

BARRIERS TO ENTRY

DECOMPOSING THE GENDER GAP IN JOB SEARCH IN URBAN PAKISTAN

*Elisabetta Gentile, Nikita Kohli, Nivedhitha Subramanian,
Zunia Tirmazee, and Kate Vyborny*

NO. 707

December 2023

**ADB ECONOMICS
WORKING PAPER SERIES**

ADB Economics Working Paper Series

Barriers to Entry: Decomposing the Gender Gap in Job Search in Urban Pakistan

Elisabetta Gentile, Nikita Kohli, Nivedhitha Subramanian, Zunia Tirmazee, and Kate Vyborny

No. 707 | December 2023

The *ADB Economics Working Paper Series* presents research in progress to elicit comments and encourage debate on development issues in Asia and the Pacific. The views expressed are those of the authors and do not necessarily reflect the views and policies of ADB or its Board of Governors or the governments they represent.

Elisabetta Gentile (egentile@adb.org) is a senior economist at the South Asia Department, Asian Development Bank. Nikita Kohli (nikita.kohli@duke.edu) is a PhD student and Kate Vyborny (katherine.vyborny@duke.edu) is associate director of the DevLab@Duke and a research associate at the Department of Economics, Duke University. Nivedhitha Subramanian (nsubrama@bates.edu) is an assistant professor at Bates College. Zunia Tirmazee (zuniatirmazee@gmail.com) is an assistant professor at Lahore School of Economics.



Creative Commons Attribution 3.0 IGO license (CC BY 3.0 IGO)

© 2023 Asian Development Bank
6 ADB Avenue, Mandaluyong City, 1550 Metro Manila, Philippines
Tel +63 2 8632 4444; Fax +63 2 8636 2444
www.adb.org

Some rights reserved. Published in 2023.

ISSN 2313-6537 (print), 2313-6545 (electronic)
Publication Stock No. WPS230551-2
DOI: <http://dx.doi.org/10.22617/WPS230551-2>

The views expressed in this publication are those of the authors and do not necessarily reflect the views and policies of the Asian Development Bank (ADB) or its Board of Governors or the governments they represent.

ADB does not guarantee the accuracy of the data included in this publication and accepts no responsibility for any consequence of their use. The mention of specific companies or products of manufacturers does not imply that they are endorsed or recommended by ADB in preference to others of a similar nature that are not mentioned.

By making any designation of or reference to a particular territory or geographic area, or by using the term “country” in this publication, ADB does not intend to make any judgments as to the legal or other status of any territory or area.

This publication is available under the Creative Commons Attribution 3.0 IGO license (CC BY 3.0 IGO) <https://creativecommons.org/licenses/by/3.0/igo/>. By using the content of this publication, you agree to be bound by the terms of this license. For attribution, translations, adaptations, and permissions, please read the provisions and terms of use at <https://www.adb.org/terms-use#openaccess>.

This CC license does not apply to non-ADB copyright materials in this publication. If the material is attributed to another source, please contact the copyright owner or publisher of that source for permission to reproduce it. ADB cannot be held liable for any claims that arise as a result of your use of the material.

Please contact pubsmarketing@adb.org if you have questions or comments with respect to content, or if you wish to obtain copyright permission for your intended use that does not fall within these terms, or for permission to use the ADB logo.

Corrigenda to ADB publications may be found at <http://www.adb.org/publications/corrigenda>.

Note:

In this publication, “PRs” refers to Pakistan rupees and “\$” refers to United States dollars. ADB recognizes “China” as the People’s Republic of China.

ABSTRACT

Gender gaps in labor market outcomes persist in South Asia. An open question is whether supply or demand side constraints play a larger role. We investigate this using matched data from three sources in Lahore, Pakistan: representative samples of jobseekers and employers; administrative data from a job matching platform; and an incentivized resume rating experiment. Employers' gender restrictions are a larger constraint on women's job opportunities than supply-side decisions. At higher levels of education, demand-side barriers relax, allowing women to qualify for more jobs but at lower salaries. On the supply side, educated women become more selective in their search.

Keywords: gender, discrimination, job search, jobs platform, vacancies, applications

JEL codes: J16, J22, J23, R23

We are grateful for valuable feedback from seminar, conference, and workshop participants at the Asian Development Bank, Duke University, the Diversity and Human Capital Workshop at Exeter, the European Economic Association, the IZA Matching Workers and Jobs Online Workshop, the LACDev Conference, the ADE Conference at the Lahore School of Economics, the North East Universities Development Consortium Conference, the Maine Economics Conference, the Yale-ADBI Gender-Sensitive Economic Recovery Conference, and the Yale EGC workshop. We thank Livia Alfonsi, Emily Conover, Rohini Pande, Danila Serra, and Niharika Singh for comments. We thank the entire Job Talash team at CERP and Duke for excellent research assistance and programmatic support. We thank Erica Field and Rob Garlick for helpful comments on the draft and collaboration on the broader Job Talash program of research. We gratefully acknowledge funding for this project from ADB, and for the broader Job Talash research program from the GLM-LIC, 3IE, IGC, J-PAL, PEDL, and the National Science Foundation (SES 1629317). This research received ethics approval from Duke University (2019-0067). Previously presented as "A Leaky Pipeline: Decomposing the gender gap in job search in urban Pakistan."

1 Introduction

Vast gender gaps in employment, stemming from low levels of women's employment, persist in many low- and middle-income economies, particularly in South Asia, the Middle East, and North Africa (Addati *et al.*, 2016). A growing literature documents that supply side factors such as self-selecting into occupations that conform to gender identity, differing preferences for job attributes, and gendered social norms about time use constrain women's labor supply in these contexts. (Akerlof and Kranton, 2000; Cortes and Pan, 2017; Delfino, 2022; Dean and Jayachandran, 2019; Subramanian, 2021; Mas and Pallais, 2017; Fletcher *et al.*, 2018). However, demand-side rather than supply-side factors might form the binding constraint to women's labor for a larger fraction of the population in contexts with low female labor force participation (FLFP). Indeed, a smaller body of work demonstrates that demand-side factors such as firm gender preferences can contribute to gender gaps in employment (Kuhn and Shen, 2013; Goldin and Rouse, 2000; Hangartner *et al.*, 2021; Ozen *et al.*, 2019).

Despite the wealth of research on low female employment, we have limited evidence quantifying the relative size of demand and supply factors that give rise to this phenomenon. In addition, much of the existing evidence focuses on specific populations (e.g. educated women) or sectors (e.g. government work). Knowing whether supply or demand constraints are binding in the broader labor market is important to target policy to either supply-side interventions, such as provision of financial services to women or exposure to female role models (Field *et al.*; Ahmed *et al.*, forthcoming), or to demand-side interventions such as eliminating gender criteria in job ads, hiring quotas, wage subsidies, physical infrastructure such as women's toilets, or employer-based childcare (Card *et al.*, 2021; Miller *et al.*, 2022; Kuhn and Shen, 2023).

We address this gap in the literature by combining data from a job search platform and incentivized resume rating experiment from Lahore, Pakistan, to quantify the relative importance of demand-side and supply-side sources of the gender gap. We document

that at low education levels demand-side constraints are much larger than supply-side factors, but this demand-side gap in quantity of job opportunities closes as education levels increase and jobs become more “white-collar”.

Our empirical approach uses a novel combination of matched data from representative surveys, administrative data on job search, and experimental data to overcome four key challenges to quantifying the size of demand and supply side factors contributing to the the gender gap in the labor market. The first empirical challenge is that survey data from representative samples of households or firms can be used to quantify the realized gender gap in equilibrium, but do not allow the analyst to decompose how much of this gap comes from men’s and women’s willingness to supply labor versus firms’ demand for male and female labor. Second, alternative data sources such as job platform data allow researchers to observe the details of search activity by jobseekers and firms; however, such data often have limitations (Nomura *et al.*, 2017). They typically do not allow the researcher to observe jobseeker preferences directly, only to infer them from application choices, which are also influenced by other factors such as the vacancies available on the platform. Similarly, they typically do not allow the researcher to observe a well-defined choice set of vacancies observed by the jobseeker, making it difficult to disentangle whether the decision to apply to a given vacancy is a function of the characteristics of that vacancy or of search effort in browsing the platform (e.g. Belot *et al.* (2018); Wheeler *et al.* (2021); Jones and Sen (2022); Banfi *et al.* (2019); Matsuda *et al.* (2019)). Third, selection of both firms and workers into search and the use of job search platforms limits the extent to which results can be extrapolated to the population as a whole (Kureková *et al.*, 2015).

