

Why Don't Jobseekers Search More?

Barriers and Returns to Search on a Job Matching Platform*

Kate Vyborny (r) Robert Garlick (r) Nivedhitha Subramanian (r) Erica Field (r) †

June 22, 2026

Abstract

Understanding specific barriers to job search and returns to relaxing these barriers is important for economists and policymakers. An experiment that changes the default process for initiating job applications increases applications by 600% on a search platform in Pakistan. Perhaps surprisingly, the marginal treatment-induced applications have approximately constant rather than decreasing returns. These results are consistent with a directed search model in which some jobseekers miss some high-return vacancies due to psychological costs of initiating applications. These findings show that small reductions in search costs can substantially improve search outcomes in environments with some relatively inactive jobseekers.

*We are grateful for valuable feedback from the editor; from seminar participants at Bowdoin, Duke, Maine, Oxford, and Virginia; from conference participants at AEA, AFE, CREB ADE, GLM-LIC, IZA Firm-Worker Matching, MIEDC, PacDev, SOLE, and WB Jobs and Development; and from Martin Abel, Michele Belot, Jenn Kades, Jeremy Magruder, Simon Quinn, Todd Sarver, Jeffrey Smith, Jessi Streib, Duncan Thomas, and Chris Woodruff. We thank Gustavo Acosta, Leonor Castro, Aiman Farrukh, Harmalah Khan, and Sahil Nisar for outstanding research assistance in preparing this draft. We thank the entire Job Talash team at CERP, in particular Maniha Aftab, Tehneiat Amjad Butt, Noor Anwar Ali, Rukhshan Arif Mian, Sarah Hussain, Hareem Fatima, Uzair Junaid, Lala Rukh Khan, Alieha Shahid, Mahin Tariq and Zoha Waqar for excellent assistance for this research program. This study is made possible by the generous support of the American people through the United States Agency for International Development (USAID) through Development Innovation Ventures (DIV). We gratefully acknowledge funding for this project from J-PAL GEA. We are also thankful for funding support for the broader Job Talash research program from the Asian Development Bank, GLM-LIC, 3IE, IGC, IPA, J-PAL, PEDL, the National Science Foundation (SES #1629317), the Gates Foundation, and the World Bank Umbrella Fund for Gender Equality (UFGE). This research received ethics approval from Duke University (#2019-0067). The experimental design, intervention, and primary outcomes are registered on the AEA's RCT registry at <https://doi.org/10.1257/rct.5997-1.0>. The authors' views do not necessarily represent the views of USAID or the US government.

† Author order is randomized using the AEA's randomization tool with confirmation code [J1rjReJNEuDt](#). Vyborny: World Bank, kvyborny@worldbank.org; Duke University, robert.garlick@duke.edu; Subramanian: Bates College, nsubrama@bates.edu; Field: Duke University, BREAD, NBER, field.ERICA@duke.edu

1 Introduction

Job search is a central feature of labor markets, and search frictions can have important economic consequences. For instance, in macroeconomic models, frictional search can help to explain both employment levels and the productivity of firm-worker matches (Shimer, 2010). Microeconomic research has documented many specific job search frictions ranging from pecuniary search costs to incomplete information (e.g. Abebe et al. 2021a,b; Abel et al. 2019; Bandiera et al. 2021; Belot et al. 2018; Franklin 2017). Recent work has shown that behavioral factors such as present bias, reference-dependence, and motivated reasoning can also impact search decisions (e.g. Cooper & Kuhn 2020; DellaVigna et al. 2022; Mueller & Spinnewijn 2022).

We study behavioral barriers to job search effort on a search and matching platform. The platform sends text messages about relevant new vacancies to jobseekers, who must call the platform to apply. Adding follow-up calls inviting jobseekers to immediately start applications, which reduces the initiative required to apply, substantially increases their propensity to apply. Moreover, returns to the additional applications, measured in terms of interview invitations, are approximately constant rather than decreasing. This raises the question of why jobseekers don't apply more in the absence of calls. To explain this, we propose a model with heterogeneous psychological costs of initiating applications that can be high enough to deter some applications to even high-return vacancies, resulting in suboptimally low search effort.

To generate experimental evidence on this search barrier, we work with a novel job search platform in Lahore, Pakistan.¹ Platform data allow us to observe all vacancy characteristics, job application decisions, application materials, and interview outcomes for roughly 1.1 million matches between vacancies and jobseekers. The 9,800 jobseekers are recruited from a city-wide representative household listing. Thus, they have a wide range of education, from incomplete primary to graduate levels, and a wide range of baseline labor force attachment, from employed and searching to non-employed and non-searching. This sample breadth is unusual in experimental job search studies (Poverty Action Lab, 2022), partly because of the household listing and partly because using the platform requires only basic literacy, a simple phone, and almost no airtime, generating

¹Job platforms have become a central feature of many labor markets. In Pakistan in 2021, Rozee, LinkedIn, and Bayt had respectively 9.5, 7.5, and 3 million users. Bayt had 39 million users in 2021 across the Middle East, North Africa, and South Asia. LinkedIn had > 10 million users in 2022 in at least 8 developed and 10 developing countries.

very few technological and pecuniary barriers to search.

Our main experimental treatment changes how jobseekers initiate applications on the platform, moving them from an active role to a passive role. Specifically, all users receive monthly text messages listing new vacancies that match the qualifications and preferences they report at sign-up. Users in the control group must call the platform to initiate applications, while users in the treatment group also receive a phone call inviting them to apply, so they do not need to initiate calls to apply. The experimental design holds constant other parts of the search process: the phone call has negligible effects on pecuniary and time costs of applying, provides no direct encouragement or pressure to apply, and provides no extra information about vacancies. Hence, we interpret treatment as primarily reducing the psychological cost of initiation.

Our two key findings are that phone calls dramatically increase the job application rate, and that the average return to additional applications is roughly constant rather than decreasing. Treatment increases the share of jobseeker-vacancy matches receiving applications by seven-fold, from 0.2 to 1.5%.² Using treatment as an instrument for applying shows that marginal applications submitted due to treatment have a 5.9% probability of yielding interviews, which is neither substantively nor statistically different from the 6.4% probability for applications from the control group. This implies that returns to job search are roughly constant over this large increase in applications. The same pattern holds when we weight interviews by their desirability in terms of salary, hours, commuting, and non-salary benefits. An additional experiment shows that this finding is not explained by differences in the quality of jobseekers who submit marginal versus inframarginal applications. We also develop tests to show that the constant returns finding is robust to potential complications around the exclusion and monotonicity conditions in our instrumental variables analysis.

The finding of roughly constant returns is surprising. We might expect jobseekers to prioritize applying to vacancies with the highest combination of expected interview probabilities and desirable attributes, and hence that extra applications would have decreasing returns. This behavior would be consistent with many models of ‘directed’ job search, reviewed by [Wright et al. \(2021\)](#). The constant returns finding by itself is consistent with canonical models of ‘random’ job search (e.g. [Pissarides 2000](#)) but we show later that our other results are inconsistent with random search.

²It is unsurprising that most matches do not generate applications. A match simply means the jobseeker qualified for the job and is interested in that occupation. In any search environment, jobseekers will apply to only a small subset of such jobs. The same pattern holds on some other platforms that economists have studied (Appendix A).

To explain our two key findings, we propose a modified directed search model. As in many models, in each period jobseekers apply to those vacancies with expected return above the cost of applying. Our key assumption is that application costs vary – within jobseeker through time and/or between jobseekers – and can be large enough that some jobseekers submit no applications in periods in which they face high costs, even to high return vacancies. For example, a jobseeker may face a high psychological cost of initiating applications when they are stressed by illness, domestic responsibilities, or work. The phone call treatment reduces application costs, leading naturally to more applications. However, these marginal applications come from two sources: jobseekers facing already low costs apply to additional vacancies, which will have lower average returns than their inframarginal applications, and jobseekers facing high costs – who would not have applied to any vacancies in that period without treatment – now apply to *some* vacancies. Because the second type of marginal applications can have higher returns than inframarginal applications, the average return to marginal applications – averaged across treated jobseekers facing high and low cost draws – can equal the return to inframarginal applications.

This framework shows how the common assumption of decreasing returns to additional search *for each individual jobseeker in each period* can lead to constant returns to additional search *averaged over jobseekers and periods*, provided some jobseekers are not actively searching in some periods. This model’s predictions match both our two key findings and our secondary results about which jobseekers submit marginal applications and where they direct them.

Given the importance of application costs in the experiment and the model, we explore in detail what types of costs jobseekers face. We show that pecuniary and time costs of applying on our platform are low, and that additional experiments that directly reduce the pecuniary or time costs of applying have little effect on applications. This leaves psychological costs of initiating applications as the most likely category of cost addressed by the phone call treatment. Within this category, the existing literature suggests multiple candidates, including the cognitive cost of paying attention to text messages and mentally processing their content ([Gabaix, 2019](#)), fear of applications being rejected ([Köszegi et al., 2022](#)), and present bias ([Ericson & Laibson, 2019](#)), all of which can vary through time. Our findings and interpretation are consistent with research showing that eliminating the need to initiate decisions can raise financial and health investment, reviewed by [DellaVigna \(2009\)](#). Our key modeling assumption of heterogeneous psychological

costs borrows from behavioral models that seek to explain low adoption of seemingly high-return investments (Carroll et al., 2009; Duflo et al., 2011).

We can test and reject several plausible alternative explanations. Perhaps most importantly, the constant returns finding by itself is consistent with prominent models of ‘random’ job search, the main alternative to directed search models. In these models, vacancies are homogeneous and jobseekers randomly choose where to apply (Pissarides, 2000). But random search models do not match other results we find: not only do we observe that jobseekers on our platform direct applications to vacancies with desirable attributes, but we also run an additional experiment to encourage random search that generates sharply *decreasing* returns to marginal applications.

We can also reject some specific behavioral explanations – reminders or explicit encouragement or pressure – because these are inconsistent with the platform design and results from additional experiments. Information- or belief-based explanations — e.g. more information about matches or higher perceived returns to applications — are also inconsistent with the platform design, results from additional experiments, and survey measures of beliefs.

We do not find evidence that this additional search has negative spillovers on other jobseekers. Individual jobseekers’ interview probabilities are unaffected by competing against more treatment-induced applications. We treat 50% of jobseekers on the platform, increasing total search enough that large spillovers are possible. But our scope for detecting spillovers is limited by the fact that the majority of firms’ job applications come from sources other than the platform.

