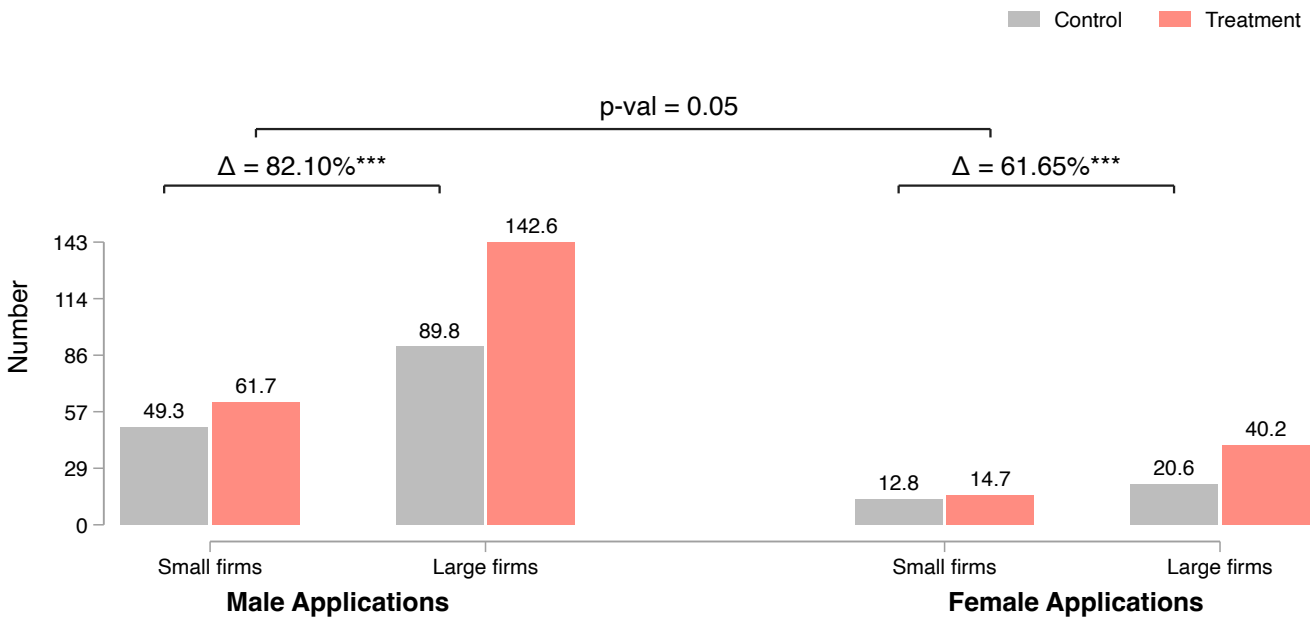
(b) Job with visible salaries



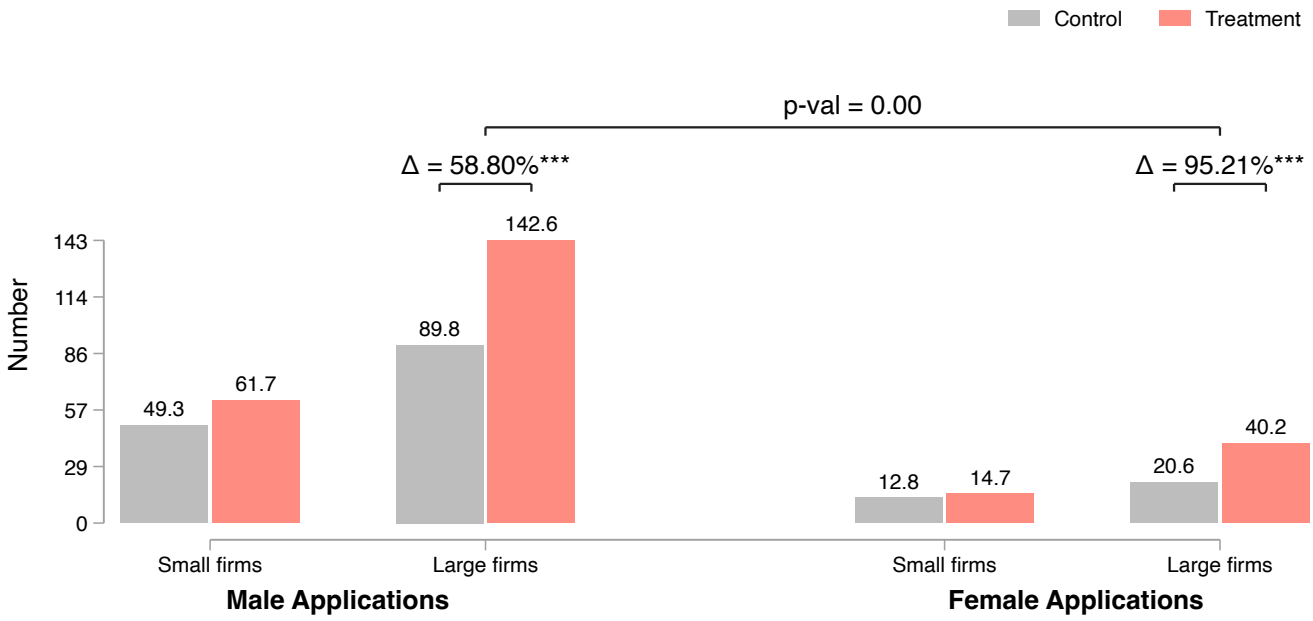
Notes: The figure presents initial previews of job postings on the platform that are visible to workers when they browse for jobs. Panel (a) illustrates a job posting where salary information is hidden, while Panel (b) shows a job posting with visible salary details. Specifically, it shows that the salary 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 29]

Figure A.17: Treatment impacts on job applications by gender (corresponding Table: B.6)