Addressing these first three challenges, our research partners at the Centre for Economic Research in Pakistan developed a new job matching platform, Job Talash, and offered it as a free service to representative listings of thousands of households and thousands of firms in a single urban labor market. We emphasize the importance of

studying the former group, economically inactive “latent workers,” who might be interested in working but are economically inactive due to lack of opportunities, who represent the population with the largest potential benefit from reductions in labor market barriers. This group is particularly important for understanding gender differences in settings such as Pakistan, where survey data suggest that a high fraction of the female population are latent workers. Female labor force participation in Pakistan was 21% in 2020 compared to a male labor force participation rate of 78% (International Labour Organization, 2019a,b); however, a quarter of women are not working but report they would like to work if they could find a suitable job (Field and Vyborny, 2016). The research design involved development, piloting, and refinement of a high-frequency job matching service that lists jobs and delivers information to respondents about them through text message and a call center. Job Talash allows us to precisely observe each step of job search activity on the supply and demand side (Field and Vyborny, 2022; Subramanian, 2021; Field *et al.*, 2023). The platform works by matching each jobseeker to open vacancies based on whether they satisfy minimal criteria set by the firm for the vacancy, and occupational preferences set by the jobseeker. The platform sends information to the jobseeker about all the vacancies that meet all criteria (we refer to such a pairing that satisfies all criteria as a “match”), and the jobseeker can decide whether to apply to each one. Thus, the platform generates high-frequency, detailed data on both the supply and demand sides of the labor market for millions of potential job matches between firm and respondent. The platform does not have a search function, which means that we observe the full set of vacancies that the jobseeker sees, and the full set of candidates sent to the firm. Because we provide information to both sides, we observe exactly the same information as both sides of the market up to the point of an interview. The representative recruitment of both jobseekers and firms (rare in the literature (Kureková *et al.*, 2015)) and the rich administrative data observed on the

platform provide us unique leverage to help pinpoint the supply-side versus demand-side constraints to women's versus men's job search.

The fourth challenge in quantifying supply versus demand side constraints is that even if the initial sample of jobseekers and firms is representative, when observing downstream outcomes in the job search process such as interviews and hires, selection problems arise again, because male and female candidates who choose to apply for a given vacancy may differ systematically from each other and from non-applicants. To address this challenge, we combine the administrative data with an incentivized resume rating (IRR) experiment which we conducted with firms in the Job Talash sample (Kessler *et al.*, 2019). We show employers on the platform a series of pairs of curriculum vitae (CV) and in each pair ask the respondent to select the one that they would be most likely to hire, with the incentive that this could help inform the applicant pool sent to them through the Job Talash platform. CVs for this exercise were constructed using the actual job applicant data from the Job Talash pool, making them a realistic representation of the candidates the firm might see on the platform; we randomly varied the gender of the applicant on the CV to identify firm preferences over gender, holding constant potential confounders such as differences in levels of education and experience between men and women in the pool, and differential selection into application.

Our first key finding is that gender gaps in employment are greater in magnitude than gender gaps in search. This complements the literature showing that women's labor supply is elastic to the introduction of good jobs even in the South Asian context, where overall female employment rates are low (Jensen, 2012; Heath and Mobarak, 2015). Women are 89% less likely than men to be working at baseline; however, they are only 53% less likely to complete sign-up for the Job Talash platform, an investment of time in the telephone based sign-up process which indicates willingness to search. The gender gap in both work and willingness to search narrows as education levels rise. At higher education levels, the gender gap in completing the signup process falls by 65%. These

findings suggest that many women, particularly educated women, are “latent workers”—pointing to key constraints on the labor demand side.

Our second key finding is that for less educated jobseekers, firm gender criteria, an entirely demand-side constraint, are more binding for women than men, and are also a larger constraint than supply-side decisions. Women in our setting are 53% less likely than men to satisfy the explicit gender requirements for any given vacancy. These patterns persist even when we restrict to vacancies where the individual met the education and experience criteria, and expressed interest in the occupation: demand-side criteria are the binding constraint on opportunities available to women. In fact, in the set of vacancies where individuals satisfied all basic criteria and were eligible to apply, women apply at a higher rate than men, overall.

Our third key finding is that the demand side gap in quantity of job opportunities substantially closes as education levels rise, while on the supply side women become more selective. The gender gap in satisfying the gender criteria for a position shrinks by 70% for the minority of women with secondary education and effectively disappears for the third of women with a tertiary education. We document this through both the administrative data and the IRR experiment. We find that firms’ gender criteria and the educational requirements of the job are mirrored; vacancies with “blue collar” characteristics such as manual labor and longer and late work hours are more likely to exclude women and more common among jobs with low education requirements, even conditional on industry and occupation fixed effects. Additionally, firms’ gender criteria and the education level they seek to hire reflect existing infrastructure at the firm: firms that have restrooms or a separate prayer space for women are both more likely to be willing to hire women and more likely to be hiring at a high education level. This connects to a broader literature pointing to the role of non-wage characteristics of firms and vacancies in gender gaps in the labor market (Mas and Pallais, 2017; Flory *et al.*, 2015; Field and Vyborny, 2022; Goldin and Katz, 2016; Chiplunkar and Goldberg, 2021; Miller

et al., 2022). Strikingly, among those with a tertiary education, women are more selective than men in their job search. At this high education level, women are slightly less likely than men to have selected the occupation of a given vacancy, and are slightly less likely to apply to a vacancy. But this is likely driven by differences in the quality of vacancies by gender; indeed, we find that among those with a tertiary education, women are more likely than men to qualify for the vacancies at the lowest quintile of the salary distribution.

We advance the literature in two key ways. First, our novel combination of matched data allow us to separately quantify the role of demand and supply side decisions in the gender gap in job opportunities. Second, we are able to document these results in a more representative sample than most studies.

Papers that use data from job search and matching platforms increasingly are able to study both sides of the labor market simultaneously, and do find differences in job search by gender. Women in Chile, Nigeria, and Denmark, respectively, have been shown to be more selective than men in their job applications, to be more qualified for the jobs they do apply to conditional on applying, and to apply to lower-wage jobs (Banfi *et al.*, 2019; Archibong *et al.*, 2022; Fluchtmann *et al.*, 2021). Like other existing platform-based studies, these papers include a sample of jobseekers who take the initiative to sign up to the platform. In contrast, we start with a representative listing and approach all adults to sign up for the platform. This allows us to characterize selection into the platform sample and quantify its importance; and our signup process enrolls many women with “latent” labor supply who are not searching at baseline, who would not appear in standard platform samples. These studies do not observe the gender preferences of employers, and can only infer them from interview decisions, which are a result of application decisions and selection of which jobseekers apply. In contrast, our platform design and IRR experiment allow us to quantify the extent to which employer gender preferences affect the opportunities available to women.

This links to a literature that has documented explicit firm gender criteria in labor markets such as the People's Republic of China (PRC) and India. Gender criteria are common on job ads on internet job boards in the PRC, and also in part determine the gender mix of applicants (Kuhn and Shen, 2013; Kuhn *et al.*, 2020). Such gender criteria explain some of the gender wage gap on a job portal in India (Chaturvedi *et al.*, 2022). Going further, Kuhn and Shen (2023) and Card *et al.* (2021) document that when policy changes led to gender criteria being removed in the PRC and Austria, the gender composition of applications increased, without sacrificing match quality. Relative to these papers, we are able in a single setting to observe the exact vacancies that each jobseeker is matched to, observe whether the jobseeker satisfies the minimum requirements for each vacancy, and additionally observe interviews as an outcome through the administrative data and hypothetical interviews through the incentivized resume rating experiment. This allows us to compare the relative magnitude of gender criteria (demand-side) versus jobseeker (supply-side) decisions. Furthermore, we advance the literature by starting with a sample representative of a broader urban labor market, not restricted by employment or search status.

There is an established theoretical basis for a U-shaped relationship between education levels and women's labor supply (Gaddis and Klasen, 2014; Goldin, 1995). Empirical work shows that the pattern of the relationship between educational attainment and labor force participation varies greatly for women across low- and middle-income countries around the world (Aromolaran, 2004; Cameron *et al.*, 2001; Klasen, 2019). Within India, a context similar to ours with low female labor force participation, women's own education is positively correlated with labor force participation in urban settings but the opposite in rural settings (Afridi *et al.*, 2017; Klasen and Pieters, 2015).

An ancillary contribution of our study is methodological. We combine large-scale, representative and administrative data from real choices in the field (the signup and platform data) with a controlled, lab-in-field style experiment (the incentivized resume

reporting experiment). This builds on approaches combining the advantages of large-scale, naturalistic field data with small, controlled lab style experiments (Garlick *et al.*, 2023; Cortes *et al.*, 2022).

We describe the context of our study, including data collection and the IRR experiment design in Section 2. We report results on gender gaps in work and job search platform sign-up in Section 3, on gender gaps in outcomes on the platform in Section 4, and how vacancy characteristics are correlated with demand-side gender barriers in Section 5, all using the administrative data. We present results from the IRR experiment in Section 6. Section 7 concludes.