On the platform, search outcomes are measured using interviews and interview attributes. Interviews are an important search outcome because they are a necessary condition for job offers and impose non-trivial costs on both job applicants and firms. Hence their widespread use in some areas of labor economics such as audit studies. Using interviews or even applications as final outcomes is relatively common in the literature studying search on platforms (e.g. Belot et al. 2018), either for power reasons or because most platforms do not track data on job offers or employment.³

Our paper makes three contributions. First, by studying psychological costs of job search, we

³Banfi et al. (2019), Belot et al. (2022b), Faberman & Kudlyak (2019), He et al. (2021), Marinescu (2017a), and Marinescu & Wolthoff (2020) use applications as their final outcomes. Fewer papers use employment as a final outcome and these largely use administrative employment data in high-income countries (Behaghel et al., 2023; Belot et al., 2022a; Ben Dhia et al., 2022; Fernando et al., 2021; Marinescu & Skandalis, 2021). Online gig work platforms provide employment data, but for a very different type of work (e.g. Agrawal et al. 2015). A related set of papers study platform users using survey data but have limited data on platform use (e.g. Kelley et al. 2021; Wheeler et al. 2022).

contribute to a nascent literature on behavioral job search, reviewed by [Cooper & Kuhn \(2020\)](#). Existing work shows patterns of job search that are consistent with present bias, motivated reasoning, and reference dependence (e.g. [DellaVigna et al. 2017, 2022](#); [Mueller & Spinnewijn 2022](#); [Paserman 2008](#)), but does not isolate the psychological costs of initiating job applications.⁴

Second, our results have clear and novel policy implications for addressing behavioral barriers to search. [Babcock et al. \(2012\)](#) suggest multiple ways to harness behavioral economics to encourage and improve job search. However, there are few evaluations of policies designed to directly target behavioral factors, all of these focus on helping jobseekers make plans to search more, and none compares returns to marginal and inframarginal search ([Abel et al., 2019](#); [Caria et al., 2023](#); [Sanders et al., 2019](#)). We extend this work by running multiple field experiments to show how different changes to the job search environment can produce substantially different impacts on search effort and different returns to search. Behavioral channels may be relevant for many other job search policies: motivated reasoning might affect how jobseekers process and use new information, present bias and reference dependence might influence how jobseekers spend subsidies, and relationships between caseworkers and jobseekers might have behavioral components.⁵ However, research into these job search policies has not sought to pin down behavioral components.

Third, we provide a direct estimate of returns to additional search effort caused by reducing behavioral barriers. Returns to search effort, typically interpreted as job applications, are central to canonical job search models ([Pissarides, 2000](#)) and are important for evaluating policies such as search subsidies or search requirements for recipients of unemployment insurance. Direct estimates of returns to search are very rare, making it difficult to understand variation in the effects of search-related policies – e.g. is this variation due to different effects on search or returns to search? – or to design search promotion policies – e.g. how many applications should be subsidised?

To make this third contribution, we combine experimental variation in search costs with data on both individual applications and the outcomes of those applications. This is a very rare combination in the literature. Many papers study the effect on employment of search subsidies or

⁴Related work studies links between job search and locus of control ([Caliendo et al., 2015](#); [McGee, 2015](#)) and behavioral job search in labs ([Brown et al., 2011](#); [Falk et al., 2006a,b](#); [Fu et al., 2019](#); [McGee & McGee, 2016](#)).

⁵[Abebe et al. \(2021a,b\)](#), [Abel et al. \(2020\)](#), [Altmann et al. \(2018, 2022\)](#), [Bandiera et al. \(2021\)](#), [Bassi & Nansamba \(2020\)](#), [Beam \(2016\)](#), [Behaghel et al. \(2023\)](#), [Belot et al. \(2018, 2022a\)](#), [Boudreau et al. \(2022\)](#), [Carranza et al. \(2021\)](#), [Dammert et al. \(2015\)](#), [Kiss et al. \(2023\)](#), [Spinnewijn \(2015\)](#), and [Subramanian \(2021\)](#) study information. [Abebe et al. \(2019, 2021a\)](#), [Banerjee & Sequeira \(2020\)](#), [Field et al. \(2024\)](#), and [Franklin \(2017\)](#) study subsidies. [Arni & Schiprowski \(2019\)](#), [Bolhaar et al. \(2020\)](#), [Lechner & Smith \(2007\)](#), and [Schiprowski \(2020\)](#) study caseworkers.

requirements for receipt of government benefits, but do not observe actual search effort (reviewed by [Card et al. 2010, 2018](#); [Filges et al. 2015](#); [Marinescu 2017b](#)). A smaller, more recent literature studies the effect of search subsidies or requirements on online search effort, but without observing any outcomes of search ([Baker & Fradkin, 2017](#); [Marinescu, 2017a](#); [Marinescu & Skandalis, 2021](#)). Other recent papers experimentally shift search strategies or search technologies, but do not isolate the role of search effort and mostly rely on low-frequency survey data that cannot link specific search actions to outcomes.⁶ The closest work to our own shows how additional policy-induced job applications affect unemployment duration ([Arni & Schiprowski, 2019](#); [Lichter & Schiprowski, 2021](#)). While we do not observe administrative data on employment, we extend this work by using application-level data that allow us to describe how marginal and inframarginal search effort is directed, and to compare the outcomes of marginal and inframarginal applications. Our findings about how jobseekers direct applications to specific vacancies and miss applying to some high-return vacancies link to a growing literature on directed job search.⁷

In Section 2, we describe the context, sample, platform, and experimental design. In Section 3, we present the treatment effects on job applications and interviews and the implied effect of marginal job applications on interviews. We describe our preferred and alternative interpretations in Section 4. Section 5 discusses spillover effects.

2 Economic Environment

2.1 Context

Our experiment takes place on Job Talash (“job search” in Urdu), a job search and matching platform in Lahore, created by our research partners at the Center for Economic Research in Pakistan. Lahore is a city of about 10 million with an adult labor force participation rate of 49% and employment rate of 47%, both substantially higher for men than women (Table A.1). Gender is an important feature of Lahore’s labor market ([Gentile et al., 2023](#)) but we do not focus on gender in

⁶See the preceding footnote for examples. In particular, our work differs from recent papers studying the effects of encouraging enrollment on job search platforms (e.g. [Afridi et al. 2022](#); [Chakravorty et al. 2023](#); [Jones & Sen 2022](#)). Joining a platform is a bundled experience that might shift factors ranging from wage expectations ([Kelley et al., 2021](#)) to information about specific vacancies ([Wheeler et al., 2022](#)). These papers have substantially different interpretations to our treatment, as does the effect of access to (faster) online job search ([Bhuller et al., 2019](#); [Chiplunkar & Goldberg, 2022](#); [Gurtzgen et al., 2020](#); [Hjort & Poulsen, 2019](#); [Kuhn & Skuterud, 2004](#); [Kuhn & Mansour, 2014](#)).

⁷[Alfonso Naya et al. \(2020\)](#), [Behaghel et al. \(2023\)](#), [Belot et al. \(2018, 2022b\)](#), [Kiss et al. \(2023\)](#), [Gee \(2019\)](#), [He et al. \(2021\)](#), and [Marinescu & Wolthoff \(2020\)](#) also study the role of information about vacancies in job applications.

this paper because all of our main experimental results hold for both women and men. Job search and matching platforms are a growing feature of Pakistan’s labor market, particularly in major cities such as Lahore, as we describe in footnote 1.

2.2 Samples of Jobseekers and Firms

We recruited participants by conducting a household listing from a random sample of 443 enumeration areas across Lahore between October 2016 and September 2017. This provides a representative listing of 49,504 households and 182,591 adults. We invited each adult household member to sign up for the Job Talash platform and 46,572 expressed interest. The Job Talash call center called each of these people to collect information on their education, work experience, job search, and occupational preferences. The 9,838 people who completed sign-up comprise our main sample.

This sampling process is designed to include participants with different levels of education and labor market attachment, including those who are neither employed nor searching. This is relatively unusual in experimental work in labor economics.⁸ This allows us to show that the search barrier we identify affects many different types of active and potential jobseekers. This breadth is also important because the distinction between non-employed searchers and non-searchers is loose and transient in many developing economies (Donovan et al., 2018).

Column 1 of Table 1 presents descriptive statistics for the control group in our study sample. At baseline, 20% of the sample were employed and searching through some channel other than Job Talash, 35% were searching but not employed, 14% were employed but not searching, and 31% were neither employed nor searching. Network search was the most common method (40% of the sample) followed by visiting establishments to ask about vacancies (23%) and applying formally (15%).⁹ Only 4% had used a job search assistance program or online platform other than Job Talash. The average respondent had 7.9 years of work experience with an interdecile range of 0-16. Respondents’ education levels also vary widely: 15% had no education, 15% had completed secondary school, and 25% had a university degree. 30% were female and the average age was 31, with an interdecile range of 20-45.

⁸Of the 29 experimental job search studies reviewed by Poverty Action Lab (2022), only 8 construct samples from household listings, while another 12 sample from unemployment registries and 4 from job search assistance services, whose participants are required or strongly encouraged to search.

⁹The prevalence of network-based search matches patterns in other developing economies (Caria et al., 2024; Carranza & McKenzie, 2024). In Lahore’s Labor Force Survey, direct applications are slightly more common than network-based search (Table A.1). But this survey does not measure on-the-job search, unlike our own survey.

The study sample starts from a representative household listing, but only 5.3% of adults from this listing completed registration on Job Talash and became part of the study sample. In Table A.1, we compare our study sample to the population of Lahore, captured by both the official Labor Force Survey and our household listing. Our study sample is slightly younger, more male, more educated, less likely to be employed and much more likely to be searching for work. This selection means that our findings should not be extrapolated to the population of Lahore, but rather speak to a population who registers on a new job search platform. See [Gentile et al. \(2023\)](#) for additional discussion on selection from the household listing to registration on this platform.

We enrolled firms through a door-to-door listing in commercial areas of Lahore, described in more detail in Appendix A. We invited firms to list any current vacancies during enrollment and recontacted them several times each year to invite them to list more vacancies. For each vacancy, we collected the job title, location, occupation, salary, benefits, and hours. Vacancies cover a wide range of education and experience levels and occupations, including computer operator, makeup artist, salesperson, sweeper, security guard and HR manager. Column 1 of Table 2 shows that the average vacancy offers a monthly salary of 14,420 Pakistani Rupees (432 USD PPP) and is posted by a firm with 27 employees that hired 5.5 people in the last year.¹⁰ At baseline, only 22% of firms had ever advertised a vacancy on a job search platform, while 69% had recruited through referrals, 35% from CVs dropped off by jobseekers, and 11% through newspapers or other traditional media.

2.3 Job Talash Platform

The Job Talash service is free to both jobseekers and firms. It requires only literacy and access to a phone with call and text message functionality. This allows broad access to the platform and easy scaling because 97% of urban households in Lahore’s province have mobile phones ([MICS, 2018](#)).

After signing up, jobseekers are matched to each listed vacancy using a simple algorithm: the jobseeker must have at least the required years of education and experience, match any gender requirement, and have indicated interest in this occupation.¹¹ We refer to each jobseeker-vacancy

¹⁰These summary statistics weight each vacancy by the number of jobseekers who match with the vacancy. We define a jobseeker $\times$ vacancy match in the next subsection. The mean salary offer is roughly 60% of the mean salary in the Labor Force Survey data for Lahore (Figure A.1) and roughly 60% of the mean salary for vacancies posted during the same period on Rozee, Pakistan’s largest job search portal ([Matsuda et al., 2019](#)). However, this does not necessarily indicate negative selection into our sample of vacancies, as the Labor Force Survey data are not restricted to starting salaries and Rozee caters mainly to highly educated jobseekers.

¹¹Of the vacancies listed on this platform, 20.2% are open only to women and 45.3% are open only to men. Explicitly gender-targeted job listings are common in Lahore’s labor market and in other settings ([Kuhn & Shen, 2013](#)).

pair, for which the respondent qualifies and has indicated interest in the occupation, as a *match*. We study 1,116,952 matches generated by the platform over four years. The average jobseeker received 113 matches (1.8 per month) with interdecile range 7-271.

Importantly, there is substantial heterogeneity in proxies for the quality of these jobseeker-vacancy matches. Column 1 of Table 2 shows summary statistics for match attributes in the control group. For example, the jobseeker has education and work experience that are an exact match for the employer's preferences in only 18 and 13% of matches respectively.¹² Furthermore, 85% of jobseekers indicate interest in multiple occupations, with the median jobseeker interested in six occupations. These patterns show heterogeneity in how much firms might value jobseekers matched to their vacancies and how much jobseekers might value the vacancies to which they are matched. This creates the potential for heterogeneous returns to applications, which is important for interpreting our experimental results.