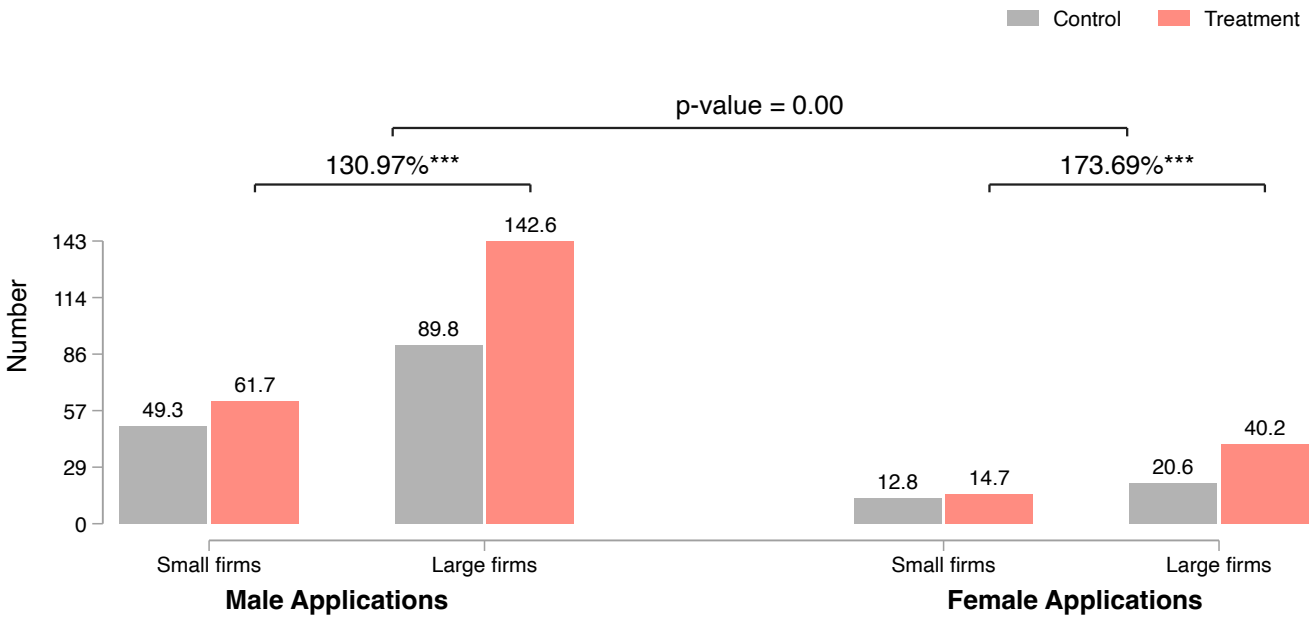
(a) Women's directed search to large firms is weaker than men's