2 Context and Data

Our study is set in Pakistan, the world's fifth most populous country (United Nations Population Division, 2023). Male labor force participation across Pakistan is 78%, while female labor force participation is much lower near 20% (International Labour Organization, 2019b,a). These figures are similar to countries in the Middle East, North Africa, and South Asia (International Labour Organization, 2019b,a). This project takes place in Lahore, the second largest city in Pakistan, with a population of about 11 million. The male employment rate is 83% while the female employment rate is just under 10% (author calculations from the Pakistan Labor Force Survey). In contrast with the gender gap in employment, women and men in this setting have similar levels of educational attainment. About 71% have at most primary education, 12% have at most secondary education, and 15% have tertiary education.

2.1 Job Talash platform

We use administrative data from Job Talash, a free job search and matching platform developed by our research partners at the Center for Economic Research in Pakistan, to serve the district of Lahore. The team began with a representative household listing across Lahore, fielded between October 2016 and September 2017, which yielded a starting

sample nearly identical to the population of Lahore, in terms of age, gender, education, and employment rates. This is shown in the comparison of columns 1 and 2 in Appendix Table A.1.¹

The area covered in the listing for both employers and households is a single metropolitan commuting area; the mean distance between jobseekers and firms in our sample is 11 kilometers. The representative household listing comprised approximately 180,000 individuals. From here, we are able to decompose gender gaps at every stage of the search process, allowing us to isolate supply-side versus demand-side constraints to women's versus men's employment.

In the representative household listing, the service first offered every adult in the household free sign-up onto the Job Talash platform. Job Talash followed up by telephone and gathered the information about their work history and education to help them construct a CV used for job applications through the platform. At the stage of signing up for the service and constructing the CV, individuals specified the occupations in which they wanted to search for jobs. Nearly 10,000 individuals registered for the platform and constructed a CV to facilitate job applications through the platform.

Job Talash also conducted a representative listing of firms across Lahore. The team listed a representative sample of approximately 10,000 firms across the metropolitan area, using a cluster-randomized selection of Enumeration Blocks followed by listing of all firms in each selected block. A team of enumerators presented the Job Talash service to firms, offering them the opportunity to enroll to list vacancies immediately or later.

Appendix Table A.2 examines selection into platform use on the employer side. 3.4% of firms approached in the listing signed up and posted at least one ad over the course of the study. Firms that did so are larger in terms of number of employees, frequency of recruitment, and physical size. They are also more likely to have any women employees at baseline: Only 8% of firms that did not post an ad had any female employees at

¹The appendix tables and figures are available at <http://dx.doi.org/10.22617/WPS230551-2>.

baseline, compared to 21% among those who did post. Among those who reported details of physical infrastructure, firms that use the platform are also more likely to have a separate toilet and separate prayer space for women. These patterns suggest that our results on demand-side barriers are a lower bound.

Each ad posted through the platform specifies the education and experience required for the position. This process generated 758 ads on the platform, placed between August 2017 and September 2022. Given the nature of the labor market in Lahore, firms also could specify if the vacancy is open only to men (59.6% ads), women (14.9% ads), or open to any gender (25.5% ads). This is a phenomenon observed in labor markets in many countries (Kuhn and Shen, 2013; Card *et al.*, 2021; Chaturvedi *et al.*, 2022).

The platform matches individuals to open vacancies, based on four criteria. If all four criteria were satisfied, then the platform sends the “match” to the individual, for the individual to decide whether to apply. The first criterion is whether the vacancy is within the set of occupations that the individual wanted to be matched with. The second and third criteria are whether the individual matches the minimum education and experience qualifications for the vacancy, as set by the firm. The fourth criterion is whether the individual satisfies the gender criteria for the position, if the firm imposed such a restriction. Individuals receive text messages for each of these job ads that satisfy all four criteria (“matches”); messages are sent in batches, approximately once a month. See Appendix Figure A.1 for a sample text message. The text messages contain the job title, firm name, firm location, and salary of each match, along with the deadline to apply. Jobseekers only learn about vacancies to which they match, as the platform does not have a search function. Participants can ask to pause or stop receiving job ads at any time. For each job ad, the individual decides whether to apply. The platform is completely free and calls back prospective applicants, so the monetary cost of application is minimal (a maximum of PRs5.00 or \$0.03, in purchasing power parity term, less than 1% of a day’s earnings at minimum wage), so financial cost is unlikely to affect

gender and education patterns in search on platform. If the jobseeker chooses to apply, Job Talash sends their CV to the firm; the firm decides whether to invite the applicant for an interview. The platform calls each firm a few weeks after the CVs are delivered and follows up as needed to confirm which applicants were interviewed.

Crucially, we observe choices by both the individual and the firm separately. This distinguishes our data from typical labor force data in which the researcher only observes an equilibrium outcome, such as the occupation in which a woman is employed. The latter could be an outcome of the woman's preferences for a certain occupation, employers' preferences to hire women into that occupation, or both. In contrast, we observe the constraints placed by both sides on their search: occupations that individuals select to receive as matches; qualifications that the individual has and the firm requires; and the firm's explicit gender preferences. We construct a dataset of every possible jobseeker-job ad dyad within the potential list of occupations offered to the jobseeker's broad education level, regardless of whether the dyad actually satisfied all of the criteria placed by both the jobseeker and firm.² Since we observe the firm's and the jobseeker's criteria separately, we are able to observe whether each dyad satisfied all jobseeker and firm criteria and was shown to the jobseeker (referred to as a "match") for them to decide whether to apply. For dyads that do not meet all criteria, and thus the jobseeker did not see the vacancy to decide whether to apply, we observe whether this was due to the individual not meeting the firm's constraints, vice versa, or both. We can then observe for each match sent whether the jobseeker applies, and ultimately is invited for an interview. We also observe all information that the firm and jobseeker have about each other up until the point of an interview.

²To reduce the length of the signup process, jobseekers with less than a 10th grade education are offered a list of occupations appropriate to a low-education jobseeker pool, while jobseekers with above a 10th grade education are offered a high-education occupation list (occupation lists are provided in Appendix Table A.6). We construct dyads only within occupations that jobseekers could select. Our main results are robust to instead including all the infeasible dyads; results available on request.

This dataset contains over 3.5 million jobseeker-job ad dyads, of which 18.6% result in a match sent to the jobseeker. We use this to further pinpoint patterns in men's and women's job search behavior and success.

2.2 Incentivized Resume Rating Experiment

The administrative data has the advantage of starting out from a representative sample, and showing us gender gaps at fine-grained decision-points through the job search process. However, at the stage of the interview decision, the relationship between applicant gender and the outcome is affected by self-selection into application. To address this, we combine this data with an incentivized resume rating experiment, following methods developed by Kessler *et al.* (2019), in order to isolate employer preferences for gender versus other CV characteristics and thus shed light on what patterns might drive gender gaps on the demand side.

We implemented this experiment with employers signing up for the Job Talash service over a part of the sign-up period (January 2019 to December 2020); an enumerator presented the respondent with a series of three pairs of CVs, and advised the respondent that while the choices are hypothetical, their answers could help inform the applicant pool for future ads they place on the platform.³

CVs for this exercise were constructed using the actual jobseeker CV data from the Job Talash pool. We selected a random sample of 176 unique CVs to span educational levels ranging from no formal education to a Master's degree; and with no more than five years of work experience, to avoid including CVs that were too specialized in a particular field to be relevant for the broad based pool of employers in our sample. We stratified the sample by each combination of the level of education (less than secondary, secondary, or tertiary) and years of experience (0–2 years or 3–5 years) regardless of gender. We then randomly

³The script used to present this exercise is as follows: "We will now show you two sample CVs. Take your time to browse through them. Out of the two, choose a CV which you will hire if these were the candidates presented to you. This choice will help us determine which CVs to send you for your opening."

selected pairs such that within each pair, the two CVs are no more than one education level apart, and no more than two years of experience apart. Personal information such as applicant name and address was removed. CVs were assigned fictitious names out of a list of common names based on the gender of the applicant. Extremely common Pakistani names such as Muhammad Ali or Ayesha Ahmed (the local equivalents of John Smith and Jane Doe) were used, to avoid any risk that the fictitious CV in some way is associated with a real individual. We randomly selected characteristics including gender, educational institution and standardized exam scores to be swapped between the two CVs in a pair to ensure exogenous variation in these characteristics. We used a series of independent randomizations for each trait to determine whether they would be swapped between the two CVs in the pair; thus a pair may have had all three traits swapped, some of them, or none of them.

Appendix Figure A.2 is an illustration of the swapping exercise for a CV pair, and Appendix Table A.3 summarizes the design. Because the applicant gender is randomized across CVs, traits are balanced across male- and female-named CVs, as shown in Appendix Table A.4.