The platform collects new vacancy listings from firms every 1-2 months and sends jobseekers text message updates if they have matched to any vacancies in that month. See Figure A.2 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 website that lists vacancies. Jobseekers on average receive a text every 2.8 months. Conditional on receiving any matches in that month, the average jobseeker receives 3.1 matches with interdecile range 1-6.

If a jobseeker wants to apply to any of these vacancies, the platform forwards her CV to the firm. (We describe the application process in Section 2.5.) The CVs are constructed by the platform by populating a template with respondent-specific demographics, education, and work experience, so there is no variation in CV design. The platform sends all applications to the firm in a packet at the application deadline, so application timing does not affect interview probability. If the firm wants to interview the jobseeker, they contact the jobseeker directly to arrange the interview. The Job Talash team surveys each firm a few weeks after the application deadline to ask which applicants they interviewed.

¹²For each vacancy, the platform collects both the *required levels* and *preferred types* of education and experience. Jobseekers are only matched to vacancies if they have the required levels of experience and education, e.g. complete high school and five years of work experience. They can be matched even if they do not have the preferred types of education and experience, e.g., their work experience might be in a non-preferred field. We use the alignment between jobseekers' education and experience and vacancies' preferred types as a measure of match quality.

The platform design has two key advantages for our research, relative to other job search environments. First, we observe all information available to both sides of the market. We observe the same information about vacancies that jobseekers receive through the text messages, and the same information about jobseekers that firms receive through the CVs. We also gather a CV quality score from the hiring managers for a subset of jobs on the platform for the CVs of the 1,470 jobseekers matched to those jobs. Second, respondents see only the vacancies to which they match. This generates a well-defined jobseeker-vacancy unit of analysis that we use throughout the paper, and refer to as a *match*. This is not possible on platforms that allow unrestricted search, as every jobseeker can apply to any vacancy on the platform and the researcher may not observe which vacancies the jobseeker has seen. This makes it difficult to distinguish between vacancies a jobseeker sees but decides not to apply to and vacancies she has not seen at all.

The set of matches jobseekers receive are based on information collected during platform sign-up. However, jobseekers can contact the platform to update their education, experience, or occupation preferences at any time, including after treatment occurs. They can also ask to pause or stop receiving matches. This might in principle create a sample selection problem for the match-level dataset. But we show in Appendix B.4 that updates are rare, so there is little selection and correcting it does not affect our findings.

2.4 Platform Use

We highlight four important patterns of platform use, using the control group statistics in Tables 1 and 2. First, most matches do not generate applications: the average jobseeker submits only 0.23 applications and applies to 0.2% of matches they receive. The application count is unsurprisingly right-skewed: 74% of jobseekers submit zero applications and 5% submit more than 5 applications. Column 4 of Table 1 shows that, within our sample, jobseekers who do and do not actively use the platform differ on baseline characteristics. We discuss what this implies for interpreting our experimental results in Section 4.4. This application rate may strike some readers as surprisingly low. However, it is unsurprising that most matches do not generate applications. A match simply means the jobseeker qualified for the job and is interested in that occupation. In any search environment, jobseekers will apply to only a small subset of such jobs. The application rate is comparable to some other platforms in countries ranging from France to Mozambique. Furthermore, our sample

deliberately includes people who were not actively searching at baseline and includes all registered platform users. In contrast, some studies of job search platforms restrict their samples to only “active” users, defined in various ways, which naturally generates much higher application rates. See Table A.3 for details on application rates on other search platforms.

Second, the interview rate is low, but mainly because the application rate is low. The average jobseeker receives 0.014 interviews through the platform but each application has a 6.4% probability of generating an interview.¹³

Third, there is substantial variation in match value, and applications are directed to relatively high-value matches. For example, the standard deviation of monthly salary is roughly 9,200 Pakistani Rupees (275 USD PPP) and higher-salary vacancies get more applications (Table 2, column 2, row 1 and Figure C.3, panel A). At the match level, jobseekers are more likely to apply to vacancies where their work experience is a closer match (Table 2, column 2, row 5). Combining our available proxies for vacancy and match value in a single summary index shows that applications are substantially more likely for high-value matches (Table 2, column 2, row 8). This confirms that jobseekers can and do apply to higher-value matches, rather than randomly picking where to apply from relatively homogeneous matches, as random search models assume.

Fourth, however, control group jobseekers miss applying to many high-value matches. For example, jobseekers apply to only 00.45% of the matches in the top quintile of their within-jobseeker salary distributions (Figure C.3, panel A). This pattern also holds for the summary index of match value (Figure C.3, panel B).

These patterns naturally motivate our research. On the one hand, the facts that job applications are rare, even to high-value matches, and that applications have reasonably high interview probabilities suggest that lowering application costs could lead to more applications and substantially more interviews. On the other hand, the facts that jobseekers seem to choose strategically where to apply and that pecuniary and time costs of applying are already very low suggest that additional applications could go to relatively low-value matches and yield few interviews. Our experiment is designed to adjudicate between these two possibilities, both by identifying returns to additional

¹³As a benchmark, [Belot et al. \(2018\)](#) find that 3.6% of job applications submitted on a Scottish platform generate interview invitations. Other studies of platform-based job search do not report this ratio. Studies of off-platform job search in developing economies find > 10% of applications generate interviews, although we might expect a higher ratio for more expensive off-platform search ([Abebe et al., 2021a](#); [Banerjee & Sequeira, 2020](#); [Carranza et al., 2021](#)).

applications and by understanding which barriers deter additional applications in this setting.

2.5 Experimental Design and Interpretation

Our primary experiment varies a single element of communication with jobseekers in order to reduce the cost of applying for jobs on the platform: whether the platform initiates the application phone call or the jobseeker must do so. The platform sends text messages to all jobseekers, irrespective of treatment status, at the same time at the start of each monthly “matching round.” The text messages list the job title, firm name, firm location, and salary of each match received by the jobseeker that month and tell jobseekers to call the Job Talash number by a stated deadline if they want to apply. The deadline is on average ten days after the text message, with some variation between matching rounds due to operational factors such as platform staff capacity. When a jobseeker calls the platform, they are offered a free call back on the same day to complete the application process. The financial cost of placing the call to initiate the application is a maximum of 5 Pakistani rupees (0.03 USD PPP, less than 1% of a day’s earnings at minimum wage).

In the treatment condition, the call center *additionally* makes two attempts to phone each jobseeker and ask if they would like to initiate the application process. Roughly 50% of jobseekers are assigned to treatment for the duration of the experiment. Assignments are balanced on baseline jobseeker characteristics (Table 1, column 5).¹⁴ Treated jobseekers are called in a random order, starting as soon as the text messages are sent. Treatment is designed to minimize anticipation effects: treated jobseekers are told in initial matching rounds that they may not receive a phone call in every round, and should contact the call center if they wish to apply.

Importantly, the text message and phone call scripts contain identical information. The phone call scripts are also identical for the treatment and control groups. The only difference between the two is that the call center initiates the call for the treatment group. Call center agents are trained to not encourage or pressure jobseekers to apply at any moment during the call, and a supervisor audits the recording of at least one call per call center agent per matching round to ensure agents are following the script. Jobseekers can ask for more information about jobs on the calls but call center agents had access to no additional information in most matching rounds and we show in Section 4.4 that our findings are robust to omitting rounds when they had access to more information.

¹⁴Randomization took place within 82 strata based on the time that each geographic area completed household listing, platform sign-up, and the first round of matching.

We interpret treatment as a reduction in the cost of applying for jobs on the platform. In principle, these costs might be monetary (of airtime to initiate a call), time (of waiting for their call to get answered), or psychological (e.g. cognitive costs of processing vacancy information or fear of rejection). However, the platform is already designed to minimize the monetary and time costs jobseekers incur to initiate applications, and we show in Section 4.3 that additional experiments further reducing monetary and time costs produce substantially smaller effects on applications. Hence the most plausible interpretation of the phone call treatment is a reduction in the *psychological* cost of initiating an application. We develop this interpretation in more detail in Section 4.1, showing what this implies for treatment effects on applications and the returns to treatment-induced applications. We show in Section 4.4 that we can rule out several other interpretations based on the platform design, additional experiments we run, and additional survey measures.

Both the treatment and control conditions on Job Talash have many similarities to other large job search platforms. On Job Talash and these platforms, users can choose to receive notifications about jobs that match their qualifications and preferences and can apply using platform-generated CV templates. On most other platforms, users submit applications online or using phone apps. These are different technologies to Job Talash’s text-and-phone approach but they also allow scope for higher or lower psychological costs of initiating applications. For example, platforms can present information about matched jobs in ways that impose higher or lower attention costs. See Table A.2 for a more detailed comparison of application processes on different platforms.

3 Search Effort and Returns to Search

In this section we first show that the phone call treatment substantially increases the number of job applications and interviews. We then combine these results in a two-stage least squares framework to show that marginal applications submitted due to treatment yield interviews with roughly the same probability as inframarginal applications submitted without treatment, and yield interviews for vacancies of similar quality. These results imply roughly constant returns to the additional search effort induced by the treatment.

3.1 Treatment Effects on Search Effort and Search Outcomes

We run all analysis at the level of the jobseeker $\times$ vacancy match. As described in Section 2, each jobseeker only learns about vacancies that match their occupational preferences, education, and work experience, so these matches provide a well-defined unit of observation. We first estimate:

$$Y_{jv} = T_j \cdot \Delta + \mu_b + \epsilon_{jv}, \quad (1)$$

Y_{jv} is either an indicator for jobseeker j applying to vacancy v or an indicator for jobseeker j being invited to an interview for vacancy v . μ_b is a fixed effect for the stratification blocks within which treatment was randomized (see footnote 14). We estimate heteroskedasticity-robust standard errors clustered by jobseeker, the unit of treatment assignment.

Treatment leads to a large increase in job applications. Treated respondents apply to 1.32 percentage points (p.p.) more matches with standard error 0.08 p.p. (Table 3, column 1). This effect is seven times larger than the control group’s application rate of 0.19%. Treatment effects decline through time but remain positive for at least four years after jobseekers register for the platform. As a result, at the jobseeker level, treatment shifts the entire distribution of the number of applications to the right (Figure B.1). In particular, treatment increases the proportion of jobseekers who ever apply to a vacancy on the platform from 12 to 38%.

Treatment also increases the probability of getting an interview by 0.077 p.p. with a standard error of 0.009 p.p. (Table 3, column 2). This effect is nearly seven times larger than the control group’s 0.012% share of jobseeker $\times$ vacancy matches that generate interviews. At the jobseeker level, treatment also shifts the entire distribution of the number of interview invitations to the right (Figure B.1). The interview data are collected from firms, not jobseekers, and firms are unaware of respondent-level treatment assignments. So using firm reports of interview invitations minimizes measurement error from experimenter demand effects.¹⁵

The treatment effects on both applications and interview invitations are broad-based. Treatment substantially raises job application and interview rates for people who were employed and not employed at baseline, searching and not searching at baseline, and with above- and below-median education and age (Table B.1). This suggests that the economic behavior driving the treatment

¹⁵A few firms do not provide the list of jobseekers they interviewed. We assume no jobseekers matched to these vacancies are interviewed. Our main results are unchanged if we instead code these interview values as missing.

effects, which we discuss in Section 4, occurs across many types of jobseekers.