(b) Reallocation of search is larger for women

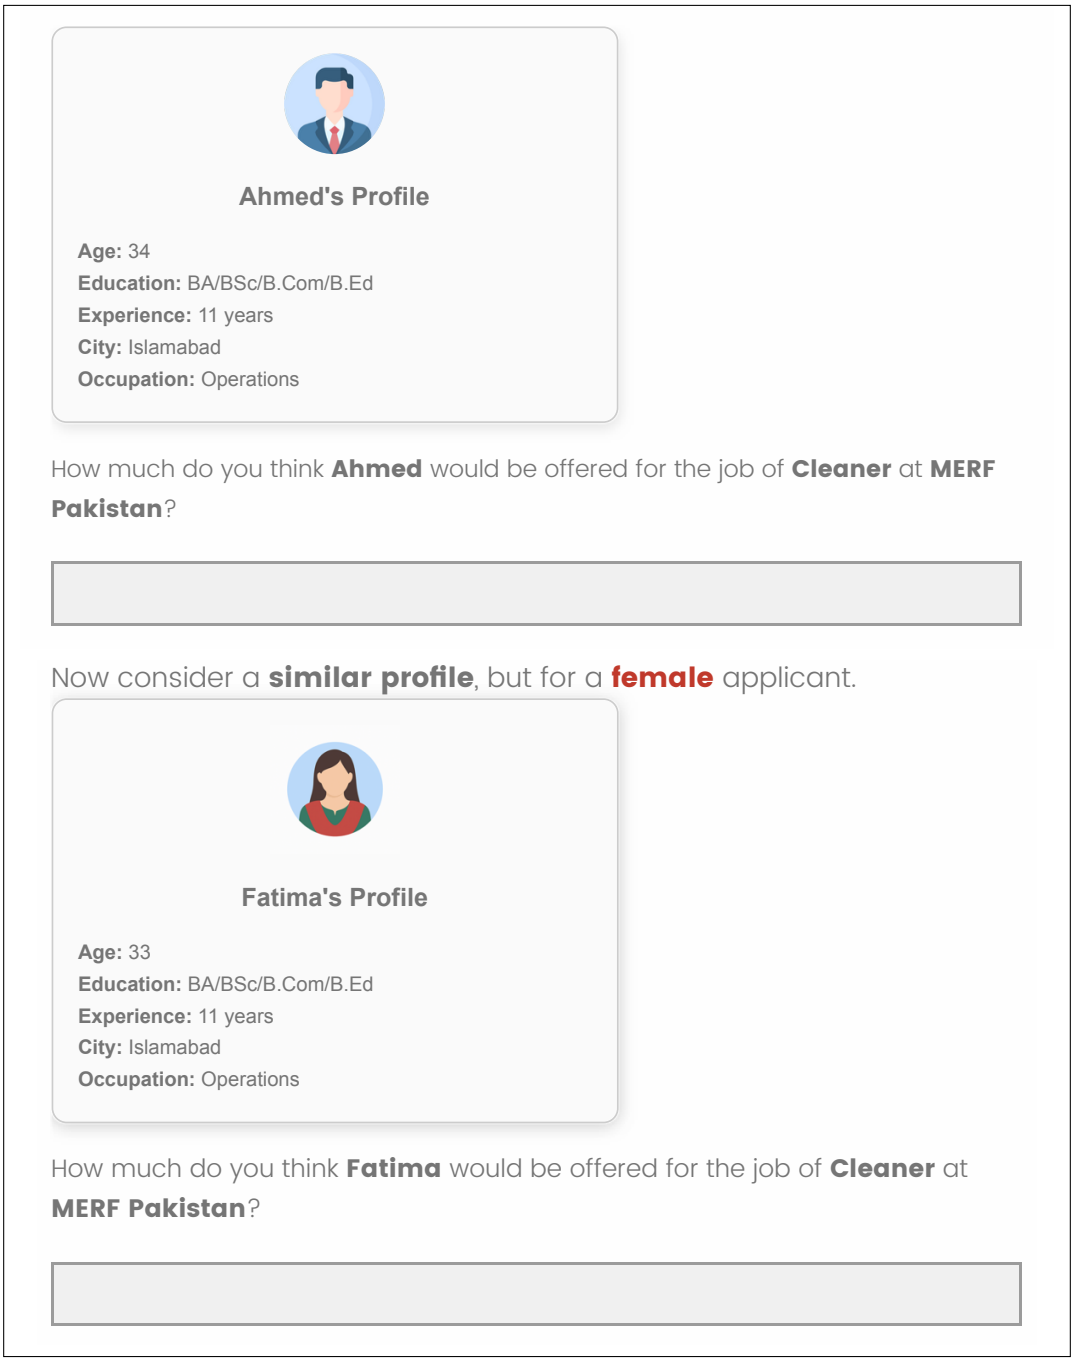


(c) Treatment closes gender gaps



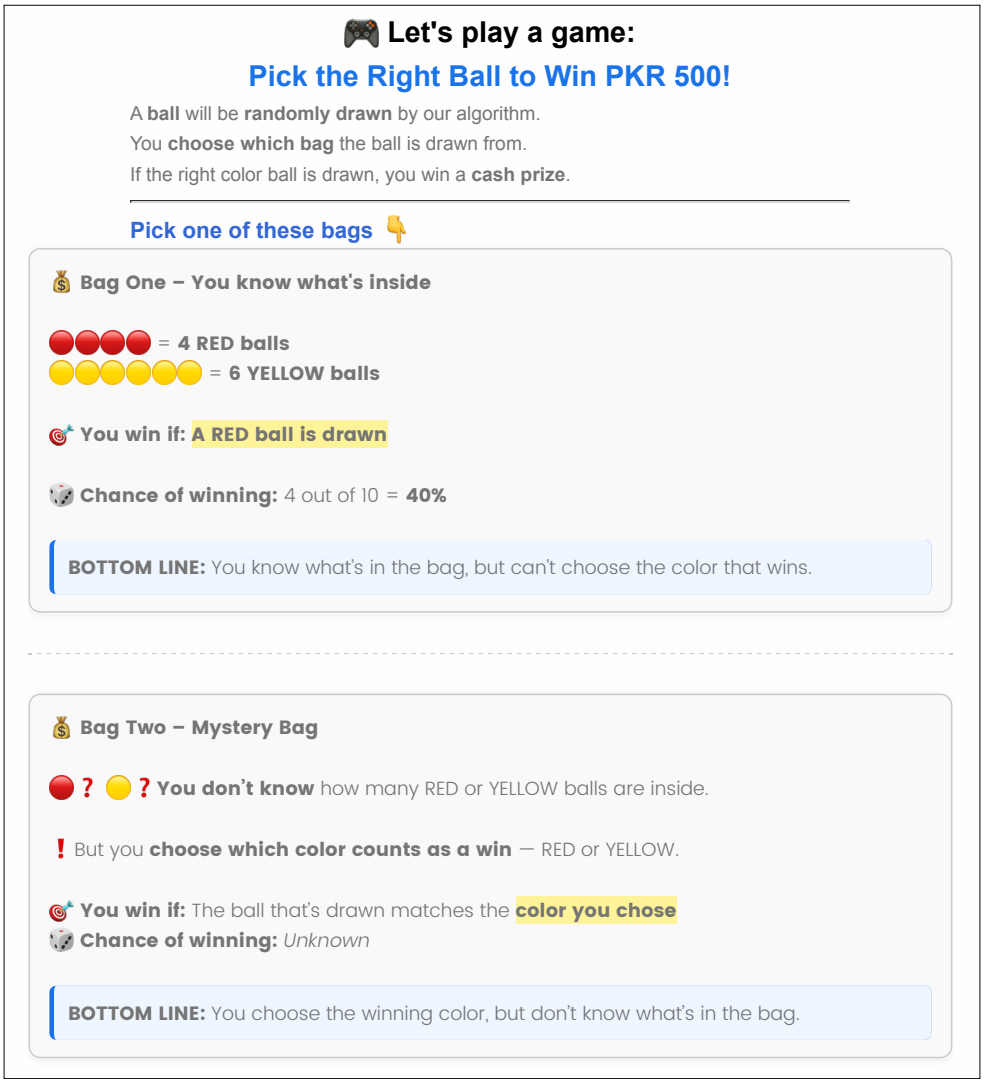
Notes: The figure shows treatment effects on applications with heterogeneity by firm size and gender. A tabular version of these results can be found in Appendix **Table** B.6. A firm is defined as large if it has more than 50 employees. The bars represent control and treatment group means. The lower 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 31 or 99]

Figure A.18: Hypothetical male and female profiles for salary offer elicitation



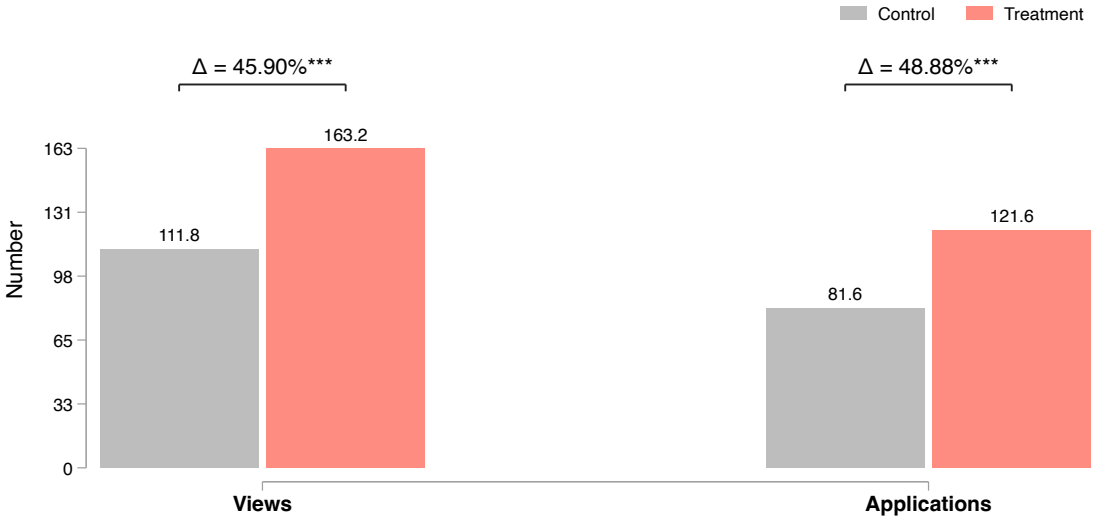
Notes: This figure shows an example of the hypothetical candidate profiles presented to survey respondents to elicit beliefs about gender discrimination in salary offers. Respondents were first shown a male candidate profile (Ahmed) with specific demographics, followed by an otherwise identical female candidate profile (Fatima). Both profiles matched the respondent's own education, age, city, and occupation. The profiles were customized for each respondent based on their survey responses. Respondents were asked to estimate the salary offer each candidate would receive for a specific job at a large firm whose ad they had previously viewed. [Return to page 40]