Some firms dropped out of the survey before reaching this module; Appendix Table A.5 compares characteristics of firms that participated versus those that did not participate in the IRR. Participation in the IRR reflects interest in the platform and posting a job on the platform. During the period this experiment was active, 392 firms from the representative listing signed up; due to partial survey nonresponse by firms, 189 of these firms responded to the CV choice module. We have a total of 447 binary choices in the full sample. For this analysis, we drop 232 binary choices in which both candidates are the same gender; thus the resulting estimation sample consists of 430 CVs (215 binary choices) shown to 136 firms.

3 Gender Gaps in Work and Job Search Platform Sign-up

We begin by estimating gender gaps in employment and willingness to search at the individual level, using the representative listing and administrative data. We regress the outcome on an indicator for whether the individual is female, and report heteroskedasticity-robust standard errors. Panel A of Table 1 reports the results of the estimation and the gender gap in percentage terms: the ratio of the coefficient on the indicator for female against the constant term. Women are 89% (63 percentage points or pp) less likely to be working (col 1). In column 2, we show the gender gap in indicating interest in Job Talash at the time of household listing. A key respondent was interviewed for the household, with 83% of respondents female. For adult household members who were not present at the time of the interview, the respondent was asked to indicate whether she thought the individual would be likely to be interested. The gender gap in whether the respondent thinks the individual would be interested is 32% (9.5 pp). This is a far smaller gap than the gap in employment. Since this outcome in column (2) is not necessarily directly from the individual, in column (3) we consider whether individuals complete the sign-up process. There is still a large gender gap, but again smaller than the gender gap in employment. Women are 53% (3.9 pp) less likely than men to complete the sign-up process (column 3). Baseline survey data from individuals who do sign up reveal that 63% of women who sign up are neither working nor searching at baseline as compared to 37% of men, again reinforcing the idea of latent female labor supply.⁴

In Table 1 Panel B, we examine how these patterns shift with education. We estimate:

$$Y_i = \beta_0 + \beta_1 F_i + \beta_2 F_i \times S_i + \beta_3 F_i \times T_i + \beta_4 S_i + \beta_5 T_i + \varepsilon_i \quad (1)$$

⁴The absolute magnitude of sign-up is 7.4% for men, and 3.5% for women. While sign-up does not have a monetary cost, there is a time cost to providing information to construct a CV.

Overall, gender gaps close as education levels rise. The gender gap in work closes by 20% (13.4 pp) for women with a tertiary education compared to women with less than a secondary education. The gender gap in interest closes completely with tertiary education, and and sign-up to the job search platform closes by 66% (2.9 pp).

Table 1: Gender gaps in work and interest in search

Panel A: Overall			
	(1)	(2)	(3)
	Working at baseline	Interested in Job Talash	Completed signup
β_1 : Female _i	-0.632*** [0.002]	-0.095*** [0.002]	-0.039*** [0.001]
β_0 : Constant	0.713*** [0.001]	0.302*** [0.002]	0.074*** [0.001]
β_1/β_0	-0.89	-0.32	-0.53
N	182,491	182,491	182,491
Panel B: By Education levels			
	(1)	(2)	(3)
	Working at baseline	Interested in Job Talash	Completed signup
β_1 : Female _i	-0.669*** [0.002]	-0.121*** [0.002]	-0.044*** [0.001]
β_2 : Female _i × Secondary Ed _i	0.131*** [0.006]	0.039*** [0.006]	-0.005 [0.004]
β_3 : Female _i × Tertiary Ed _i	0.134*** [0.005]	0.126*** [0.006]	0.029*** [0.003]
β_4 : Secondary Ed _i	-0.111*** [0.005]	0.026*** [0.005]	0.027*** [0.003]
β_5 : Tertiary Ed _i	0.016*** [0.004]	-0.011** [0.004]	0.012*** [0.003]
β_0 : Constant	0.724*** [0.002]	0.301*** [0.002]	0.069*** [0.001]
P-value: $\beta_1+\beta_2=0$	0.00	0.00	0.00
P-value: $\beta_1+\beta_3=0$	0.00	0.28	0.00
N	182,491	182,491	182,491

Notes: The unit of observation is the individual in the household listing. Education variables are mutually exclusive and exhaustive indicators. Primary Education is the omitted category. Primary Education includes no education, completed primary or secondary (0-10 years). Secondary education refers to completed higher secondary (12 years). Tertiary education refers to completed tertiary education (16 years or more). The constant is the mean for males (Panel A) or for males with a primary education (Panel A). Robust SEs in brackets. * $p < .1$, ** $p < .05$, *** $p < .01$.

Source: Authors' estimates.

4 Gender Gaps In Outcomes on Platform

We now turn to gender gaps arising after selection of individuals and firms onto the platform. To decompose these gaps, we construct all possible jobseeker (i) and vacancy (j) dyads; including all individuals who completed the Job Talash sign-up process. We begin by analyzing all possible dyads and examining whether the individual's qualifications, the individual's choices, and the firm's requirements resulted in the individual and vacancy being matched, such that the individual sees the vacancy ad and can make a decision of whether to apply. This allows us to decompose gender gaps on supply-side and demand-side margins.⁵ We estimate gender gaps in whether each dyad satisfied each of the criteria such that the vacancy became available for the individual to decide to apply: whether the individual selected the occupation of the vacancy, whether the individual met the education, experience criteria, and conditional on being matched to the position, whether the individual chose to apply. We also estimate gender gaps in whether the individual met the vacancy's explicit gender criteria, and whether the individual was selected to interview, conditional on applying to the position. The unit of observation is now the jobseeker-job ad dyad. For all jobseeker-job ad dyad regressions we cluster standard errors on the jobseeker and job ad.

Overall, 18.6% (657,312) of dyads satisfy all four criteria and thus convert to a "match," a job ad sent to the jobseeker, that the jobseeker can see and decide whether to apply. 31.4% (1,112,609) of dyads are not shown to the jobseeker only because they are in occupations that the jobseeker did not select, and 17.5% (619,454) of dyads are not shown to the jobseeker only because the jobseeker did not meet the firm's education, experience, or gender criteria. 32.5% (1,152,557) are not shown because neither condition was met.

⁵Individuals with less than a secondary education and those with at least a secondary education were shown different sets of approximately 20 occupations to choose from as shown in Appendix Table A.6. In our main analysis we construct dyads only from occupations that were in the potential option set of occupation list shown to the jobseeker (e.g. we do not construct a dyad for a college graduate and a janitor job, as this occupation was not in the option list offered to this jobseeker); our central results are quantitatively and qualitatively similar if we include such dyads in the analysis.

In Table 2 we examine gender gaps in each component of the matching process. The first takeaway is that demand-side barriers, namely explicit firm-imposed gender criteria, are quantitatively the largest barrier preventing a jobseeker-vacancy dyad from meeting all of the criteria and converting to a match shown to the jobseeker, who can then decide whether to apply. We see this in columns 1–4 of Panel A in Table 2. Women and men have statistically indistinguishable probabilities of choosing the occupation of a given vacancy (col 1), which is a purely supply-side decision.⁶ There is no statistically significant difference in the likelihood of meeting educational requirements for a given vacancy between women and men (col 2). However, women are 20% less likely (17.5 pp; col 3) to have met the experience requirements for the vacancy. The latter is unsurprising since low employment rates for women overall in the context are consistent with women having less work experience compared to men. In column 4, the last column that explores criteria under which a dyad would or would not be shown to a jobseeker, women are 53% (45.8 pp) less likely than men to meet the gender criteria for a vacancy. In Panel B, columns 1–4, of the same table, we see that women at higher levels of education are more selective in choosing occupations. Among jobseekers with tertiary education, the gender gap in qualifying by education favors women, and the gender gaps in qualifying by experience and gender narrow.

All together, women are 59% less likely (13.2 pp; Table 2 col 5) to be matched to a vacancy, meaning that they satisfy all four criteria and are shown the vacancy/able to choose whether to apply. Again, this gender gap closes as education levels rise.

The second major takeaway is that women with less than a secondary education are more likely to apply to a vacancy they have been shown than men, but that this reverses at higher education levels. Conditional on satisfying all four criteria and being shown the vacancy, women are 34% (0.2 pp) *more* likely to apply to a vacancy than men (Table

⁶While women and men choose different occupations, this does not change the average number of vacancies to which they match from the firms from the representative listing who choose to post vacancies on the platform.

2, Panel A, column 6). In Panel B of Table 2 (column 6) we see that among those with less than a secondary education, women are more likely than men to apply to any given vacancy, but that this pattern reverses at higher levels of education.

In column 7 of Table 2, we see that overall there are no statistically significant gender gaps in being invited for an interview, among vacancies to which the jobseeker submitted an application. Gender gaps in interview selection do not vary detectably by education, but these results are imprecise (Panel B). In addition, the sample of applicants may also differ systematically in other characteristics between men and women. This motivates the use of the IRR experiment in Section 6.