The treatment effects on applications and interviews are robust to a range of checks we show in Appendix B.2, including different ways of handling fixed effects, conditioning on baseline covariates, reweighting the data to give equal weight to each jobseeker rather than each jobseeker $\times$ vacancy match, and accounting for pauses in receiving matches that some jobseekers request.

3.2 Returns to Inframarginal Search and Treatment-Induced Marginal Search

To evaluate the returns to search, we estimate the relationship between the treatment effects on applications and interviews using an instrumental variables approach. We estimate the system:

$$\text{Apply}_{jv} = T_j \cdot \alpha + \mu_b + \epsilon_{jv} \quad (2)$$

$$\text{Interview}_{jv} = \text{Apply}_{jv} \cdot \beta + \eta_b + \varepsilon_{jv} \quad (3)$$

β recovers the local average effect of a treatment-induced application on the probability of an interview (LATE) under four conditions: treatment should be independent of all other factors influencing applications and interviews (independence), influence applications (strength), influence interviews only through applications (exclusion), and increase the probability of application for all respondents (monotonicity). The independence condition holds by random assignment and the preceding results show that the strength condition holds. We discuss potential complications with the monotonicity and exclusion conditions and how we address them at the end of this subsection.

Marginal applications submitted due to treatment have roughly the same return as inframarginal applications, measured in terms of interview invitations. The LATE estimate shows that the average treatment-induced application has a 5.9% probability of an interview invitation with standard error 0.5 (Table 3, column 4, row 2). This is very similar to the 6.4% mean interview probability for control group applications and we cannot reject equality of the probabilities ($p = 0.573$). As we discuss further below, this is not a consequence of low power.

Marginal and inframarginal applications also have equal returns measured in ‘value-weighted’ interviews. This finding is important, as the return to an application, and the decision to apply, reflects both the probability of an interview P and the value of an interview V . To show this, we construct a proxy index V_{vm} for the value of each match a jobseeker receives: an inverse-covariance weighted average of positive attributes of the vacancy and match, such as salary and commuting

distance, defined in detail in the note below Table 3. We estimate the system (2)-(3), replacing the second stage outcome with an interaction between the interview invitation indicator and the proxy index. This gives us the local average treatment effect on $P \cdot V$. The returns to inframarginal and marginal search using this measure are again very similar: respectively 0.21 and 0.24, with $p = 0.413$ for the test of equality (Table 3, column 5). We repeat this value-weighting exercise using each individual proxy for interview value and fail to reject equality of marginal and inframarginal applications' value-weighted interview outcomes for all eleven proxies (Table B.2).

The finding of roughly constant returns on both interviews and value-weighted interviews is not a mechanical consequence of a matching algorithm or labor market that ensures homogeneous returns. Instead, as we explain in Section 2, most jobseekers are matched with vacancies from multiple occupations and with firms that prefer different types of work experience and education. This creates scope for heterogeneous returns from applying to different types of matches. Furthermore, Table B.1 shows that the constant returns finding also holds for jobseekers with above-median education and who were employed at baseline. They match to a broader set of jobs, giving them more scope to direct applications widely, making the constant returns finding more surprising.

The finding of roughly constant returns is also not a consequence of low power. The return to marginal applications is precisely estimated, with a 95% confidence interval of 4.8 to 6.9 p.p. for interview invitations. Relative to the interview rate of 6.4% for inframarginal applications, we can reject decreases of more than 1.5 p.p. and increases of more than 0.5 p.p. Even the lower bound of the confidence interval implies a decrease of only $1.5/6.4 = 24\%$ in the average interview probability over a 615% increase in the application rate, implying a slowly decreasing return to search effort. A similar pattern holds for the returns measured in value-weighted interviews. We do not, of course, claim that returns would be constant over all possible levels of search effort and acknowledge that returns may be substantially lower with very high search effort.¹⁶

Before proceeding, we briefly discuss an extensive battery of robustness checks on the constant returns finding, shown in detail in Appendices B.2 - B.4. First, we address the possibility that treatment increases applications from some jobseekers and decreases applications from others, which

¹⁶As a very speculative back-of-the-envelope calculation, we can estimate a linear returns curve using the control group means and treatment effects for the application and interview rates. We can then use the estimated curve to extrapolate the marginal interview probability at even higher application rates. The estimated curve is relatively flat. For example, if the share of matches generating applications increased 25 fold, from 0.185% to 4.625%, then the linear extrapolation implies that the interview probability for the marginal application would only drop from 5.4 to 3.2%.

would violate the monotonicity condition used in our IV analysis. To do this, we derive a bound on the bias from violations of monotonicity in these data, following [De Chaisemartin \(2017\)](#). This implies that a bias-corrected LATE of applications on interviews is bounded between 4.4 and 5.9%. Second, we address the possibility that treatment affects both the quantity and quality of applications, which would complicate the exclusion restriction used in our IV analysis. All application content is sent by the Job Talash platform using template CVs. We show that treatment effects on measures of application quality that jobseekers can change by updating information used in their CV templates are close to zero. Third, we address the possibility that treatment affects which matches jobseekers receive, which would create a sample selection problem because we use each jobseeker $\times$ vacancy match as a unit of analysis. This can only occur if treatment causes jobseekers to update the information used to match them to vacancies: their occupational preferences, education, or experience. We show that treatment has little impact on updating this information and that our key results are unchanged when we use a sample consisting of the counterfactual set of matches that would have been generated in the absence of these updates. Fourth, we use a non-IV approach to compare the returns to marginal and inframarginal applications under different assumptions, which also generates similar estimates of returns. Finally, we show that our key findings are robust to different ways of handling fixed effects and conditioning on baseline covariates, including allowing interactions between treatment assignment and the fixed effects.

We focus on interviews and value-weighted interviews as outcomes because these take advantage of the strengths of the platform we study. The platform gives us detailed data at the level of jobseeker $\times$ vacancy matches: all vacancy characteristics observed by the jobseeker, all jobseeker characteristics observed by the firm, application decisions, and interview invitations. These data allow us to precisely describe how search decisions are made and the consequences of those decisions up to the interview stage. Interviews are also a key search outcome because they are a necessary condition for job offers, impose non-trivial costs on both job applicants and firms, and provide learning opportunities for jobseekers. Hence their widespread use as central outcomes in areas of labor economics such as audit studies, highlighted in the review by [Neumark \(2018\)](#).

The disadvantage of platforms is that they do not generally provide data on employment outcomes, so evaluations relying on employment outcomes require off-platform data. To this end, we survey jobseekers about their employment and find a treatment effect on self-reported employ-

ment of 1.1 percentage points, with a standard error of 2 p.p (Table B.7, column 4). However, this estimate should be interpreted very cautiously for three reasons. First, the surveys take place on average 40 months after randomization, so they capture the effects of multiple years of ongoing exposure to treatment rather than immediate effects. Second, despite extensive tracking effort, the survey response rate is 47% and differs between treatment and control groups, which could produce sample selection bias. To address this, we randomize some features of the survey data collection, e.g., number of call attempts, and use this to create instruments for a sample selection correction term, following [DiNardo et al. \(2021\)](#). We describe the selection correction method and how the randomized survey features influence response rates in detail in Appendix B.6. Third, and perhaps most importantly, we are underpowered to study treatment effects on employment at the scale of this experiment. In particular, the 95% confidence interval on the estimated treatment effect covers -2.9 to 5.0 p.p. (Table B.7, column 4). Even with a 100% survey response rate, the minimum detectable effect size on employment would be 2.4 p.p. The phone call treatment increases the share of jobseekers receiving any interview invitations by 4.7 p.p., so an employment effect of 2.4 p.p. would be achieved in the possible but unlikely event that half of these jobseekers converted their additional interview into a job.¹⁷

These calculations suggest that the strength of light-touch treatments like this is the possibility of modestly raising employment rates on larger platforms at very low marginal costs. For example, Pakistan’s Rozee has 9.5 million users, 1000 times the size of our platform. *If* a treatment like the one we study could raise the share of respondents getting interviews by the same 4.7 p.p. and *if* only 5% of these additional interviews converted into offers, that would lead to roughly 22,000 offers. As [Kircher \(2022\)](#) notes, many other studies of interventions on job search platforms either do not study employment effects at all (see examples in footnote 3) or use samples of hundreds of thousands of jobseekers to detect effects of 1 percentage point or smaller (e.g. [Behaghel et al. 2023](#); [Le Barbanchon et al. 2023](#)).¹⁸

¹⁷To derive this minimum detectable effect size, we assume: 80% test power, 5% test size, mean employment of 73% (equal to the control group’s reported employment rate in the survey), and that covariates can absorb 10% of the outcome variation (roughly what we see in the interview invitation regressions). We preregistered employment and employment characteristics as trial outcomes because we did not know at the time (July 2020) how much COVID-19 would constrain platform operation, survey data collection, and hence power. We did not collect survey data on employment characteristics or proxies for match quality once COVID-19 made it apparent that we would not be powered to study treatment effects on these outcomes.

¹⁸The latter studies are based exclusively in high-income countries where data can be linked between government-

Our survey also asked jobseekers about their off-platform search behavior, after 40 months of treatment exposure. Treatment effects on any search, number of applications, and number of search methods used are negative and 4-17% of the control group means but with wide confidence intervals that include zero. So we again recommend great caution when interpreting the estimated effects (Tables B.7 and B.8).

4 Explaining Marginal Returns to Search

Our finding of roughly constant returns to job search raises a puzzle: why do jobseekers not apply to more jobs in the absence of treatment, especially given the seemingly low cost of applying on the platform? In this section, we develop a simple conceptual framework that can explain both the large treatment effect on applications and the roughly constant returns to treatment-induced applications. We then present several empirical results to support this framework, better understand the nature of application costs, and argue against alternative frameworks. We summarize the empirical analysis relatively briefly in this text and provide detailed explanations of the methods and results in Appendix C.

4.1 Conceptual Framework

Here we present a brief, intuitive discussion of our conceptual framework, with the formal model left to Appendix C.2. This paper’s contribution is empirical rather than theoretical, so the framework is deliberately simple and stylized.¹⁹ This framework shows how the common assumption of decreasing returns to marginal applications for *each individual jobseeker in each period* can lead to constant returns *averaged over jobseekers and periods*, provided some jobseekers are not actively searching in some periods.

The platform sends each jobseeker a monthly batch of matches. We begin with a standard assumption (A1) that the jobseeker applies to all matches whose expected gross return, PV , exceeds the cost of applying. P is the probability of an interview conditional on applying. V is the gross value of getting an interview, which captures the expected present value of the flow of future utility from the interview, including the potential for a job offer.

run job search platforms and unemployment benefit registers. This is not currently feasible in any developing country, including the one we study.

¹⁹This framework has a similar spirit to recent models of ‘partially directed search,’ in which jobseekers want to apply to the highest-return matches but miss some high-return matches due to frictions (Lentz et al., 2022; Wu, 2021).

Our key assumption (A2) is that the cost of applying varies within jobseeker through time and/or between jobseekers, and can be high enough that some untreated jobseekers choose not to apply in some matching rounds. We write this section in terms of costs that vary within jobseeker through time to simplify the writing but the full explanation of the framework in Appendix C.2 allows both types of cost variation. Figure 1 shows application behavior by untreated jobseekers under assumptions (A1) and (A2): jobseekers facing low costs in that month apply to matches with PV above PV_{L0} (the light-shaded area in panel A), while jobseekers facing high costs apply to no matches (panel B).