Figure A.19: Measuring ambiguity aversion with the Ellsberg Paradox



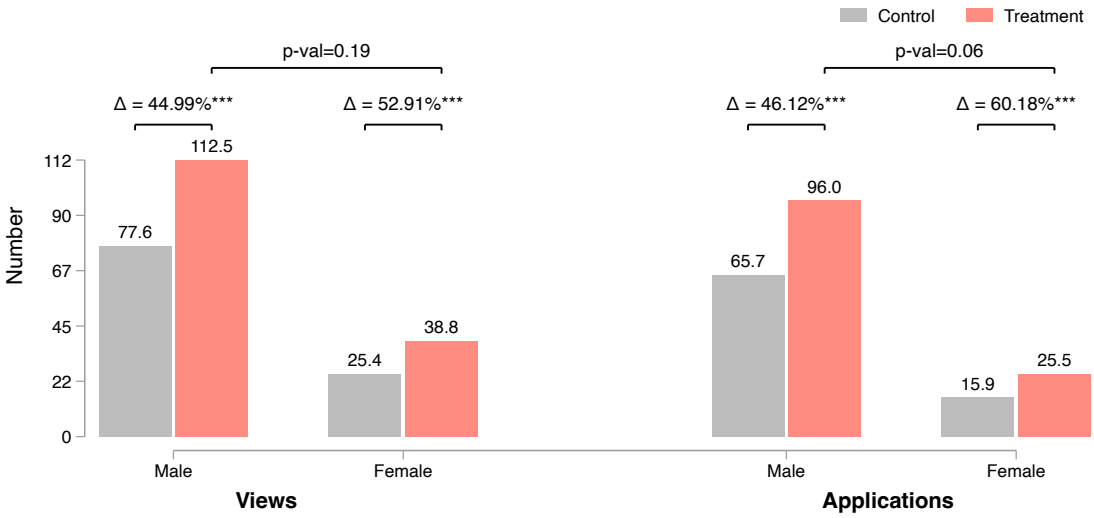
Notes: This figure illustrates the classic Ellsberg Paradox used to measure ambiguity aversion in the workers' survey. Bag One contains a known distribution of balls (4 red, 6 yellow), offering a 40% chance of winning if red is drawn. Bag Two contains an unknown distribution of red and yellow balls. Risk-neutral individuals should be indifferent between the two bags if they believe each color is equally likely in Bag Two. Ambiguity-averse individuals prefer Bag One despite identical expected payoffs, revealing a preference for known probabilities over unknown probabilities. In the survey, respondents choose between these two bags to reveal their degree of ambiguity aversion. [Return to page 42]

Figure A.20: Treatment effect on views and applications



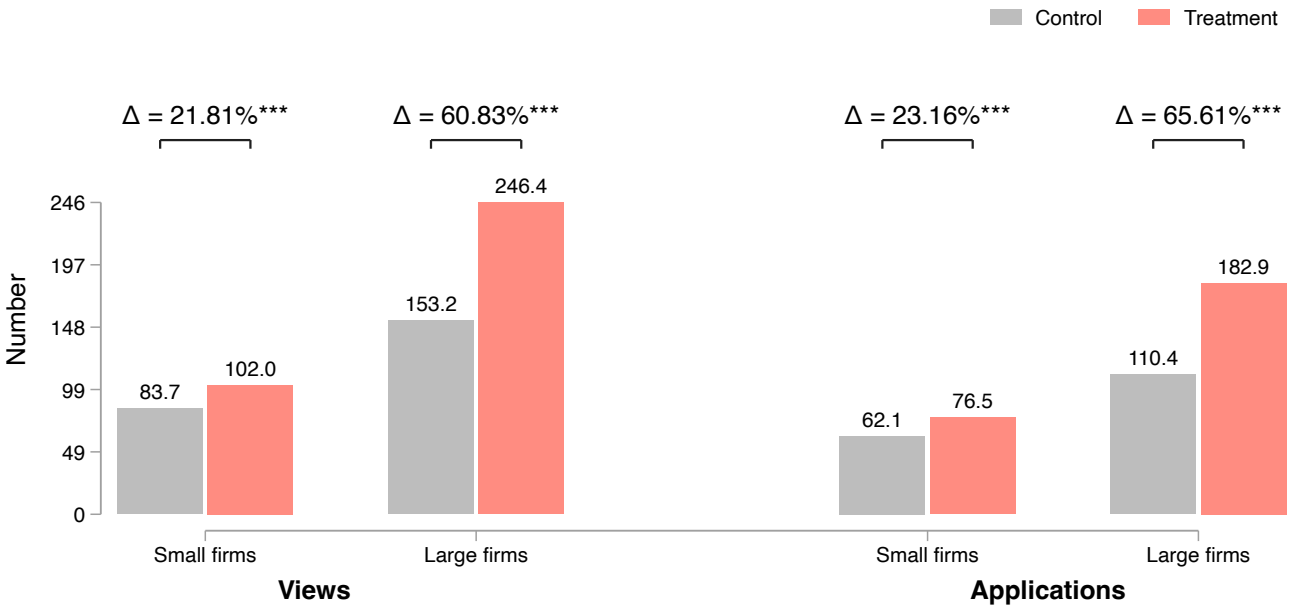
Notes: The figure shows treatment effect on job ad views and applications. The bars represent control and treatment group means, while the brackets represent percentage change from a poisson regression. [Return to page 59]

Figure A.21: 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 59]

Figure A.22: 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 60]

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 30]

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: Family-friendly amenities and firm size - baseline data

Panel A: Firm level

	Flex hours		Remote work		Need vehicle		Overnight travel		Transport support		Female supervisor	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Large firm	-0.10***	-0.08***	-0.06***	-0.04***	0.07***	0.06***	0.05***	0.04***	0.03**	0.01	-0.01	-0.01
	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Small firm	0.56	0.56	0.26	0.26	0.06	0.06	0.24	0.24	0.24	0.24	0.09	0.09
Observations	7,038	7,038	7,038	7,038	7,038	7,038	7,038	7,038	7,038	7,038	7,038	7,038
Industry FE		✓		✓		✓		✓		✓		✓

Panel B: Job level

	Flex hours		Remote work		Need vehicle		Overnight travel		Transport support		Female supervisor	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Large firm	-0.12***	-0.10***	-0.07***	-0.06***	0.03***	0.02***	0.05***	0.00	-0.01**	-0.01**	-0.01***	-0.00
	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Small firm	0.45	0.45	0.19	0.19	0.07	0.07	0.21	0.21	0.23	0.23	0.09	0.09
Observations	35,382	35,382	35,382	35,382	35,382	35,382	35,382	35,382	35,382	35,382	35,382	35,382
Industry FE		✓		✓		✓		✓		✓		✓
Occupation FE		✓		✓		✓		✓		✓		✓
Education		✓		✓		✓		✓		✓		✓
Career level		✓		✓		✓		✓		✓		✓
Experience		✓		✓		✓		✓		✓		✓

Notes: For a limited time, the platform mandated firms to report information on the amenities listed in the table. The sample consists of job advertisements (and corresponding firms) for which the platform collected this data. Panel A reports analysis at the firm level, where the dependent variable is the modal amenity value across all jobs posted by each firm. It is 1 if the firm offers the amenity in the majority of its job postings and 0 if the firm does not offer the amenity in the majority of its job postings. Panel B reports analysis at the job level, where the dependent variable is the amenity value for each individual job posting. It is 1 if the job offers the amenity and 0 if the job does not offer the amenity. In both panels, odd-numbered columns present unconditional differences by firm size, while even-numbered columns control for industry fixed effects (Panel A) or industry fixed effects, occupation fixed effects, education requirements, career level, and experience requirements (Panel B). Missing control variables are replaced with a placeholder value and a missing indicator is included. "Small firm" reports the mean of the dependent variable for small firms (fewer than 50 employees). $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 13 or 35]

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 30]

Table B.6: Treatment impacts by gender and firm size (Corresponding graph: Figure A.17)

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.20 for views, and Figure A.17 for applications. Robust standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. [Return to page 31 or 90]

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**
	(0.024)	(0.064)	
Low salary	0.827***	0.785***	-0.041
	(0.054)	(0.067)	
Medium salary	0.820***	0.831***	0.011
	(0.053)	(0.064)	
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
	(0.031)	(0.090)	
Low salary	0.664***	0.593***	-0.072
	(0.091)	(0.122)	
Medium salary	0.659***	0.672***	0.013
	(0.084)	(0.101)	
High salary	0.831***	0.720***	-0.111
	(0.062)	(0.102)	
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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. [Return to page 35]

Table B.8: 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 32].