As a robustness check, we estimate Panel A, columns (2) and (4) for the full representative sample, for whom we have education and gender information. Results are reported in Appendix Table A.7. In the representative sample, the gender gap in qualifying by education is small and favoring men at 2% (1.3 pp). Unlike in the analysis sample, this gap is statistically significant, though as in the analysis sample, the magnitude is small. The gender gap in qualifying by gender in the representative sample is 53%. These results are nearly identical to the main results among jobseekers who signed up for the platform.

Table 2: Supply and demand side gender gaps

Panel A: Overall							
	(1) Selected occup.	(2) Qualify educ	(3) Qualify exper.	(4) Qualify gender	(5) Matched	(6) Apply matched	(7) Interview apply
β_1 : Female _i	-0.006 [0.009]	-0.001 [0.005]	-0.175*** [0.008]	-0.458*** [0.028]	-0.132*** [0.010]	0.002** [0.001]	0.022 [0.023]
β_0 : Constant	0.361*** [0.007]	0.799*** [0.010]	0.866*** [0.006]	0.864*** [0.013]	0.225*** [0.007]	0.006*** [0.000]	0.071*** [0.012]
β_1/β_0	-0.02	-0.00	-0.20	-0.53	-0.59	0.34	0.31
N	3,541,932	3,541,932	3,541,932	3,541,932	3,541,932	606,579	3,548
Panel B: By Education levels							
	(1) Selected occup.	(2) Qualify educ	(3) Qualify exper.	(4) Qualify gender	(5) Matched	(6) Apply matched	(7) Interview apply
β_1 : Female _i	0.002 [0.011]	-0.028*** [0.006]	-0.211*** [0.011]	-0.623*** [0.030]	-0.179*** [0.010]	0.007*** [0.002]	0.021 [0.035]
β_2 : Female _i × Secondary Ed _i	-0.015 [0.019]	0.013 [0.013]	0.123*** [0.017]	0.439*** [0.041]	0.119*** [0.016]	-0.009*** [0.002]	0.030 [0.044]
β_3 : Female _i × Tertiary Ed _i	-0.038** [0.018]	0.049*** [0.009]	0.121*** [0.016]	0.558*** [0.051]	0.151*** [0.021]	-0.012*** [0.002]	0.000 [0.044]
β_4 : Secondary Ed _i	0.033*** [0.012]	0.013 [0.017]	-0.099*** [0.013]	-0.084*** [0.021]	-0.032*** [0.011]	0.004*** [0.001]	-0.041** [0.018]
β_5 : Tertiary Ed _i	0.013 [0.015]	0.134*** [0.014]	-0.043*** [0.012]	-0.112*** [0.027]	-0.012 [0.014]	0.004*** [0.001]	-0.036* [0.021]
β_0 : Constant	0.356*** [0.008]	0.782*** [0.012]	0.882*** [0.006]	0.886*** [0.014]	0.230*** [0.008]	0.005*** [0.000]	0.084*** [0.017]
P-value: $\beta_1 + \beta_2 = 0$	0.43	0.24	0.00	0.00	0.00	0.24	0.09
P-value: $\beta_1 + \beta_3 = 0$	0.01	0.00	0.00	0.14	0.14	0.00	0.44
N	3,541,932	3,541,932	3,541,932	3,541,932	3,541,932	606,579	3,548

Notes: The unit of observation is a jobseeker-job dyad, for all jobseekers who sign up and all jobs posted on the platform. The dependent variable in column 4 is an indicator for whether the jobseeker meets any gender requirements for the vacancy; it equals 1 for all jobseekers for vacancies that are open to both men and women. The dependent variable in column 5, "matched," is an indicator for whether the algorithm identified job j as a potential match for jobseeker i and sent the vacancy ad to the jobseeker; this occurs if and only if the jobseeker selected the relevant occupation category (column 1), meets the minimum education and experience qualifications (columns 2-3) and meets the firm's gender restrictions (column 4). Education variables are mutually exclusive and exhaustive indicators. Primary Education is the omitted category. Primary Education includes no education, completed primary or secondary (0-10 years). Secondary education refers to completed higher secondary (12 years). Tertiary education refers to completed tertiary education (16 years or more). The constant is the mean for males (Panel A) or for males with a primary education (Panel A). Robust SEs in brackets, two-way clustered by jobseeker and vacancy. * $p < .1$, ** $p < .05$, *** $p < .01$.

Source: Authors' estimates.

In the next set of results, we delve further into the firm-side hiring constraints. We saw in Table 2 that women are 53% less likely than men to satisfy a firm's explicit gender criteria, with much smaller gaps coming from women's occupational choices and qualifications. But it is possible that to some extent firms' ex ante gender restrictions screen out women from jobs they would not have qualified for or selected in any case.

To what extent do firms' explicit gender restrictions form the *binding* constraint on the number of opportunities available for women to apply to? We explore this in Table 3. Column 1 of Table 3 replicates column 4 of Table 2: the dependent variable is an indicator for whether the jobseeker meets the gender requirements of the vacancy. In column 2, we restrict the sample to dyads where the individual met the vacancy's education requirements; and in column 3, dyads in which the individual met both education and experience requirements. Finally, in column 4, we restrict to dyads where the individual met the vacancy's education and experience requirements, and additionally, the individual selected the occupation of the vacancy for matching. In Panel A, column 2, we see that women are 55% (48.4 pp) less likely than men to meet the firm's gender criteria, among dyads where they qualify based on education; in column 3, we see that women are 52% (45.5 pp) less likely than men to meet the firm's gender criteria, among dyads where they qualify based on education and experience. In column 4, we see that women are 49% (42.7 pp) overall less likely than men to meet the firm's gender criteria, for dyads for which they qualify based on education and experience, and additionally selected into the occupation of the vacancy. In all of these cases, the firm's gender criteria are the binding constraint on potential matches that female jobseekers would otherwise have received. In Appendix Table A.7, column (3), we construct the gender gap in qualifying based on gender, for the subset of matches where the individual would have qualified based on education, analogous to column (2) in Table 3, as a robustness check using the full representative sample, for whom we have education and gender information. Here, we see that among dyads that would have qualified for the posting based on education, women are 69% less likely to qualify based on gender than men. Across the board, it is the firm side criteria rather than education, experience or occupation preferences that restrict women's access to these potential job opportunities. This suggests again that demand-side gender criteria are a key constraint.

In Panel B, we break down these patterns by education level. Within the set of dyads where they qualify for the vacancy based on education and experience (col 3), women with less than secondary education are 65.4 pp less likely to qualify based on firm gender criteria compared to men with less than secondary education. At the tertiary level, the gender gap is nearly eliminated (col 3). When the sample is restricted to dyads in which the individual also selected the occupation of the vacancy, the gender gap closes completely (column 4).

Table 3: Role of firm-side gender restrictions in gender gap


Panel A: Overall				
	Qualify based on gender			
	(1)	(2)	(3)	(4)
β_1 : Female _i	-0.458*** [0.028]	-0.484*** [0.028]	-0.455*** [0.031]	-0.427*** [0.035]
β_0 : Constant	0.864*** [0.013]	0.881*** [0.012]	0.878*** [0.013]	0.878*** [0.014]
β_1/β_0	-0.53	-0.55	-0.52	-0.49
Sample	Full Sample	Qualify educ	Qualify educ+exp	Qualify educ+exp +select occp
N	3,541,932	2,827,515	2,317,189	841,114
Panel B: By Education levels				
	Qualify based on gender			
	(1)	(2)	(3)	(4)
β_1 : Female _i	-0.623*** [0.030]	-0.680*** [0.029]	-0.654*** [0.032]	-0.643*** [0.037]
β_2 : Female _i × Secondary Ed _i	0.439*** [0.041]	0.454*** [0.042]	0.440*** [0.046]	0.500*** [0.053]
β_3 : Female _i × Tertiary Ed _i	0.558*** [0.051]	0.615*** [0.051]	0.605*** [0.055]	0.668*** [0.062]
β_4 : Secondary Ed _i	-0.084*** [0.021]	-0.101*** [0.020]	-0.106*** [0.022]	-0.131*** [0.026]
β_5 : Tertiary Ed _i	-0.112*** [0.027]	-0.138*** [0.028]	-0.144*** [0.029]	-0.170*** [0.033]
β_0 : Constant	0.886*** [0.014]	0.909*** [0.013]	0.906*** [0.014]	0.914*** [0.013]
P-value: $\beta_1+\beta_2=0$	0.00	0.00	0.00	0.00
P-value: $\beta_1+\beta_3=0$	0.14	0.16	0.31	0.65
Sample	Full Sample	Qualify educ	Qualify educ+exp	Qualify educ+exp +select occp
N	3,541,932	2,827,515	2,317,189	841,114

Notes: The unit of observation is a jobseeker-job dyad, for all jobseekers who sign up and all jobs posted on the platform. The dependent variable is an indicator for whether the jobseeker meets any gender requirements for the vacancy; it equals 1 for all jobseekers for vacancies that are open to both men and women. The sample in Column 1 includes all dyads, as in Table 2; in Column 2 includes only those in which the jobseeker qualified for the vacancy in terms of education; in Column 3 includes only those in which the jobseeker qualified for the vacancy in terms of both education and experience; and in Column 4 includes only those who qualified and also selected the occupation (i.e. met all other criteria for being “matched” to the vacancy other than the gender restriction). Education variables are mutually exclusive and exhaustive indicators. Primary Education is the omitted category; it includes no education, completed primary or secondary (0-10 years). Secondary education refers to completed higher secondary (12 years). Tertiary education refers to completed tertiary education (16 years or more). The constant is the mean for males (Panel A) or for males with a primary education (Panel A). Robust SEs in brackets, two-way clustered by jobseeker and vacancy. * $p < .1$, ** $p < .05$, *** $p < .01$.