In this framework, there are two types of marginal applications induced by treatment. The first type of marginal applications comes from jobseekers facing low costs at the time, who would apply to at least one match in that round even without treatment. Treatment lowers their cost of applying, so they apply to matches with PV above PV_{L1} (the light- and dark-shaded areas in panel A). These marginal applications have strictly lower returns than the inframarginal applications. The second type of marginal applications comes from jobseekers facing high costs at the time, who would not apply to any matches in that round without treatment. Treatment lowers their cost of applying, so they apply to matches with PV above PV_{H1} (the dark-shaded area in panel B). These marginal applications will have higher returns than the inframarginal applications if the cost reduction due to treatment is small relative to the cost variation within the control group.

The treatment effect on applications and return to marginal applications are averages across these two types, weighted by their relative size. The large effect on applications relative to the control group mean suggests that many more jobseekers face high application costs at each time than low. The roughly equal returns to marginal and inframarginal applications can occur if the lower marginal return to applications from low-cost jobseekers (panel A) are offset by the potentially higher marginal return to applications from the more numerous high-cost jobseekers (panel B). We show this formally in Appendix C.2 and explain that the framework does not require the simplifying assumption of only two cost types.²⁰

This framework provides a clear economic interpretation of the LATE we estimate in Section 3.2: it is the average effect of applying on interview invitations, for applications sent due to a

²⁰This framework allows the possibility of decreasing returns to marginal applications for treatments that decrease the application cost by more. These would lead to very large increases in application rates and to $PV_{H1} < PV_{L0}$.

treatment-induced drop in the cost of applying. In this framework, marginal applications come from jobseekers who face higher costs of applying in the absence of treatment, relative to jobseekers submitting inframarginal applications. The constant returns finding shows that these higher costs are not associated with lower returns to applications.²¹

4.2 Additional Tests of the Conceptual Framework

This framework delivers three additional predictions that we can test. First, *control group jobseekers will not apply to some high-value vacancies*, because some of them face high application costs, either permanently or during some matching rounds. Second, *treatment and control group applications will go to vacancies with similar average values*, because treatment will induce applications to a mix of higher- and lower-return matches whose average value is similar to the control group.²² Third, *treatment group applications will go to vacancies with more dispersed values*, as shown by the wider range of PV in the light+dark region versus the dark-only region in Figure 1.

We show evidence that is not inconsistent with all three predictions, summarized here with the detailed methods and results in Appendix C.4. We use two proxies for the value of each jobseeker $\times$ vacancy match: the salary, an admittedly narrow proxy but one that is easily interpretable and valued by all jobseekers; and the inverse covariance-weighted average of many positive attributes of a match (e.g. salary, benefits, commute distance) that is defined in Section 3.2. Figures C.2 and C.3 show evidence consistent with the first prediction: that many high-value matches receive few or no control group applications. For example, under half of all control group applications are sent to matches in the top quintile of each value proxy, and under 0.1% of matches in the top quintile receive applications. Figures C.2 and C.3 also show evidence consistent with the second prediction: that treatment and control group applications will go to vacancies with similar average values. Specifically, they show that the shares of control group applications sent to each of quintiles 1, ..., 5 are equal to the shares of treatment group applications sent to each of these quintiles, with formal test results reported in the footnotes below the figures. As an additional test of the second prediction, we use the same proxies to calculate the mean values of matches that get appli-

²¹This echoes a finding in education research that costs of education and returns to education are weakly correlated over individuals in some applications (Card, 2001).

²²Technically, this prediction holds in the special case of the framework where returns to marginal and inframarginal applications are equal, as we see in our data. When returns to marginal and inframarginal returns differ, then treatment and control applications may be sent to vacancies with different average values, as we explain in Appendix C.2.

cations and compare these between the treatment and control groups. Table C.5 shows that there are some differences between mean values of observed characteristics between treatment and control group applications but that these differences do not show consistently higher values in either group, consistent with the second prediction. For example, control group applications go to jobs that offer slightly higher salaries, slightly less flexible hours, similar values of the summary index V_{vm} , and similar latent interview probabilities (which we predict using a LASSO-based approach). Third, we use the same proxies to calculate the dispersion of values of matches receiving applications and compare these between the treatment and control group. Table C.6 shows that treatment raises the variance and lowers or leaves unchanged the 10th percentiles for the value proxies. This qualitatively matches the model predictions but the effects are small and not close to statistically significant, giving limited support for the prediction that treatment leads to greater dispersion in salaries.

4.3 Understanding Application Costs

The roughly constant returns to marginal applications shown in Section 3.2 and the patterns of matches that receive applications shown in Section 4.2 are both consistent with treatment reducing the cost of initiating applications. In this subsection, we report the results of several additional experiments designed to shed light on what types of application costs are likely to fall with treatment. We show more details on all methods and results in Appendix C.5. Each treatment in these experiments is assigned to a small share of the sample, and controlling for these assignments and their interactions has no impact on the estimated effects of the main phone call treatment (Table C.1).

First, *pecuniary costs* of applications are unlikely to explain our main results. Job applications on the platform are very cheap, as the control group can call the platform to apply for < 1% of the daily minimum wage, and mobile phone providers in Pakistan allow users to fund phone calls with short-term loans when they have no airtime credit. We also randomly select some control group jobseekers in one round to receive a text message reminder that they can ask the platform to call them back about a job posting, saving the cost of calling back entirely. This free callback reminder treatment has an effect one hundredth of the size of the effect of the main phone call treatment, suggesting a very small role for pecuniary costs of job search (Table C.7).

Similarly, *time costs* of applications are unlikely to explain our main results. We randomly offer some control group jobseekers in some rounds the option to text the platform and ask for a callback at a specific time of their choice. This eliminates time waiting on hold or connecting to a call center agent. The effect of this callback request treatment is one quarter the size of the main phone call and statistically significantly smaller. This suggests that while time costs matter, they matter much less than the overall effect of the phone call treatment (Table C.7).

Given the limited role for pecuniary and time costs of applying, we view psychological costs of initiating applications as the most likely explanation for our main results. As we discuss in the introduction, the existing literature suggests five types of psychological costs that might be reduced by the phone call treatment. First, control group jobseekers might ignore text messages due to *attention costs* (Gabaix, 2019). Second, control group jobseekers might not initiate applications due to *fear of rejection* (Köszegi et al., 2022). Third, *present bias* might lead control group jobseekers to repeatedly postpone applications until the deadline passes (Ericson & Laibson, 2019). Fourth, phone calls may function as *reminders* to apply. Fifth, phone calls may *encourage or pressure* jobseekers to apply. There is existing empirical evidence that some of these factors can influence job search decisions (e.g. DellaVigna & Paserman 2005; Zizzamia 2023), as well as a larger body of research reviewed by DellaVigna (2009) showing that eliminating or reducing the need to initiate decisions can raise financial and health investments. We show in Appendix C.2 how each of these factors can enter our model.

While we cannot pin down exactly how the phone call treatment reduces psychological costs of initiating applications, we can test and largely reject two of these five possible psychological explanations. First, treatment does not simply increase application rates by providing a *reminder effect* that offsets forgetfulness. We randomly send some control group jobseekers in some rounds a second text message listing their matched vacancies as a reminder. The effect of the reminder is one-fourteenth as large as the effect of the phone call (p -value of difference < 0.001 , Table C.8). Furthermore, the phone call treatment has a smaller effect on the application rate when phone calls take place longer after text messages and when the window between text messages and application deadlines is shorter (Table C.9). This pattern is not consistent with an important role for reminder effects, as earlier calls and shorter application windows allow less time for forgetting and hence less scope for reminder effects.

Second, we find some evidence against an explanation that treatment increases applications because call center agents *encourage or pressure jobseekers to apply*. Agents are trained not to *explicitly* encourage or pressure jobseekers, and regular audits of call recordings verified that they followed their scripts. It remains possible that jobseekers feel *implicit* pressure to apply because they have been called or because they are interacting with a person. However, if treated respondents did feel pressure to apply when called, they could easily avoid this pressure by avoiding calls after the first few rounds of matching. Instead, we find that jobseekers who answer calls in the first few rounds of matching are actually more likely to answer calls in subsequent rounds, conditional on observed characteristics (Table C.10). Furthermore, treated and control jobseekers are equally likely to apply to the first job listed on their call or text message, showing that treated jobseekers do not simply apply to the first listed job to end a pressure-inducing call quickly (Figure C.7).

This collection of results implies that the phone call treatment produces more applications at roughly constant returns by reducing the psychological rather than the pecuniary or time costs of applying, and probably not by providing reminders, encouragement, or pressure. However, we acknowledge that we cannot pin down exactly which psychological cost(s) are reduced by the phone call treatment, so alternative explanations remain possible.

4.4 Evaluating Alternative Explanations

In this section we summarize evidence against five alternative explanations for our two key findings: that the phone call treatment substantially increases the job application rate and that the returns to these marginal applications are approximately constant, in terms of interview probabilities. We show more details on all methods and results in Appendix C.6. First, treatment-induced job applications are not sent to systematically *less competitive or desirable vacancies* than control group job applications. Instead, we have already shown in Section 4.2 that treatment and control group job applications are sent to equally desirable matches on average. This allows us to reject an alternative explanation in which marginal applications and inframarginal applications might have roughly equal returns if marginal applications are sent to vacancies that are both systematically worse matches for the jobseekers (leading to lower interview probabilities) and systematically less competitive (leading to higher interview probabilities).

Second, treatment-induced job applications do not come from *systematically better-qualified*

jobseekers than control group applications. Instead, applications in the treatment and control groups come from jobseekers with roughly equal values of observed characteristics such as education, work experience, and CV quality scores (Table C.11). We also use a data-driven approach to estimate the latent probability that an application sent to each match by each matched jobseeker would yield an application, based on the jobseeker characteristics that both we and the firm observe. Matches that receive applications from the treatment and control groups have similar values of this latent probability, suggesting that treatment does not shift the selection of which types of jobseekers submit applications (Table C.11). Furthermore, our main findings hold when we control for time-invariant characteristics in two ways. We use an additional “crossover” experiment that randomly moves some control group jobseekers to the treatment group in some rounds, which allows us to estimate treatment effects conditional on jobseeker fixed effects (Table C.12).²³ And we repeat our main analysis controlling for observed baseline characteristics.²⁴ This allows us to reject an alternative explanation in which marginal and inframarginal applications might have roughly equal returns because each individual jobseeker experiences decreasing returns to additional search effort but treatment-induced applications come from jobseekers who are positively selected on education, experience, etc.

Third, it is unlikely that phone calls provide *more information about specific jobs*. Call center agents are trained to read specific scripts with no additional information about jobs, general labor market conditions, or assessments of the jobseeker’s prospects, and regular audits of call recordings verified that they followed their scripts. Additionally, in 80% of matching rounds, we gave call center agents no additional information about the jobs and our results are almost unchanged when we restrict analysis to these rounds (Table C.14). It is possible that jobseekers might be more likely to receive phone calls than text messages, perhaps if some text messages are blocked or go unread. However, when we survey jobseekers and ask if they remember receiving a recent job match from the platform by either phone call or text message, treatment and control jobseekers

²³In particular, we cannot reject equality of the interview rates or quality-adjusted interview rates for inframarginal applications and marginal applications submitted due to the crossover treatment ($p > 0.430$). 16% of jobseekers have at least one match affected by this crossover treatment, allowing precise estimation of the treatment effect conditional on the fixed effects. But only 0.65% of matches are affected by this treatment, so it has almost no impact on our estimates of the overall treatment effect (Table C.13).