Table B.9: 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 32].

Table B.10: 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	0.028	0.027	0.040	-0.041
	(0.026)	(0.029)	(0.026)	(0.030)
Male Mean	0.070	0.090	0.065	0.135
Observations	456	456	456	456

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 37]

Table B.11: 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	-0.015	-0.005	0.008	-0.009
	(0.019)	(0.018)	(0.018)	(0.018)
Male Mean	0.050	0.040	0.035	0.040
Observations	456	456	456	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 37]

Table B.12: Accuracy of worker beliefs about firm size

	Guess firm size correctly:	
	(1) Small firm	(2) Large firm
Female	0.037	-0.017
	(0.049)	(0.045)
Male Mean	0.347	0.728
Observations	402	402

Notes: This table reports the probability that job-seekers 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." The number of observations is lower relative to Tables B.10 and B.11 because this question was asked later in the survey and some respondents dropped out of the survey before it was asked. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 38]

Table B.13: Gender differences in preferences for salary levels

Salary Level	Odds Ratio		Difference (Female-Male)	
	Female	Male	Percentage Points	p-value
	(1)	(2)	(3)	(4)
Medium (vs. Low)	2.17***	2.48***	-1.67	0.670
	(0.45)	(0.61)	(5.42)	
High (vs. Low)	3.41***	2.89***	3.17	0.756
	(1.18)	(1.10)	(7.29)	
Observations	1021	773		

Notes: This table reports conditional logit estimates from a discrete choice experiment in which respondents choose between pairs of stylized jobs. The analysis is restricted to choices where salary information was displayed. When visible, displayed salary ranges were constructed by randomly selecting a minimum from {0.7, 0.75, 0.8} and a maximum from {1.1, 1.15, 1.2} times the respondent's previously reported expected salary. Salary levels are defined by the midpoint of the displayed range: low (90-92.5% of expected salary, the baseline), medium (95-97.5%), and high (100%). Columns 1-2 show odds ratios for the effect of medium and high salary levels (relative to a low midpoint) on the likelihood of a job being chosen, estimated separately for women and men, controlling for other job attributes (remote work, diversity language, transport subsidy, and firm size). Column 3 reports the gender difference in marginal effects (female minus male) in percentage points, i.e., the difference in choice probability when salary level increases. Column 4 reports p-values from tests of whether the utility weights for salary levels differ significantly between women and men. Standard errors are clustered at the respondent level and shown in parentheses. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 39]

Table B.14: Gender differences in networks

	(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 39]

Table B.15: Beliefs about discrimination by large firms

	(1) **Perceive discrimination**	(2) **Ratio: Own offer to Perceived max pay**
Female	0.054	0.086
	(0.041)	(0.080)
Male Mean	0.224	0.908
Observations	444	444

Notes: This table reports gender differences in beliefs about salaries and discrimination at large firms. Respondents viewed a large-firm job ad in their occupation without disclosed salaries and were asked to guess the minimum and maximum salary range. They were then shown two otherwise identical candidate profiles (one male, one female) and asked to estimate the salary offer each would receive. Column (1) equals one if respondents expected men to receive higher offers than women. Column (2) reports the ratio of the salary respondents thought they would personally receive to the maximum salary they guessed for the job. Male means are reported in the row labeled "Male Mean." The female coefficient reports the gender difference in women's answers relative to men's. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 40]

Table B.16: 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 41]

Table B.17: 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 41]

Table B.18: Ambiguity and bargaining preferences

	Gender gap in preferences		Salary visibility preference by bargaining aversion (DCE)	
	(1) Ambiguity averse	(2) Bargaining averse	(3) Odds Ratio	(4) Percentage Points
Female	0.032	0.182***		
	(0.048)	(0.049)		
Salary visible × Bargaining averse			2.726***	0.181***
			(0.411)	(0.024)
Salary visible × Prefer bargaining			2.826***	0.188***
			(0.583)	(0.034)
Male Mean	0.588	0.434		
p-value (difference)			0.888	0.888
Observations	419	419	1,896	1,896

Notes: This table reports gender differences in ambiguity and bargaining preferences, and examines whether these preferences predict salary visibility preferences in a discrete choice experiment (DCE). Column (1) reports whether respondents are ambiguity averse in an incentivized urn game following Ashraf et al. (2009). Column (2) reports whether respondents prefer a non-negotiable job offer over a negotiable one, coded as bargaining aversion. Columns (3) and (4) estimate how salary disclosure affects job choice among women, separately for those who are bargaining averse versus those who prefer bargaining. Column (3) reports odds ratios from a conditional logit model, where coefficients represent the multiplicative effect on the odds of choosing a job when salary is disclosed. Column (4) reports the corresponding marginal effects in percentage points (for ease of interpretation), showing the change in the probability of choosing a job when salary is disclosed. The p-value tests whether the underlying log-odds coefficients generated by the logit model are significantly different by bargaining aversion. Male means are reported in the row labeled "Male Mean." $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. [Return to page 42]

Table B.19: 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 48].

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 cutoff ($\theta \geq$

θ_{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:

$$\pi(w) = [1 - F(\theta_{\min})^{n(w)}] \left(\bar{\theta}(w) - w\right) - c(n(w)) \tag{9}$$

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:

$$U_\theta = F(\theta)^{n-1}[1 - F(\theta_{min})] \left(\lambda\theta_i + w\right) - V(\theta) - s \geq 0 \tag{10}$$

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.23 .