Source: Authors' estimates.

The analysis in Tables 2–3 focuses on the number of potential job matches for which a jobseeker qualifies. What about the quality of those opportunities? The firm-side gender criteria may give rise to a gender wage gap (Chaturvedi *et al.*, 2022; Nomura *et al.*, 2017; Matsuda *et al.*, 2019). In Figure 1 we repeat the dyadic analysis above, splitting the sample by jobseeker education level and quintiles of salary distribution within each education level. The dependent variable is an indicator for whether the jobseeker meets the vacancy’s gender requirements. The results are striking. At the primary education level (Panel A), the gap by gender is large, mirroring the results in Table 2; as the salary level of the vacancy rises, the gap widens. This pattern changes dramatically as the jobseeker’s education level rises (Panels B to C). Overall, the gap in gender restrictions shrinks, again mirroring the results in Table 2. At the tertiary education level, for vacancies with lower posted salaries, women actually meet the gender criteria for *more* jobs than men do; this difference is statistically significant at the 5% level. The result reverts to the previous pattern at higher salary levels, with men qualifying for significantly more vacancies based on gender alone. While the gender gap in the *quantity* of opportunities due to firm-side gender constraints closes with higher education, there is still a gender gap in quality which does not disappear. In contrast, the gender gap in satisfying education and experience criteria are smaller and stable across salary levels (Appendix Figure A.3).

Figure 1: Qualify gender across salary quintiles, by education



Continued on the next page



Notes: This figure shows the results of repeating the dyad level analysis in Table 2, Panel A, with separate estimations on samples for each education level and within level, each quintile of the posted vacancy salary. The unit of observation is jobseeker-vacancy dyad. The outcome variable is an indicator for whether the jobseeker meets gender criteria for the vacancy. Robust SEs two-way clustered by jobseeker and vacancy; 95% confidence intervals shown. Stars shown alongside coefficients denote P-values testing equality between female and male jobseekers. * $p < .1$, ** $p < .05$, *** $p < .01$.

Source: Authors' estimates.

5 Vacancy Characteristics Correlated with Demand-Side Gender Barriers

The analysis in Section 4 demonstrates that demand-side *ex ante* gender restrictions play a dominant role in affecting the opportunities available for women. This motivates further investigation of vacancy and firm characteristics.

As shown in Appendix Table A.2, our listing of employers has substantial gender segregation at the firm level. As a result, 70% of vacancies in the sample are at firms that have no female employees at baseline, while only 10% are at firms with majority or all women employees. In Appendix Table A.8 we explore how this firm-level gender segregation mediates the *ex ante* gender restrictions and their impact on the number of opportunities available to women. The outcome variable is again an indicator for whether

the individual qualified for the firm's explicit gender criteria; the unit of observation is again the individual-vacancy dyad. We find that the gender criteria reflect existing gender composition of the firm. At all-male firms, women are 74% (69 pp) less likely to qualify based purely on gender (column 1). This gap only narrows very slightly when restricting the sample to vacancies where the individual met the education and experience qualifications (column 2), and additionally the individual's selected occupations (column 3). The gap closes dramatically for firms that have any women, and reverses at firms that are at least half female at baseline.

Gender segregation is thus a symmetric phenomenon; but because male-dominated firms represent the vast majority of the market (Appendix Figure A.4), this results in many opportunities being closed to women. One possible explanation for firm level gender segregation is that firms need to incur fixed costs to integrate female employees (Miller *et al.*, 2022). These costs could be social or related to physical infrastructure, and could matter via firm-side decisions even though jobseekers cannot directly observe these attributes through the platform. In fact, only 43.8% of firms responding to survey questions at baseline reported having a women's toilet, and 57.3% reported having a place for women to pray, both key accommodations for female workers in the Pakistani context. Appendix Table A.9 shows that firms which have more of these features are more likely to open opportunities to women *and* select women applicants for interviews. However, firms' investment in such accommodations is endogenous to their preference for hiring women.

It is noteworthy, however, that a quarter of firms without any female employees did post an ad open to women—so baseline composition does not fully determine demand side gender restrictions at the posting stage (Appendix Figure A.4). This raises the question of whether industry or occupation might drive the patterns that firms that already have women are more likely to hire women. Appendix Figures A.5 to A.6 show the composition of gender requirements of job ads by industry and occupation. While these categories

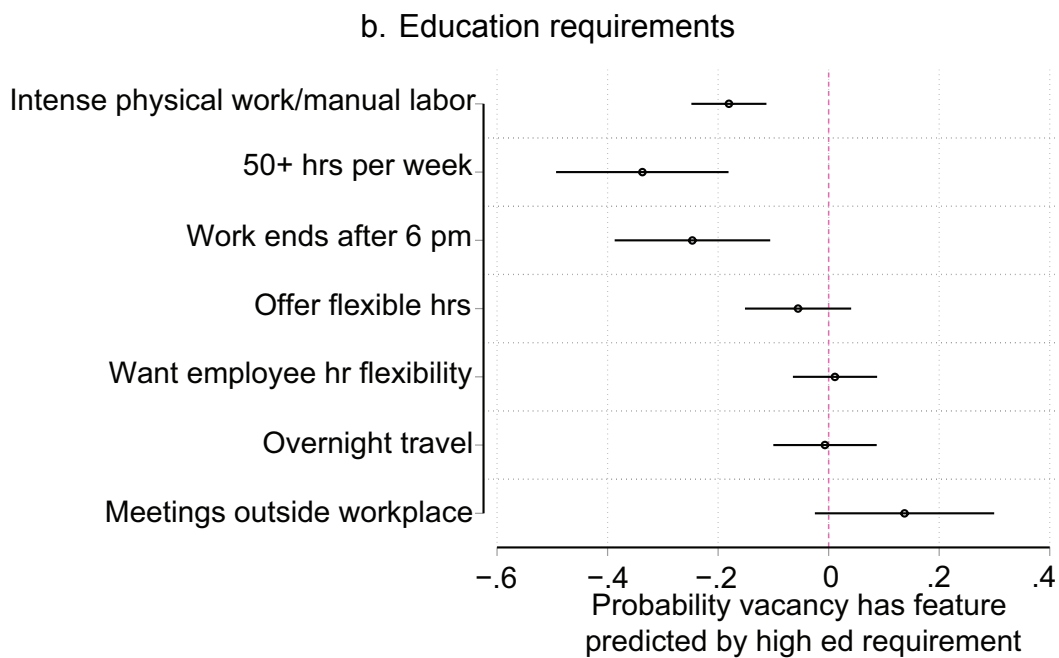

predict whether ads are open to women to some degree, there is a substantial degree of variation within industry and occupation in the gender restrictions employers place on ads. In our data, we find that industry and occupation explain 39% or less of the variation in whether a vacancy will accept applications from women (Appendix Table A.10). This is very similar to data from the Labor Force Survey 2018 for Pakistan; here, industry and occupation explain at most 37% of the variation in whether a worker is female.⁷ This is not enough to explain the full gender gap in whether an individual matches the explicit gender criteria for a given vacancy.

Since occupation and industry do not fully explain the variation in whether a vacancy is open to women, we further explore vacancy characteristics. Vacancies that require intense manual labor and more working hours (50+ hours/week) are significantly less likely to be open to women (Figure 2, Panel A). These patterns persist with occupation and industry fixed effects (Appendix Figure A.7), suggesting that firm-level variation is important. In Panel A of Figure 2, we see the analog of Panel A, but comparing characteristics of vacancies seeking candidates with greater or less than secondary education. The pattern mirrors Panel A. Panels C and D show a similar pattern for firm-level characteristics; vacancies at firms with a separate toilet or prayer space for women are also more likely to have a high education requirement. These results suggest that one mechanism through which education may help to close the gender gap is that “white collar” jobs are less likely to have requirements that lead employers to consider them unsuitable for women. We cannot rule out that employers who wish to hire women set shorter or earlier hours in order to attract female applicants; however, there is no similar pattern on offering flexible hours to employees.