²⁴To show this, we repeat our analysis of the main experiment using a post-double selection LASSO to control for an extensive set of jobseeker baseline characteristics, following [Belloni et al. \(2014\)](#). Table B.3 shows that the point estimates and standard errors are almost identical.

are equally likely to report that they received matches, with or without adjusting for survey non-response (Table C.15). This allows us to reject an alternative explanation in which the phone call treatment might provide information about specific jobs, leading to higher application rates and enabling jobseekers to target better-matched vacancies that have higher interview probabilities, keeping the returns to marginal applications as high as inframarginal applications.

Fourth, it is unlikely that the phone call treatment increases job application rates by *shifting jobseekers' beliefs about the value of applying*. It is possible that calls from a professional recruiting service might be taken as signals that platform firms are larger or wealthier and thus able to provide more benefits or opportunities for advancement (higher V), or that these are jobs to which the jobseeker is particularly well-matched and likely to get an interview (higher P). But this explanation does not match several patterns in the data. We directly test this by collecting data on jobseekers' beliefs about P and V and estimating treatment effects on these two belief measures.²⁵ Treatment effects on both these measures are close to zero, with or without adjusting for survey non-response (Table C.16). Furthermore, if phone calls influence job applications because a jobseeker views them as informative about the quality of a specific match, then phone calls should have larger effects on applications when the jobseeker views the phone call as unusual than when she views it as part of normal platform operations. But when control group jobseekers receive occasional phone calls as part of our within-jobseeker “crossover” experiment, their response is very similar to the main phone call treatment. See details on the experiment two paragraphs above and results in Table C.12. These patterns suggest that the phone call is unlikely to shift application decisions by signaling that these are unusually high-value matches. This allows us to reject an alternative explanation in which the phone call treatment might increase application rates by raising the perceived value of applying.

Fifth, the main experimental results do not arise because jobseekers search randomly, which might lead to constant returns to additional applications. Random job search may seem implausible but it is often assumed in canonical search models (e.g. [Pissarides 2000](#)) and may be plausible when jobseekers have very limited information about labor market conditions and match quality

²⁵Translated from Urdu, these questions ask: “Suppose Job Talash sends you one hundred job ads in the next year. Based on your past experience with our job matching service, how many of these jobs do you think would be desirable for you?” and “Suppose you apply for all of these jobs that you think are desirable. How many do you think would make you an offer?” Our main treatment assignment is time-invariant, so these questions are asking jobseekers about jobs sent by the mode of communication used in their treatment group.

(Behaghel et al., 2023; Belot et al., 2018). However, random search is not consistent with the pattern we showed in Sections 2.4 and 4.2 that applications are directed toward vacancies with more desirable attributes. We also run an experiment to directly induce random job search and show that this produces different results to the phone call treatment. Specifically, in 20% of rounds we randomize the order in which vacancies are listed in both text messages and phone calls, which generates a random increase in the rate of applications to vacancies listed first. Unlike jobs to which individuals are encouraged to apply because of the phone call treatment, applications induced by jobs being listed first have decreasing rather than constant returns. The average interview probability for these marginal applications is 2.6%, substantially lower than the 6.3% for inframarginal applications (Table C.17). This allows us to reject an alternative explanation in which the phone call treatment generates marginal applications with constant returns because jobseekers are searching randomly, so marginal and inframarginal applications are sent to similar vacancies.

5 Spillover Effects

Increased search effort by some jobseekers may affect firms and other jobseekers. The sign of this effect is theoretically ambiguous. For firms, getting more applications can raise the probability of getting an application from a well-matched applicant and hence making a hire. But it can also generate congestion costs if firms need to review many poorly-matched applications. For other jobseekers, competing against more applications can lead to crowd-out. But the magnitude of crowd-out may be small and offset if firms are more likely to hire when they get more applications.

We can identify spillover effects using variation in the vacancy-level treatment rate: the share of users matched to each vacancy who are treated. This share is random because matches are determined by pre-treatment characteristics (education, work experience, and occupational preferences). This approach works well because the platform’s matching structure fully determines the set of platform users who can compete with each other for each vacancy. This approach is not feasible for jobseeker-facing experiments on most platforms, where users can search and apply for many different jobs. On such platforms, it is difficult to define how much each user is competing with other users without a full model of the job search process.

We briefly summarize our methods and results here and provide many more details in Appendix D. Within each of the 1,342 vacancies, we estimate the share of jobseekers matched to that vacancy

who are treated and the treatment effect on interview invitations. We create a vacancy-level dataset with these two statistics and regress the treatment effect on the treatment share, conditional on other vacancy-level characteristics. We find no evidence of a negative relationship: a vacancy exposed to the 75th percentile of the treatment rate rather than the 25th percentile would have a 0.018 percentage point higher treatment effect on interviews (standard error = 0.011, $p = 0.097$). This shows that treatment effects on individual jobseekers' interview probabilities are not smaller when they face more treatment-induced competition, suggesting they do not face negative spillovers. Similarly, we find no evidence of negative spillovers when we allow the relationship to be nonlinear or allow spillovers to have different effects on treated and control group jobseekers.

These results suggest negligible between-jobseeker spillovers on interview invitations. However interpretation of these patterns is complex and requires some caveats. Spillovers might be negligible because firms report filling only 60% of the vacancies posted on the platform, so more applications might lead to more well-matched applicants and hence fill more vacancies, in line with findings by [Fernando et al. \(2021\)](#). But spillovers might also be negligible because firms report receiving 70% of applications from outside the platform so the majority of competition that jobseekers face is unaffected by treatment. Finally, spillovers may be very different at the interview stage versus the hiring stage, which we do not observe. See Appendix D for more discussion.

6 Conclusion

We show that job search effort can be dramatically increased by reducing the psychological cost of initiating job applications. Perhaps surprisingly, returns to the additional search effort, measured in terms of interview invitations, are constant rather than decreasing. This pattern is consistent with a model in which marginal applications in any period are a mix of lower-return applications from jobseekers who would send some applications without treatment and higher-return applications from jobseekers who would not apply in that period without treatment. These findings show that small reductions in search costs can substantially improve search outcomes in environments with some relatively inactive jobseekers. This echoes findings that changing default options to avoid initiation costs can lead to economically significant increases in financial and health investments ([DellaVigna, 2009](#)). Our findings are particularly striking because this is a platform designed to have minimal pecuniary, time, and technology barriers to use and hence to be broadly accessible to

jobseekers in a low-resource setting. Yet psychological costs of initiating applications still present a significant barrier for jobseekers on the platform.

These findings link to a broader literature around the design of job search policy and platforms. The possibility that psychological costs lead to suboptimally low search effort has implications for policies such as using caseworkers to increase jobseekers' accountability and motivation, subsidizing job search, requiring active search for unemployment insurance recipients, or automatically enrolling jobseekers in search assistance services (Card et al., 2010, 2018). Job search and matching platforms could also encourage search by simplifying the process of evaluating job listings or encouraging jobseekers to start applications, although the value of such design changes depends on existing application volumes.

References

- Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin & S. Quinn (2021a) "Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City," *Review of Economic Studies*, 88 (3).
- Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin, S. Quinn & F. Shilpi (2021b) "Matching Frictions and Distorted Beliefs: Evidence from a Job Fair Experiment," Working paper, University of Oxford.
- Abebe, G., S. Caria & E. Ospina-Ortiz (2019) "The Selection of Talent: Experimental and Structural Evidence from Ethiopia," *American Economic Review*, 111 (6).
- Abel, M., R. Burger, E. Carranza & P. Piraino (2019) "Bridging the Intention-Behavior Gap? The Effect of Plan-Making Prompts on Job Search and Employment," *American Economic Journal: Applied Economics*, 11 (2), 284–301.
- Abel, M., R. Burger & P. Piraino (2020) "The Value of Reference Letters: Experimental Evidence from South Africa," *American Economic Journal: Applied Economics*, 12 (3), 40–71.
- Afridi, A., Farzana ad Dhillon, S. Roy & N. Sangwan (2022) "Social Networks, Gender Norms and Women's Labor Supply: Experimental Evidence using a Job Search Platform," Working paper, Indian Statistical Institute.
- Agrawal, A., J. Horton, N. Lacetera & E. Lyons (2015) "Digitization and the contract labor market: A research agenda," in Goldfarb, A., S. Greenstein & C. Tucker eds. *Economic analysis of the digital economy*, 219–250: University of Chicago Press.
- Alfonso Naya, V., G. Bied, P. Caillou, B. Crepon, C. Gaillac, E. Perennes & M. Sebag (2020) "Designing labor market recommender systems: the importance of job seeker preferences and competition," Manuscript, LISN.
- Altmann, S., A. Falk, S. Jäger & F. Zimmermann (2018) "Learning about Job Search: A Field Experiment with Job Seekers in Germany," *Journal of Public Economics*, 164, 33–49.
- Altmann, S., A. Glenney, R. Mahlstedt & A. Sebald (2022) "The Direct and Indirect Effects of Online Job Search Advice," IZA Discussion Paper 15830.
- Anderson, M. (2008) "Multiple Inference and Gender Differences in the Effects of Early Interven-

- tion: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects,” *Journal of the American Statistical Association*, 103 (484), 1481–1495.
- Arni, P. & A. Schiprowski (2019) “Job search requirements, effort provision and labor market outcomes,” *Journal of Public Economics*, 169, 65–88.
- Babcock, L., W. J. Congdon, L. F. Katz & S. Mullainathan (2012) “Notes on behavioral economics and labor market policy,” *IZA Journal of Labor Policy*, 1, 2.
- Baker, S. R. & A. Fradkin (2017) “The Impact of Unemployment Insurance on Job Search: Evidence from Google Search Data,” *The Review of Economics and Statistics*, 99, 756–768.
- Bandiera, O., V. Bassi, R. Burgess, I. Rasul, M. Sulaiman & A. Vitali (2021) “The Search for Good Jobs: Evidence from a Six-year Field Experiment in Uganda,” Working paper.
- Banerjee, A. & S. Sequeira (2020) “Spatial Mismatches and Imperfect Information in the Job Search,” Discussion Paper 14414, Centre for Economic Policy Research.
- Banfi, S., S. Choi & B. Villena-Roldan (2019) “Deconstructing Job Search Behavior,” Working Paper 3323545, Social Science Research Network.
- Bassi, V. & A. Nansamba (2020) “Screening and Signaling Non-Cognitive Skills: Experimental Evidence from Uganda,” Manuscript, University of Southern California.
- Beam, E. (2016) “Do Job Fairs Matter? Experimental Evidence from the Philippines,” *Journal of Development Economics*, 120, 32–40.
- Behaghel, L., S. Dromundo, M. Gurgand, Y. Hazard & T. Zuber (2023) “The Potential of Recommender Systems for Directing Job Search: A Large-Scale Experiment,” Manuscript, Paris School of Economics.
- Belloni, A., V. Chernozhukov & C. Hansen (2014) “Inference on treatment effects after selection among high-dimensional controls,” *The Review of Economic Studies*, 81 (2), 608–650.
- Belot, M., P. Kircher & P. Muller (2018) “Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice,” *Review of Economic Studies*, 86 (4), 1411–1447.
- (2022a) “Do the Long-term Unemployed Benefit from Automated Occupational Advice during Online Job Search?”, Working paper, Cornell.
- (2022b) “How Wage Announcements Affect Job Search—A Field Experiment,” *American Economic Journal: Macroeconomics*, 14 (4), 1–67.
- Ben Dhia, A., B. Crepon, E. Mbih, L. Paul-Delvaux, B. Picard & V. Pons (2022) “Can a Website Bring Unemployment Down? Experimental Evidence from France,” Working paper.
- Bhuller, M., A. Kostøl & T. Vigtel (2019) “How Broadband Internet Affects Labour Market Matching,” Working paper.
- Bolhaar, J., N. Ketel & B. van der Klaauw (2020) “Caseworker’s discretion and the effectiveness of welfare-to-work programs,” *Journal of Public Economics*, 183, 104080.
- Boudreau, L., R. Heath & T. McCormick (2022) “Migrants, Experience, and Working Conditions in Bangladeshi Garment Factories,” Working paper.
- Brown, M., C. J. Flinn & A. Schotter (2011) “Real-Time Search in the Laboratory and the Market,” *American Economic Review*, 101 (2), 948–74.
- Caliendo, M., D. Cobb-Clark & A. Uhlendorff (2015) “Locus of Control and Job Search Strategies,” *The Review of Economics and Statistics*, 97 (1), 88–103.
- Card, D. (2001) “Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems,” *Econometrica*, 69 (5), 1127–1160.
- Card, D., J. Kluve & A. Weber (2010) “Active Labour Market Policy Evaluations: A Meta-Analysis,” *The Economic Journal*, 120, F452–F477.