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.23, 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):

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

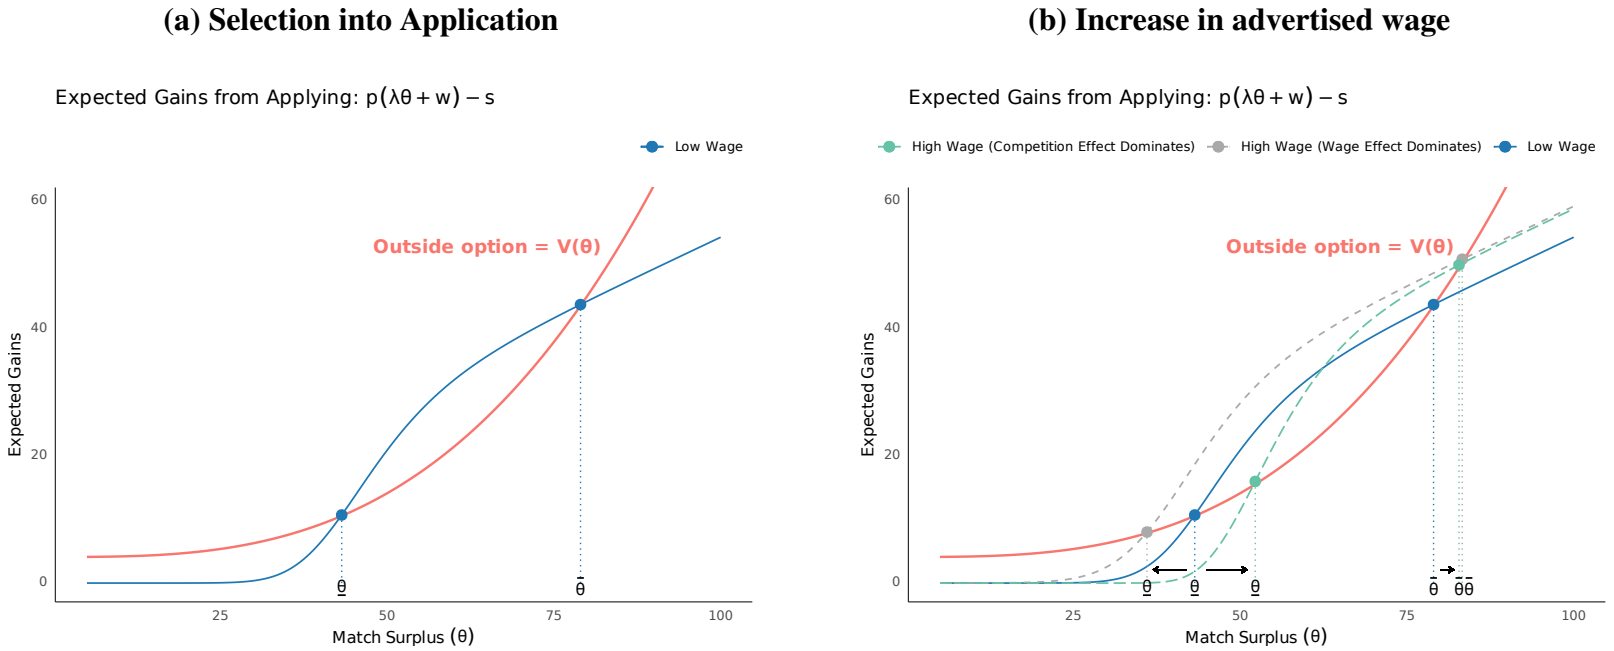
But this effect is stronger for low-ability workers because higher ability workers are now selecting into the applicant pool, reducing the likelihood that low ability workers get hired (proof in Appendix Section B.4):

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

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.23: Theoretical framework: Application thresholds



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 111]

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:

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

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:

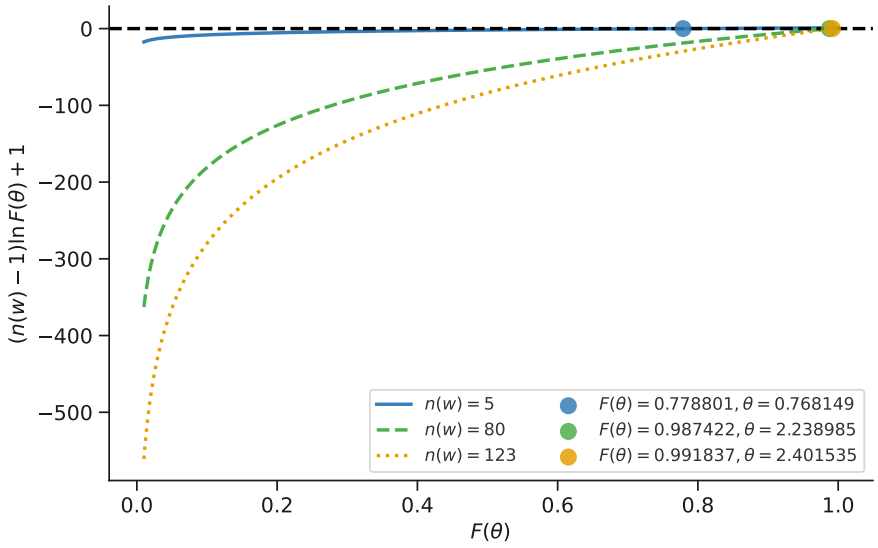
$$(n(w)-1)\ln F(\theta) > -1$$

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

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

– 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.24: 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*.

C Extension: Wage dispersion and rational expectations about firm-size wage premiums

This appendix extends the baseline framework in Section 4 by allowing wages to be draws from firm-size-specific distributions, rather than deterministic, and lets workers hold rational expectations about these distributions. The small firm discloses the realized wage attached to its posting, whereas the large firm conceals it. The extension below shows that even when large firms pay more than small firms on average, and workers know this, they may rationally choose to apply to the small firm due to risk aversion and strong amenity tastes. While wage non-disclosure forces workers to evaluate every large-firm posting at a single pooled, certainty-equivalent wage, transparency lets them instead compare the realized wage on each posting directly against their compensating threshold. As a result, workers who would have avoided the large firm under opacity become more likely to apply once a disclosed wage is high enough to compensate them.

C.1 Sorting under opacity with rational expectations

There are two job types, indexed by $j \in \{L, S\}$. As in the baseline model, large firms offer fewer family-friendly amenities,

$$a_L < a_S, \tag{11}$$

where a higher value of a_j denotes a more desirable amenity bundle. A worker with amenity taste γ_g and wage preferences θ has utility

$$u_g(w, a) = w^\theta a^{\gamma_g} \tag{12}$$

Throughout, $0 < \theta \leq 1$ and $\gamma_g \geq 0$. The parameter γ_g may vary across genders, as in the baseline model. When $\theta < 1$, utility is concave in wages and workers are averse to wage risk.

Wages at small and large firms are drawn from the distributions

$$W_S \sim F_S \qquad \text{and} \qquad W_L \sim F_L, \tag{13}$$

with $W_S > 0$ and $W_L > 0$. Throughout, uppercase W_L, W_S denote these underlying random variables, and lowercase w_L, w_S denote particular realized wages once observed; a probability or expectation is taken over any wage that is being treated as not yet revealed. Assume that large firms pay more on average:

$$\mathbb{E}[W_L] > \mathbb{E}[W_S]. \tag{14}$$

Workers have rational expectations and know both distributions. When considering a small-firm posting, the worker observes a realization w_S drawn from F_S. By contrast, the worker knows F_L but does not observe the realization of W_L before applying to the large firm. Conditional on observing w_S, the worker therefore applies to a large firm over a small firm if:

$$\mathbb{E}[W_L^\theta] a_L^{\gamma_g} \geq w_S^\theta a_S^{\gamma_g}. \tag{15}$$

Define the certainty-equivalent wage associated with the concealed large-firm wage as

$$w_L^{CE} \equiv \left(\mathbb{E}[W_L^\theta]\right)^{1/\theta}. \tag{16}$$

This is the known wage that would give the worker the same expected utility as the concealed large-firm wage distribution.