⁷Similar calculations range from as low as 1.7% in the PRC to as high as 73% in Saudi Arabia (Kuhn and Shen, 2013; Miller *et al.*, 2022).


Figure 2: Vacancy and firm characteristics, gender restrictions and education

Panel I: Vacancy characteristics



Continued on the next page

Panel II: Firm characteristics



Notes: Unit of observation is the vacancy. Panel I: Data from 332 firms who post a total of 758 job advertisements. Panel II: Data on separate toilet and prayer space for women and men at the firm come from 452 ads from 178 firms who agreed to participate in this module of the survey. Data on separate work space comes from 388 ads from 150 firms (some observations are missing due to enumerator error). In Panels A and C, an indicator for 'is the vacancy open to women?' is regressed on each characteristic in separate regressions, and the coefficient from each regression is shown. In Panels B and D, an indicator for 'does the vacancy have a high education requirement?' is regressed on each characteristic in separate regressions, and the coefficient from each regression is shown. High education requirement refers to completed tertiary education (16 years or more). Standard errors clustered at the firm level; 95% confidence intervals shown. Figures A.7 and A.8 replicate the results in Panels A and C including occupation and industry fixed effects. Source: Authors' estimates.

6 Isolating Firm Decisions with Incentivized Résumé Rating

The administrative data results suggest the importance of *ex ante* demand-side gender constraints—limitations on whether the employer will even consider female applicants—in limiting women’s opportunities on the job market. However, we are underpowered with the administrative data to detect effects on interviews. Moreover, the interpretation of gender gaps at the interview stage from platform data is complicated by self-selection of jobseekers into application, which may differ by gender. To address these constraints, we use our IRR experiment. As described in Section 2.2, employers selected between pairs of anonymized CVs from subscribers to the platform, replacing their names with generic male and female names and swapping characteristics between members of the pair at random. Employers were incentivized with the information that their selection could inform the process used to send them job applicants for future vacancies.

The unit of observation is a CV k shown to a firm for vacancy j ; this includes both CVs in a binary choice as separate observations. We first estimate a linear probability model, regressing an indicator for whether the CV was chosen on an indicator for whether the CV was randomly assigned a female name and other attributes of the CV. We cluster standard errors by the binary choice pairs of CVs. This allows us to quantify the value employers place on a female name in the CV relative to other characteristics (Neumark, 2012). Results are reported in Table A.11. Employers are 11.6 pp less likely to select a CV with a female name. In striking contrast to this, experience level and education have no detectable effect on the probability a CV is selected.

We next explore heterogeneity by firm gender composition and restrictions (Table 4). Hiring managers in all-male firms (the majority of firms) are 23 pp less likely to select a CV with a randomly assigned female name than a male name (Column 1, β_1). At firms with any female employees, this pattern reverses; hiring managers are 22 pp *more* likely to select a CV with a female name (Column 1, $\beta_1 + \beta_2$). In firms with all female employees, this rate increases further; hiring managers are 60 pp more likely to select a CV with

a female name than a male name (column 2, $\beta_1 + \beta_2$). This pattern is consistent with findings from India on female-headed firms (Chiplunkar and Goldberg, 2021). The results also reinforce the pattern of firm-level gender segregation seen in the *ex ante* gender restrictions in Appendix Table A.8. The pattern in Column 1 also shows, however, that even firms without any women currently employed are choosing a female-named CV 33% of the time (control mean minus β_1). This is consistent with recent literature showing that removing gender criteria from job ads in the PRC and Austria has been shown to increase gender diversity of hires (Kuhn and Shen, 2023; Card *et al.*, 2021).

Finally, we examine whether the pattern of the gender gap closing with education observed in the administrative data is evident in the IRR experiment (Figure 3). The pattern of gender segregation seen in Table 4 is apparent for CVs at low levels of

Table 4: Incentivized résumé reporting results by firm gender composition


	CV Chosen	
	(1)	(2)
β_1 : Female name _k	-0.230** [0.072]	-0.151* [0.066]
β_2 : Female name _k X Group _j	0.452** [0.149]	0.751** [0.287]
β_3 : Group _j	-0.226** [0.074]	-0.376** [0.143]
$\beta_1 + \beta_2$:		
Total effect of female name _k in HTE group	0.222* [0.130]	0.600** [0.279]
Outcome Control Mean	0.558	0.558
N	430	430
HTE Group	Firm has any female employees	Firm has all female employees

CV = curriculum vitae, HTE = heterogeneous treatment effect.

Source: Authors' estimates. Note: The dependent variable is an indicator for if the CV_k was chosen. K refers to an individual CV. Female name_k is 1 if the gender assigned to CV_k was female. "Group_j" refers to the interaction variable j specified in the table footer. In column 1, the interaction variable is an indicator which takes the value 1 if the firm has any female employees and zero otherwise. In column 2, the interaction variable is an indicator which takes the value 1 if all employees in a firm are female and zero otherwise. Robust standard errors in brackets, clustered by firm. * $p < .1$, ** $p < .05$, *** $p < .01$.

Source: Authors' estimates.

Figure 3: Incentivized résumé rating: Heterogeneity by CV education level



CV = curriculum vitae.

Note: This figure displays coefficients from an OLS regression run separately for full sample, for firms with no women (75% of the sample), and firms that already have at least one female employee (25% of the sample). The dependent variable 'CV Chosen' is a binary indicator equal to 1 if CV k was chosen by the respondent in the Incentivized Resume Rating exercise. The coefficients shown are for the interaction of 'Female CV' and education levels (Primary, Secondary, and Tertiary). Each observation in these regressions is one CV shown to the firm. Standard errors clustered by CV pair IDs; 95% confidence intervals shown.

Source: Authors' estimates.

education. Firms with no female employees are less likely to select CVs with female names, while firms with at least some female employees are more likely to do so. As the education level of the hypothetical candidate rises, the gender gap closes completely: women's CVs are equally likely to be selected regardless of whether women already work at the firm.

7 Conclusion

We assemble a unique dataset from a job matching platform in Lahore, Pakistan, with several advantages. First, rare in literature that studies labor markets through job search and matching platforms, we begin with a representative listing of households and firms across a large metropolitan city. Second, the nature of the platform allows us to observe fine-grained search decisions by both employers and individuals. Third, due to the matching process on this platform, we observe which vacancies are seen by individuals and which individuals are considered by firms, and consequently observe the fine-grained decisions that result in some individual-vacancy dyads converting to matches, and some not.

We use this administrative dataset to decompose supply- versus demand-side constraints that give rise to gender gaps in employment. To help isolate firm-side decisions after the ad posting stage, we combine this analysis with an incentivized resume rating experiment to shed light on how firms decide between female- and male-named CVs, holding other observables constant. Through this combination of data and methods, we first conclude that gender gaps in employment are larger than gender gaps in job search. Second, firm-side explicit gender criteria dominate other factors in determining the quantity of opportunities open to women with less than a secondary education; this gender gap in quantity of opportunities closes dramatically with education. Firms' explicit gender criteria reflect existing firm-level gender segregation as well as requirements of the job perceived unfriendly to women, such as long and late working hours. Third, at higher levels of education, women become more selective; they are less likely than men to choose the occupation of a given dyad or to apply to a given match. However, demand-side factors constrain the quality of job opportunities available to women at the highest education levels; women are more likely than men to meet gender qualifications for the lowest salary quintile of vacancies hiring those with tertiary education.

These results help contextualize a growing literature documenting specific barriers to women's employment on both the supply and demand side. Much of the recent literature that studies low female employment focuses on alleviating supply-side constraints via interventions such as overcoming information asymmetries, training in socio-emotional skills, addressing norms by engaging partners and family members, safe transport, and social protection programs that target women. Relatively less emphasis has been placed on demand-side interventions, including incentives such as tax breaks or grants, for firms to offer workplace facilities that would be inclusive to women which might in turn increase firms' willingness to hire women. Furthermore, much of the existing literature focuses on across-sector occupational gender segregation, rather than within-sector across-firm variation.

The majority of women and men in the population we study, and indeed in many settings with low levels of women's employment, have less than a secondary education. We demonstrate that in one such population, firm gender criteria are overwhelmingly the binding constraint to women's job opportunities, compared to any decisions that individuals make in their own job search. Across-sector variation does not fully explain whether firms are willing to hire women; rather within-sector differences in firm infrastructure and vacancy characteristics are correlated both with the education level at which firms are hiring and whether they are open to hiring women. Thus, our results suggest that while supply-side decisions are important, alleviating demand-side constraints to female employment might have larger impact.