- (2018) “What Works? A Meta-Analysis of Recent Active Labor Market Program Evaluations,” *Journal of the European Economic Association*, 16 (3), 894–931.
- Caria, S., G. Gordon, M. Kasy, S. Quinn, S. Shami & A. Teytelboym (2023) “An Adaptive Targeted Field Experiment: Job Search Assistance for Refugees in Jordan,” *Journal of the European Economic Association*, forthcoming.
- Caria, S., K. Orkin, A. Andrew, R. Garlick, R. Heath & N. Singh (2024) “Barriers to Search and Hiring in Urban Labour Markets,” *VoxDevLit*, 10 (1).
- Carranza, E., R. Garlick, K. Orkin & N. Rankin (2021) “Job Search and Hiring with Limited Information about Workseekers’ Skills,” Manuscript, Duke University.
- Carranza, E. & D. McKenzie (2024) “Job Training and Job Search Assistance Policies in Developing Countries,” *Journal of Economic Perspectives*, 38 (1), 221–244.
- Carroll, G., J. CHoi, D. Laibson, B. Madrian & A. Metrick (2009) “Optimal Defaults and Active Decisions,” *Quarterly Journal of Economics*, 124 (4), 1639–1674.
- Chakravorty, B., A. Y. Bhatiya, C. Imbert, M. Lohnert, P. Panda & R. Rathelot (2023) “Impact of the COVID-19 Crisis on India’s Rural Youth: Evidence from a Panel Survey and an Experiment,” *World Development*, 168.
- Chiplunkar, G. & P. K. Goldberg (2022) “The Employment Effects of Mobile Internet in Developing Countries,” Working Paper 30741, NBER.
- Cooper, M. & P. Kuhn (2020) “Behavioral Job Search,” in Zimmermann, K. F. ed. *Handbook of Labor, Human Resources and Population Economics*, 1–22, Cham: Springer International Publishing.
- Dammert, A., J. Galdo & V. Galdo (2015) “Integrating Mobile Phone Technologies into Labor-market Intermediation: a Multi-treatment Experimental Design,” *IZA Journal of Labor and Development*, 4.
- De Chaisemartin, C. (2017) “Tolerating defiance? Local average treatment effects without monotonicity,” *Quantitative Economics*, 8, 367–396.
- DellaVigna, S. (2009) “Psychology and Economics: Evidence from the Field,” *Journal of Economic Literature*, 47 (2), 315–72.
- DellaVigna, S., J. Heining, J. Schmieder & S. Trenkle (2022) “Evidence on Job Search Models from a Survey of Unemployed Workers in Germany,” *Quarterly Journal of Economics*, 137, 1181–1232.
- DellaVigna, S., A. Lindner, B. Reizer & J. F. Schmieder (2017) “Reference-Dependent Job Search: Evidence from Hungary,” *The Quarterly Journal of Economics*, 132, 1969–2018.
- DellaVigna, S. & M. D. Paserman (2005) “Job Search and Impatience,” *Journal of Labor Economics*, 23, 527–588.
- DiNardo, J., J. Matsudaira, J. McCrary & L. Sanbonmatsu (2021) “A Practical Proactive Proposal for Dealing with Attrition: Alternative Approaches and an Empirical Example,” *Journal of Labor Economics*, 39, S507–S541.
- Donovan, K., J. Lu, T. Schoellman et al. (2018) “Labor Market Flows and Development,” in *2018 Meeting Papers*, 976, Society for Economic Dynamics.
- Duflo, E., M. Kremer & J. Robinson (2011) “Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya,” *The American Economic Review*, 101, 2350–2390.
- Ericson, K. M. & D. Laibson (2019) “Intertemporal Choice,” in Bernheim, B. D., S. DellaVigna & D. Laibson eds. *Handbook of Behavioral Economics - Foundations and Applications* 2, 2 of *Handbook of Behavioral Economics: Applications and Foundations* 1, 1–67: North-Holland.

- Faberman, J. & M. Kudlyak (2019) “The Intensity of Job Search and Search Duration,” *American Economic Journal: Macroeconomics*, 3 (11), 327–357.
- (2006b) “Self-Confidence and Search,” IZA Discussion Paper 2525.
- Falk, A., D. Huffman & U. Sunde (2006a) “Do I Have What It Takes? Equilibrium Search with Type Uncertainty and Non-Participation,” Discussion paper 2531, IZA.
- Fernando, N., N. Singh & G. Tourek (2021) “Hiring Frictions in Urban Labor Markets: Experimental Evidence from India,” Working paper, Notre Dame.
- Field, E., R. Garlick & K. Vyborny (2024) “Women’s Mobility and Labor Supply: Experimental Evidence from Pakistan,” Working paper.
- Filges, T., G. Smedslund, A. D. Knudsen & A. K. Jørgensen (2015) “Active Labour Market Programme Participation for Unemployment Insurance Recipients: A Systematic Review,” *Campbell Systematic Reviews*, 11, 1–342.
- Franklin, S. (2017) “Location, Search Costs and Youth Unemployment: Experimental Evidence from Transport Subsidies,” *The Economic Journal*, 128 (614), 2353–2379.
- Fu, J., M. Sefton & R. Upward (2019) “Social Comparisons in Job Search,” *Journal of Economic Behavior & Organization*, 168, 338–361.
- Gabaix, X. (2019) “Behavioral Inattention,” in Bernheim, D., S. DellaVigna & D. Laibson eds. *Handbook of Behavioral Economics*, 261–343: Elsevier.
- Gee, L. (2019) “The More You Know: Information Effects on Job Application Rates in a Large Field Experiment,” *Management Science*, 65 (5), 2077–2094.
- Gentile, E., N. Kohli, N. Subramanian, Z. Tirmazee & K. Vyborny (2023) “A Leaky Pipeline? Decomposing the gender gap in job search in urban Pakistan,” Working paper.
- Gurtzgen, N., L. Benjamin, L. Pohlman & G. van den Berg (2020) “Does Online Search Improve the Match Quality of New Hires?”, Working paper.
- He, H., D. Neumark & Q. Weng (2021) ““I Still Haven’t Found What I’m Looking For”: Evidence of Directed Search from a Field Experiment,” Working Paper 28660, NBER.
- Hjort, J. & J. Poulsen (2019) “The Arrival of Fast Internet and Employment in Africa,” *American Economic Review*, 109, 1032–1079.
- Jones, S. & K. Sen (2022) “Labour Market Effects of Digital Matching Platforms: Experimental Evidence from Sub-Saharan Africa,” 07.
- Kelley, E. M., C. Ksoll & J. Magruder (2021) “How do Online Job Portals Affect Employment and Search Outcomes? Evidence from India,” Working paper.
- Kircher, P. (2022) “Job Search in the 21st Century,” *Journal of the European Economic Association*, 20 (6), 2317–2352.
- Kiss, A., R. Garlick, K. Orkin, L. Hensel [author order certified randomized] (2023) “Jobseekers’ Beliefs about Comparative Advantage and (Mis)Directed Search,” Working paper, Duke University.
- Kleibergen, F. & R. Paap (2006) “Generalized Reduced Rank Tests Using the Singular Value Decomposition,” *Journal of Econometrics*, 133 (1), 97–126.
- Köszegi, B., G. Loewenstein & T. Murooka (2022) “Fragile Self-Esteem,” *Review of Economic Studies*, 89, 2026–2060.
- Kuhn, P. & H. Mansour (2014) “Is Internet Job Search Still Ineffective?” *The Economic Journal*, 124 (581), 1213–1233.
- Kuhn, P. & K. Shen (2013) “Gender Discrimination in Job Ads: Evidence from China,” *The Quarterly Journal of Economics*, 128 (1), 287–336.

- Kuhn, P. & M. Skuterud (2004) “Internet Job Search and Unemployment Durations,” *American Economic Review*, 94 (1), 218–232.
- Le Barbanchon, T., L. Hensvik & R. Rathelot (2023) “How can AI improve search and matching? Evidence from 59 million personalized job recommendations,” Working Paper.
- Lechner, M. & J. Smith (2007) “What is the value added by caseworkers?” *Labour Economics*, 14 (2), 135–151.
- Lentz, R., J. Maibom & E. Moen (2022) “Competitive or Random Search?”, Working Paper.
- Lichter, A. & A. Schiprowski (2021) “Benefit duration, job search behavior and re-employment,” *Journal of Public Economics*, 193, 104–326.
- Marinescu, I. (2017a) “The general equilibrium impacts of unemployment insurance: Evidence from a large online job board,” *Journal of Public Economics*, 150, 14–29.
- (2017b) “Job Search Monitoring and Assistance for the Unemployed,” *IZA World of Labor*, 380.
- Marinescu, I. & D. Skandalis (2021) “Unemployment Insurance and Job Search Behavior,” *The Quarterly Journal of Economics*, 136, 887–931.
- Marinescu, I. & R. Wolthoff (2020) “Opening the Black Box of the Matching Function: The Power of Words,” *Journal of Labor Economics*, 38 (2), 535–568.
- Matsuda, N., T. Ahmed & S. Nomura (2019) “Labor Market Analysis Using Big Data: The Case of a Pakistani Online Job Portal,” Policy Research Working Paper 9063, World Bank.
- McGee, A. (2015) “How the Perception of Control Influences Unemployed Job Search,” *Industrial and Labor Relations Review*, 68 (1), 184–211.
- McGee, A. & P. McGee (2016) “Search, effort, and locus of control,” *Journal of Economic Behavior and Organization*, 126, 89–101.
- MICS (2018) “Multiple Indicator Cluster Survey Punjab,” Bureau of Statistics Punjab, Planning & Development Board, Government of the Punjab.
- Mueller, A. I. & J. Spinnewijn (2022) “Expectations Data, Labor Market and Job Search,” in Bachmann, R., G. Topa & W. van der Klaauw eds. *Handbook of Economic Expectations: ScienceDirect*.
- Neumark, D. (2018) “Experimental Research on Labor Market Discrimination,” *Journal of Economic Literature*, 56 (3), 799–866.
- Paserman, M. D. (2008) “Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation,” *The Economic Journal*, 118, 1418–1452.
- Pissarides, C. A. (2000) *Equilibrium Unemployment Theory*: MIT Press, 2nd edition.
- Poverty Action Lab (2022) “Reducing Search Barriers for Job Seekers,” 01.
- Sanders, M., G. Briscese, R. Gallagher, A. Gyani, S. Hanes, E. Kirkman & O. Service (2019) “Behavioural insight and the labour market: evidence from a pilot study and a large stepped-wedge controlled trial,” *Journal of Public Policy*, 41, 42–65.
- Schiprowski, A. (2020) “The Role of Caseworkers in Unemployment Insurance: Evidence from Unplanned Absences,” *Journal of Labor Economics*, 38 (4), 1189–1225.
- Shimer, R. (2010) *Labor Markets and Business Cycles*: Princeton University Press.
- Spinnewijn, J. (2015) “Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs,” *Journal of the European Economic Association*, 13 (1), 130–167.
- Subramanian, N. (2021) “Workplace Attributes and Women’s Labor Supply Decisions: Evidence from a Randomized Experiment,” Working paper.
- Wheeler, L., R. Garlick, E. Johnson, P. Shaw & M. Gargano (2022) “LinkedIn(to) Job Oppor-

- tunities: Experimental Evidence from Job Readiness Training,” *American Economic Journal: Applied Economics*, 14, 101–125.
- Wright, R., P. Kircher, B. Julien & V. Guerrieri (2021) “Directed Search and Competitive Search Equilibrium: A Guided Tour,” *Journal of Economic Literature*, 59 (1), 90–148.
- Wu, L. (2021) “Partially Directed Search in the Labor Market,” Working Paper.
- Zizzamia, R. (2023) “Ignorance is bliss? Rejection and discouragement in on-the-job search,” Working Paper.