Proposition 1 (Application choice under wage opacity)**.** *After observing a small-firm wage w_S, a worker of type γ_g applies to the large firm if and only if*

$$w_L^{CE} \geq w_S \left(\frac{a_S}{a_L}\right)^{\gamma_g/\theta}. \tag{17}$$

The worker chooses the small firm whenever the certainty-equivalent large-firm wage falls below the compensating wage in equation (17).

This highlights that workers can rationally sort toward small firms despite knowing that large firms pay more on average, because under wage opacity, they compare the observed small firm wages to the large firm certainty equivalent wage, not the large firm average wage.

Two forces can reduce this certainty equivalent wage, making it insufficient to draw workers away from small firms: aversion to wage risk and a strong preference for amenities.

Aversion to wage risk. Suppose $\theta < 1$, so that utility is concave in wages. Aversion to wage risk reduces the attractiveness of the concealed large-firm wage in two related ways.

First, greater curvature in wages lowers the certainty-equivalent value of a given non-degenerate wage distribution. The certainty equivalent in equation (16) is the power mean of W_L of order θ, which is increasing in θ. Thus, a lower value of θ lowers w_L^{CE} and makes equation (17) less likely to hold.

Second, risk aversion causes workers to penalize dispersion in concealed wages. Fix $\theta \in (0, 1)$ and consider an alternative large-firm wage distribution W_L' that is a mean-preserving spread of W_L, holding the support fixed. Concavity of w^θ implies

$$\mathbb{E}[(W_L')^\theta] \leq \mathbb{E}[W_L^\theta], \tag{18}$$

with strict inequality for a non-trivial spread. Therefore,

$$\left(\mathbb{E}[(W_L')^\theta]\right)^{1/\theta} \leq \left(\mathbb{E}[W_L^\theta]\right)^{1/\theta}. \tag{19}$$

Sufficiently high dispersion in concealed large-firm wages therefore lowers the certainty-equivalent value, making equation (17) less likely to hold.

High preference for amenities. The wage required to compensate for the large firm's amenity disadvantage is increasing in γ_g. Since $a_S/a_L > 1$, the right-hand side of equation (17) rises with the worker's preference for amenities. Thus, holding the wage distributions and θ fixed, workers with a higher γ_g are less likely to apply to the large firm. In particular, whenever $w_L^{CE} > w_S$, the worker chooses the small firm if

$$\gamma_g > \frac{\theta \log\left(w_L^{CE}/w_S\right)}{\log(a_S/a_L)}. \tag{20}$$

Differences in amenity preferences across genders can therefore generate differential sorting toward small firms even when workers hold rational beliefs about the wage distributions.

Summary. Even when workers know that large firms pay more on average and form rational expectations using the large-firm wage distribution, they may choose small firms if they are averse to the concealed wage risk or because they place a high value on amenities.

C.2 Sorting under pay transparency

Under pay transparency, the worker observes both realized wages before choosing where to apply. Assume throughout this section that W_L and W_S are independent draws. The worker applies to the large firm if and only if

$$w_L^\theta a_L^{\gamma_g} \geq w_S^\theta a_S^{\gamma_g}, \tag{21}$$

or equivalently,

$$w_L \geq w_S \left(\frac{a_S}{a_L}\right)^{\gamma_g/\theta}. \tag{22}$$

Define

$$\Delta a \equiv \left(\frac{a_S}{a_L}\right)^{\gamma_g/\theta} > 1. \tag{23}$$

Under transparency, the worker therefore applies to the large firm whenever $W_L \geq \Delta a W_S$.

Consider a worker who observes a small-firm wage w_S satisfying $w_L^{CE} < \Delta a w_S$. Under opacity, this worker does not apply to the large firm because its certainty-equivalent wage is insufficient to compensate for its amenity disadvantage. Thus, with C denoting the control condition of the

experiment where wages are opaque, the probability of application to the large firm is:

$$p_g^C(w_S) = 0. \tag{24}$$

Under transparency, however, the worker observes the realized large-firm wage and applies whenever $W_L \geq \Delta a w_S$. Across realizations of the large-firm wage, the probability of applying to the large firm is therefore

$$p_g^T(w_S) = \mathbb{P}\left(W_L \geq \Delta a w_S\right) = 1 - F_L\left(\Delta a w_S\right), \tag{25}$$

where the final equality assumes that F_L is continuous.

Let $w_L^{\max} \equiv \sup\{w : F_L(w) < 1\}$ denote the upper support of F_L. Continuity of F_L implies that $\mathbb{P}\left(W_L \geq x\right) > 0$ if and only if $x < w_L^{\max}$.

Lemma 1 (Opacity pools the large-firm wage distribution at its certainty equivalent)**.** *If F_L is non-degenerate, then*

$$\underline{w}_L < w_L^{CE} < w_L^{\max}, \tag{26}$$

so both intervals $\left(\underline{w}_L,\, w_L^{CE}\right)$ and $\left(w_L^{CE},\, w_L^{\max}\right)$ are nonempty, where $\underline{w}_L \equiv \inf\{w : F_L(w) > 0\}$ denotes the lower support of F_L.

The wedge on the right of equation (26) follows because $W_L^\theta \leq (w_L^{\max})^\theta$ with strict inequality on a set of positive probability, so $\mathbb{E}[W_L^\theta] < (w_L^{\max})^\theta$ (if $w_L^{\max} = \infty$, the wedge is immediate); the wedge on the left follows symmetrically from $W_L^\theta \geq (\underline{w}_L)^\theta$ with strict inequality on a set of positive probability.

Lemma 1 gives the sense in which opacity matters. Under opacity, every large-firm posting is evaluated as if it paid w_L^{CE} with certainty: wage variation across large-firm postings is invisible to applicants, so postings that would otherwise pay very differently are pooled at the single perceived value w_L^{CE}, and neither realizations above w_L^{CE} nor realizations below it change an applicant's behavior.

Transparency unpools these postings by revealing which of them lie in the upper tail $\left(w_L^{CE},\, w_L^{\max}\right]$ and which lie in the lower tail $\left[\underline{w}_L,\, w_L^{CE}\right)$. Both propositions below follow the same logic: whenever a worker does not apply under opacity ($p_g^C(w_S) = 0$ or $= 1$), transparency sets $p_g^T(w_S) = 1 - F_L(\Delta a w_S)$ regardless, so the comparison reduces to where $\Delta a w_S$ falls relative to w_L^{CE}.