References

- Addati, L., Bonnet, F., Ernst, E., Merola, R. and Wan, P. M. J. (2016). *Women at Work: Trends 2016*. Tech. rep., International Labour Office, Geneva: ILO.
- Afridi, F., Dinkelman, T. and Mahajan, K. (2017). Why are fewer married women joining the work force in rural India? A decomposition analysis over two decades. *Journal of Population Economics*, **31** (3), 783–818.
- Ahmed, H., Mahmud, M., Said, F. and Tirmazee, Z. S. (forthcoming). Encouraging female graduates to enter the labor force: Evidence from a role model intervention in Pakistan. *Economic Development and Cultural Change*, **0** (ja), null.
- Akerlof, G. A. and Kranton, R. (2000). Economics and identity. *The Quarterly Journal of Economics*, **CVX** (3), 715–753.
- Archibong, B., Annan, F., Benshaul-Tolonen, A., Okunogbe, O. and Oliobi, I. (2022). Firm culture: Examining the role of gender in job matching in an online African labor market. *CEPR Working Paper*.
- Aromolaran, A. B. (2004). Female schooling, non-market productivity, and labor market participation in Nigeria. *SSRN Electronic Journal*, (493763).
- Banfi, S., Choi, S. and Villena-Roldán, B. (2019). Deconstructing job search behavior.
- Belot, M., Kircher, P. and Muller, P. (2018). Providing advice to jobseekers at low cost: An experimental study on online advice. *Review of Economic Studies*, **86** (4), 1411–1447.
- Cameron, L. A., Dowling, J. M. and Worswick, C. (2001). Education and labor market participation of women in Asia: Evidence from five countries. *Economic Development and Cultural Change*, **49** (3).
- Card, D., Colella, F. and Lalive, R. (2021). Gender preferences in job vacancies and workplace gender diversity. *NBER Working Paper Series*, **29350**.
- Chaturvedi, S., Mahajan, K. and Siddique, Z. (2022). Words matter: Gender, jobs and applicant behavior. *Unpublished Manuscript*.
- Chiplunkar, G. and Goldberg, P. (2021). Aggregate implications of barriers to female entrepreneurship. *Unpublished Manuscript*.
- Cortes, P. and Pan, J. (2017). Occupation and gender.
- , —, Reuben, E., Pilossoph, L. and Zafar, B. (2022). Gender differences in job search and the earnings gap: Evidence from the field and lab. *Working paper*.
- Dean, J. T. and Jayachandran, S. (2019). Changing family attitudes to promote female employment. *AEA Papers and Proceedings*, **109**, 138–142.
- Delfino, A. (2022). Breaking gender barriers: Bringing men into the pink-collar jobs of the future.

- Field, E., Garlick, R., Subramanian, N. and Vyborny, K. (2023). *Why don't jobseekers search more? Barriers and returns to search on a job matching platform*. Tech. rep.
- , Pande, R., Rigol, N., Schaner, S. and Troyer Moore, C. (). On her own account: How strengthening women's financial control impacts labor supply and gender norms. *American Economic Review*, **111** (7), 2342–2375.
- and Vyborny, K. (2016). Female labor force participation: Pakistan country study for the Asian Development Bank.
- and — (2022). Women's mobility and labor supply: Experimental evidence from Pakistan. *ADB Economics Working Paper Series*, (655).
- Fletcher, E., Pande, R. and Moore, C. M. T. (2018). Women and work in India: Descriptive evidence and a review of potential policies. *HKS Working Paper*, (RWP18-004).
- Flory, J. A., Leibbrandt, A. and List, J. A. (2015). Do competitive workplaces deter female workers? A large-scale natural field experiment on job entry decisions. *The Review of Economic Studies*, **82** (1), 122–155.
- Fluchtmann, J., Glenny, A. M., Harmon, N. and Maibom, J. (2021). The gender application gap: Do men and women apply for the same jobs? *IZA Discussion Paper Series*, **14906**.
- Gaddis, I. and Klasen, S. (2014). Economic development, structural change, and women's labor force participation: A reexamination of the feminization U hypothesis. *Journal of Population Economics*, **27**, 639–681.
- Garlick, R., Hensel, L., Kiss, A. and Orkin, K. (2023). Jobseekers' beliefs about comparative advantage and (mis)directed search. *Working paper, Duke University*.
- Goldin, C. (1995). *The U-shaped female labor force function in economic development and economic history*, University of Chicago Press, pp. 61–90.
- and Katz, L. (2016). A most egalitarian profession: Pharmacy and the evolution of a family-friendly occupation. *Journal of Labor Economics*, **34** (3), 705–746.
- and Rouse, C. (2000). Orchestrating impartiality: The impact of "blind" auditions on female musicians. *American Economic Review*, **90** (4), 715–741.
- Hangartner, D., Kopp, D. and Siegenthaler, M. (2021). Monitoring hiring discrimination through online recruitment platforms. *Nature*, **589**, 572–576.
- Heath, R. and Mobarak, A. M. (2015). Manufacturing growth and the lives of Bangladeshi women. *Journal of Development Economics*, **115**, 1–15.
- International Labour Organization (2019a). Labor force participation rate, female (% female population ages 15+).
- International Labour Organization (2019b). Labor force participation rate, male (% male population ages 15+).

- Jensen, R. (2012). Do labor market opportunities affect young women's work and family decisions? Experimental evidence from India. *The Quarterly Journal of Economics*, **127** (2), 753–792.
- Jones, S. and Sen, K. (2022). Labour market effects of digital matching platforms: Experimental evidence from sub-Saharan Africa. *IZA Discussion Paper*, (15409).
- Kessler, J. B., Low, C. and Sullivan, C. D. (2019). Incentivized resume rating: Eliciting employer preferences without deception. *American Economic Review*, **109** (11), 3713–3744.
- Klasen, S. (2019). What explains uneven female labor force participation levels and trends in developing countries? *World Bank Research Observer*, **34** (2), 161–197.
- and Pieters, J. (2015). What explains the stagnation of female labor force participation in urban India? *The World Bank Economic Review*, **29** (3), 449–478.
- Kuhn, P. and Shen, K. (2013). Gender discrimination in job ads: Evidence from China. *The Quarterly Journal of Economics*, **128** (1), 287–336.
- and — (2023). What happens when employers can no longer discriminate in job ads? *American Economic Review*, **113** (4), 1013–48.
- , — and Zhang, S. (2020). Gender-targeted job ads in the recruitment process: Facts from a Chinese job board. *Journal of Development Economics*, **147**.
- Kureková, L. M., Beblavý, M. and Thum-Thysen, A. (2015). Using online vacancies and web surveys to analyse the labour market: A methodological inquiry. *IZA Journal of Labor Economics*, **4** (1).
- Mas, A. and Pallais, A. (2017). Valuing alternative work arrangements. *American Economic Review*, **107** (12), 3722–3759.
- Matsuda, N., Ahmed, T. and Nomura, S. (2019). Labor market analysis using big data: The case of a Pakistani online job portal.
- Miller, C., Peck, J. and Seflek, M. (2022). Missing women, integration costs, and big push policies in the Saudi labor market. *American Economic Journal: Applied Economics*, **14** (2), 51–77.
- Neumark, D. (2012). Detecting discrimination in audit and correspondence studies. *Journal of Human resources*, **47** (4), 1128–1157.
- Nomura, S., Imaizumi, S., Areias, A. C. and Yamauchi, F. (2017). Toward labor market policy 2.0: The potential for using online job-portal big data to inform labor market policies in India. *World Bank Policy Research Working Paper*, (7966).
- Ozen, E. N., Hut, S., Levin, V. and Munoz Boudet, A. M. (2019). The importance of being modest: The roles of gender and socioemotional skill signals in the labor market.

Subramanian, N. (2021). Workplace attributes and women's labor supply decisions: Evidence from a randomized experiment.

United Nations Population Division (2023). Population, total—Pakistan | data.

Wheeler, L., Garlick, R., Johnson, E., Shaw, P. and Gargano, M. (2021). LinkedIn(to) job opportunities: Experimental evidence from job readiness training. *American Economic Journal: Applied Economics*, (forthcoming).

Barriers to Entry

Decomposing the Gender Gap in Job Search in Urban Pakistan

Using matched data from three sources in Lahore, Pakistan, the paper finds that employers' gender restrictions are a larger constraint on women's job opportunities than supply-side decisions. At higher levels of education, demand-side barriers relax, allowing women to qualify for more jobs but at lower salaries. On the supply side, educated women become more selective in their search.

About the Asian Development Bank

ADB is committed to achieving a prosperous, inclusive, resilient, and sustainable Asia and the Pacific, while sustaining its efforts to eradicate extreme poverty. Established in 1966, it is owned by 68 members—49 from the region. Its main instruments for helping its developing member countries are policy dialogue, loans, equity investments, guarantees, grants, and technical assistance.



ASIAN DEVELOPMENT BANK

6 ADB Avenue, Mandaluyong City

1550 Metro Manila, Philippines

www.adb.org