Table 1: Jobseeker Summary Statistics, Selection into Applications, and Balance Tests

	Control group mean & std. dev			Selection into applying in control group	Balance checks
	All	Never apply	Ever apply	Mean for ever apply – never apply [p-value]	Mean for treated – for control [p-value]
	(1)	(2)	(3)	(4)	(5)
Employed and searching	0.200 (0.400)	0.184 (0.388)	0.292 (0.455)	0.108 [0.000]	0.034 [0.228]
Employed and not searching	0.141 (0.348)	0.148 (0.355)	0.097 (0.296)	-0.051 [0.000]	-0.028 [0.256]
Searching and not employed	0.345 (0.475)	0.338 (0.473)	0.386 (0.487)	0.048 [0.033]	0.024 [0.344]
Not searching and not employed	0.314 (0.464)	0.330 (0.470)	0.225 (0.418)	-0.105 [0.000]	-0.030 [0.307]
Search method: network	0.397 (0.489)	0.379 (0.485)	0.506 (0.500)	0.127 [0.000]	0.032 [0.476]
Search method: formal application	0.154 (0.361)	0.150 (0.357)	0.176 (0.381)	0.026 [0.147]	0.028 [0.651]
Search method: asked at establishments	0.225 (0.417)	0.211 (0.408)	0.305 (0.461)	0.094 [0.000]	0.032 [0.728]
Years of work experience	7.85 (8.88)	7.89 (8.97)	7.62 (8.25)	-0.27 [0.463]	-0.22 [0.568]
Education: none	0.146 (0.353)	0.154 (0.361)	0.083 (0.276)	-0.071 [0.000]	-0.012 [0.294]
Education: primary or some secondary	0.457 (0.498)	0.470 (0.499)	0.361 (0.481)	-0.109 [0.000]	-0.023 [0.871]
Education: complete secondary	0.148 (0.355)	0.143 (0.350)	0.180 (0.384)	0.037 [0.027]	0.002 [0.673]
Education: university degree	0.250 (0.433)	0.232 (0.422)	0.376 (0.485)	0.144 [0.000]	0.033 [0.335]
CV: excellent score	0.093 (0.291)	0.092 (0.289)	0.098 (0.298)	0.006 [0.812]	0.084 [0.868]
CV: good score	0.330 (0.471)	0.342 (0.475)	0.299 (0.459)	-0.043 [0.281]	0.032 [0.970]
CV: average or lower score	0.576 (0.495)	0.566 (0.496)	0.603 (0.491)	0.037 [0.383]	-0.116 [0.872]
Female	0.303 (0.460)	0.307 (0.461)	0.271 (0.445)	-0.036 [0.063]	0.022 [0.329]
Age	30.7 (9.7)	31.0 (9.8)	28.7 (9.1)	-2.3 [0.000]	-0.5 [0.307]
# matches sent by platform	113 (121)	107 (120)	154 (119)	47 [0.000]	-
# applications on platform	0.226 (0.863)	0.000 (0.000)	1.83 (1.76)	1.83 [0.000]	-
# interviews through platform	0.014 (0.131)	0.000 (0.000)	0.116 (0.356)	0.116 [0.000]	-

Notes: This table shows summary statistics for jobseekers' baseline characteristics and, in the last three rows, platform use characteristics. This table uses one observation per jobseeker. Column (1) shows the mean and standard deviation for the control group. Column (2) shows the mean and standard deviation for the control group sample of jobseekers who never applied to any job on the platform. Column (3) shows the mean and standard deviation for the control group sample of jobseekers who apply to at least one job on the platform. Column (4) shows the difference between the mean for the control group sample of jobseekers who apply to at least one job and those who never applied, along with the p-value for testing if this difference is zero. This shows how jobseekers who apply to jobs on the platform differ from jobseekers who do not apply to jobs on the platform. Column (5) provides balance tests by showing the difference between the mean for the treated sample and the mean for the control group sample, along with the p-value for testing if this difference is zero. This checks if the treated and control respondents have the same baseline characteristics on average. P-values are generated from regressions that use heteroskedasticity-robust standard errors and include fixed effects for the strata within which treatment was randomized (see footnote 14). We leave column (5) blank for the final three rows because applications and interviews are post-treatment outcomes and the number of matches can be influenced by post-treatment actions, although we show in Section 3.1 that this influence is irrelevant for our main results.

Table 2: Vacancy- and Match-level Summary Statistics and Selection into Applications

	(1)	(2)
	Mean T=0 (Std dev. T=0)	Selection into application Mean T=0, A=1 – Mean T=0 [p-value]
Salary	14,420 (9,152)	6,577 [0.000]
Firm # employees	26.6 (135)	61.7 [0.000]
Firm # vacancies in last year	5.50 (12.2)	6.79 [0.000]
Exact education match vacancy requires high ed	0.184 (0.387)	-0.016 [0.542]
Exact experience match vacancy requires experience	0.126 (0.331)	0.050 [0.016]
Gender preference aligned	0.700 (0.458)	-0.191 [0.000]
Short commute	0.519 (0.500)	0.021 [0.329]
V_{vm} index: proxies of value of vacancy to jobseeker	0.011 (0.896)	0.255 [0.000]
Applied	0.002 (0.045)	0.998
Interviewed	0.000 (0.011)	0.064 [0.000]


Notes: This table shows summary statistics for vacancy- and match-level characteristics. Column (1) shows the mean and standard deviation for the control group sample. Column (2) shows the difference between the mean for the control group sample of matches that resulted in applications and the mean of the full control group sample of matches, along with the p-value for testing if this difference is zero. This shows how matches that lead to applications differ from other matches. P-values are generated from regressions that control for stratification block fixed effects and use heteroskedasticity-robust standard errors clustered by jobseeker. The p-value for ‘Applied’ in column (2) is omitted because the standard error is zero by definition for the mean application rate conditional on application. Salary is in Pakistani Rupees per month. 1 Rupee $\approx$ 0.03 USD in purchasing power parity terms during the study period. Exact education match is an indicator for an exact match between the employer’s preferred field of educational specialization and the jobseeker’s field. Exact experience match is an indicator for a match in which the jobseeker has experience in the same occupation as the vacancy. These two variables are only defined for vacancies that require respectively more than basic education and some experience. These two variables use employers’ *preferred* education and experience, rather than the *required* education and experience used in the matching algorithm. The proxy index V_{vm} is an inverse covariance-weighted average (following [Anderson 2008](#)) constructed using vacancy-level characteristics log salary and indicators for offering any non-salary benefits, below-median working hours, and allowing flexible hours as well as indicators for the match-level characteristics of vacancy salary exceeding the jobseeker’s expected salary, below-median commuting distance, the jobseeker’s educational specialization exactly matching the vacancy’s preference, and the jobseeker’s work experience exactly matching the vacancy’s preference. Anderson-style indices, by construction, have zero means and hence some negative values.

Table 3: Treatment Effects on Job Search & Search Returns

	(1) Apply	(2) Interview	(3) Int. $\times V_{vm}$	(4) Interview	(5) Int. $\times V_{vm}$
Phone call treatment	0.01322 (0.00075)	0.00077 (0.00009)	0.00283 (0.00036)		
Apply				0.05852 (0.00516)	0.21372 (0.02169)
# matches	1,116,952	1,116,952	1,116,952	1,116,952	1,116,952
# jobseekers	9831	9831	9831	9831	9831
Mean outcome T = 0	0.00185	0.00012	0.00045	0.00012	0.00045
Mean outcome T = 0, Apply = 1				0.06381	0.24487
p: IV effect = mean T = 0, Apply = 1				0.573	0.413
IV strength test: F-stat				312.8	312.8
IV strength test: p-value				0.00000	0.00000

Notes: Column 1 shows the coefficient from regressing an indicator for job application on treatment assignment. Column 2 shows the coefficient from regressing an indicator for interview invitation on treatment assignment. Column 3 shows the coefficient from regressing an indicator for interview invitation weighted by a proxy index for the value of the vacancy to the jobseeker, V_{vm} , on treatment assignment. Column 4 shows the coefficient from regressing an indicator for interview invitation on job application, instrumented by treatment assignment. Column 5 shows the coefficient from regressing an indicator for interview invitation weighted by the proxy index V_{vm} on job application, instrumented by treatment assignment. The proxy index V_{vm} is an inverse covariance-weighted average (following [Anderson 2008](#)) constructed using vacancy-level characteristics log salary and indicators for offering any non-salary benefits, below-median working hours, and allowing flexible hours as well as indicators for the match-level characteristics of vacancy salary exceeding the jobseeker's expected salary, below-median commuting distance, the jobseeker's educational specialization exactly matching the vacancy's preference, and the jobseeker's work experience exactly matching the vacancy's preference. Anderson-style indices, by construction, have zero means and hence some negative values. But multiplying the interview invitation indicator by a negative value would not produce sensible results. Hence we recenter the index so it has strictly positive values. All regressions use one observation per jobseeker $\times$ vacancy match, include stratification block fixed effects, and use use heteroskedasticity-robust standard errors clustered by jobseeker, which are shown in parentheses. The p-value is for a test of equality between the IV treatment effect and the mean interview rate for control group applications. The first-stage F-statistic and p-value are for the test of weak identification from [Kleibergen & Paap \(2006\)](#).

Figure 1: Application Decisions for Treated and Control Jobseekers Facing High versus Low Costs



Notes: This figure shows the application decisions for jobseekers facing low application costs at the time they receive matches (top panel) and jobseekers facing high application costs at the time they receive matches (bottom panel). The light-shaded sections show the matches to which control group jobseekers apply. The dark-shaded sections show the additional matches to which treatment group jobseekers apply. For simplicity, we show only the right tail of the density of PV . We formally derive values for PV_{H0} , PV_{H1} , PV_{L0} , and PV_{L1} in Appendix C.2.