Proposition 2 (Transparency draws in workers priced out by the certainty equivalent)**.** *If $w_L^{CE} < \Delta a w_S < w_L^{\max}$, the worker does not apply under opacity ($p_g^C(w_S) = 0$), but transparency strictly increases the probability of applying:*

$$p_g^T(w_S) > p_g^C(w_S) = 0, \tag{27}$$

since $p_g^T(w_S) = 1 - F_L(\Delta a w_S) > 0$. *By Lemma 1, this condition holds on a nonempty interval of compensating wages: workers priced out by the pooled value* w_L^{CE} *but who would clear the threshold at some disclosed wages in the upper tail.*

Proposition 3 (Transparency screens out workers oversold by the certainty equivalent)**.** *If* $\underline{w}_L < \Delta a w_S \leq w_L^{CE}$*, the worker applies with certainty under opacity (*$p_g^C(w_S) = 1$*), but transparency strictly decreases the probability of applying:*

$$p_g^T(w_S) < p_g^C(w_S) = 1, \tag{28}$$

since $p_g^T(w_S) = 1 - F_L(\Delta a w_S) < 1$. *By Lemma 1, this condition also holds on a nonempty interval of compensating wages: workers the pooled value* w_L^{CE} *oversold, who withdraw once disclosed wages in the lower tail fall short.*

The two propositions are mirror images, governed by which side of w_L^{CE} a worker's compensating wage falls on. To characterize the overall application probability, define

$$P_g^C = \mathbb{P}\left(w_L^{CE} \geq \Delta a W_S\right) \tag{29}$$

under opacity and

$$P_g^T = \mathbb{P}\left(W_L \geq \Delta a W_S\right) \tag{30}$$

under transparency. Their difference is

$$\begin{aligned} P_g^T - P_g^C = &\mathbb{P}\left(w_L^{CE} < \Delta a W_S \leq W_L\right) \\ &- \mathbb{P}\left(W_L < \Delta a W_S \leq w_L^{CE}\right). \end{aligned} \tag{31}$$

The first term is the mass of workers induced to apply to the large firm by the disclosure of a sufficiently high wage. The second term is the mass of workers who would have applied under opacity but are discouraged by the disclosure of a low wage.

Whether transparency raises the overall application probability depends on which margin dominates. By Propositions 2 and 3, the two margins involve disjoint sets of worker–posting pairs: the exit term concerns only pairs with $\Delta a W_S \leq w_L^{CE}$, and the entry term only pairs with $\Delta a W_S > w_L^{CE}$, so each worker–posting pair contributes to at most one of the two. Because Δa is increasing in γ_g, which margin a worker is on is determined by their amenity preference. Conditional on a small-firm wage $w_S < w_L^{CE}$, equation (20) states that the worker applies to the large firm under opacity if and only if

$$\gamma_g \leq \gamma^*(w_S) \equiv \frac{\theta \log\left(w_L^{CE}/w_S\right)}{\log(a_S/a_L)}. \tag{32}$$

Proposition 4 (Amenity preference determines whether transparency induces entry or exit)**.** *Conditional on a small-firm wage $w_S < w_L^{CE}$, the effect of transparency on the probability of applying to the large firm is*

$$p_g^T(w_S) - p_g^C(w_S) = \begin{cases} -F_L\left(\Delta a w_S\right) \leq 0, & \gamma_g \leq \gamma^*(w_S), \\ 1 - F_L\left(\Delta a w_S\right) \geq 0, & \gamma_g > \gamma^*(w_S), \end{cases} \tag{33}$$

with the second case strictly positive whenever $\Delta a w_S < w_L^{\max}$. (If $w_S \geq w_L^{CE}$, every worker falls into the second case.)

The two cases are exactly Propositions 3 and 2, restated in terms of γ_g rather than w_L^{CE}: equation (20) shows the opacity cutoff $\gamma^*(w_S)$ is where $p_g^C(w_S)$ switches from 1 to 0, and $p_g^T(w_S) = 1 - F_L(\Delta a w_S)$ throughout, so subtracting gives equation (33) directly.

The key is that γ_g pins down the *sign* of the treatment effect: transparency's effect on a worker is exit if $\gamma_g \leq \gamma^*(w_S)$ and entry if $\gamma_g > \gamma^*(w_S)$, with the sign flipping discontinuously at the cutoff (the effect jumps from $-F_L(\Delta a w_S)$ to $1 - F_L(\Delta a w_S)$, an upward jump of one).

The same logic delivers the ex ante effect of transparency in terms of primitives. Let $\underline{w}_S$ and $\overline{w}_S$ denote the lower and upper support of F_S, and define

$$\gamma_{\text{low}} \equiv \frac{\theta \log\left(w_L^{CE}/\overline{w}_S\right)}{\log(a_S/a_L)} \qquad \text{and} \qquad \gamma_{\text{high}} \equiv \frac{\theta \log\left(w_L^{CE}/\underline{w}_S\right)}{\log(a_S/a_L)}, \tag{34}$$

with $\gamma_{\text{low}} \leq \gamma_{\text{high}}$. Workers with $\gamma_g \leq \gamma_{\text{low}}$ apply to the large firm at every realization of W_S under opacity, so the entry margin in equation (31) is always empty for them and $P_g^T - P_g^C \leq 0$. Workers with $\gamma_g > \gamma_{\text{high}}$ never apply under opacity, so the exit margin is always empty for them and $P_g^T - P_g^C \geq 0$, strictly whenever $\Delta a\, \underline{w}_S < w_L^{\max}$.

For intermediate amenity preferences, $\gamma_g \in (\gamma_{\text{low}}, \gamma_{\text{high}}]$, both margins are active and the sign of $P_g^T - P_g^C$ is given by the decomposition in equation (31); it is ambiguous in general. It depends on (1) how many worker–posting pairs are on the entry rather than exit side of the certainty equivalent, and (2) how often the realized large-firm wage actually changes their decision once it is revealed.

When is entry likely to dominate exit? Both terms in equation (31) split the same way: each is the probability of being on the relevant side of the cutoff w_L^{CE}, times the probability that the realized large-firm wage W_L actually clears (for entry) or fails to clear (for exit) the compensating wage $\Delta a W_S$. Entry therefore dominates exit when a large share of worker–posting pairs has a compensating wage above w_L^{CE}, and W_L has enough upper-tail mass to clear those compensating wages once revealed. Both conditions are easier to satisfy when risk aversion or wage dispersion is

high, since either lowers w_L^{CE} relative to workers' compensating thresholds: more pairs fall on the entry side of the cutoff, and a lower cutoff is easier for large-firm wages to clear.

Summary. Opacity pools heterogeneous large-firm postings at their certainty-equivalent value; transparency unpools them, generating entry at high-wage postings and exit at low-wage postings. The net effect raises the probability of applying to large firms precisely for workers whose amenity preferences placed the certainty equivalent out of reach, and lowers it for workers whose amenity preferences let the certainty equivalent clear their threshold on average. Because greater risk aversion and greater wage dispersion both lower w_L^{CE}, they push more worker–posting pairs onto the entry side of this cutoff and make entry into large firms more likely to dominate under transparency.