

Remote work and job applicant diversity: Evidence from technology startups

David H. Hsu & Prasanna Tambe*
Wharton School, University of Pennsylvania

December 2, 2023

Management Science, *Forthcoming*

Abstract

A significant element of managerial post-COVID job design regards remote work. In an era of renewed recognition of diversity, equity and inclusion, employers may wonder how diverse (gender and race) and experienced job applicants respond to remote job listings, especially for high-skilled technical and managerial positions. Prior work has shown that while remote work allows employee flexibility, it may limit career promotion prospects, so the net effect of designating a job as remote-eligible is not clear from an applicant interest standpoint, particularly when recruiting females and underrepresented minorities (URM). We analyze job applicant data from a leading startup job platform that spans long windows before and after the COVID-19 pandemic-induced shutdowns of March 2020. To address the empirical challenge that remote job designation may be co-determined with unobserved job and employer characteristics, we leverage a matching approach (and an alternative method which leverages the sudden shutdowns) to estimate how applicant characteristics differ for otherwise similar remote and onsite job postings. We find that offering remote work attracts more experienced and diverse (especially URM) job applicants, with larger effects in less diverse geographic areas. A discrete change in job posting to *remote* status (holding all else constant) is associated with an approximately 15% increase in applicants who are female, 33% increase in applicants with URM status, and 17% increase in applicant experience. Using the application data, we estimate an average estimated compensating wage differential for remote work that is about 7% of posted salaries in this labor market.

Keywords: remote work; startup technical labor markets; racial and gender diversity; COVID-19.

*Correspondence to: dhsu@wharton.upenn.edu & tambe@wharton.upenn.edu. We are grateful for feedback from David Autor, Matthew Bidwell, Peter Cappelli, Raj Choudhury, Ananya Sen, Olav Sorenson (our editor), Bledi Taska, Vilma Todri, and seminar audiences at the Asian Innovation and Entrepreneurship Association (AIEA), Boston College, Univ. California-Berkeley, Univ. Colorado-Boulder, George Mason Univ., INSEAD, Univ. Southern California, Tulane, the Wharton People & Organizations Conference, the Workshop for Information Systems and Economics (WISE), the INFORMS Conference on Information Systems and Technology (CIST), the CMU Block Center Future of Work Conference, and the MIT IDE seminar. We particularly thank Kevin Laws at AngelList for data access and Remi Gabillet from Wellfound (formerly AngelList Talent) for providing data assistance. We acknowledge research support from the Mack Center for Innovation Management at the University of Pennsylvania. Errors and omissions are ours.

1 Introduction

Remote work, also known as working-from-home (WFH), is not a new phenomenon. The 2010 US Census reported that approximately 10% of Americans in the work force conducted remote work at least one day a week, and about 5% exclusively practiced WFH.¹ A much wider share of the world experienced remote work from 2020-22, during the COVID-19 pandemic due to government-mandated shutdowns, with global remote-eligible job listings in that period rapidly scaling (Brynjolfsson et al. [2020]; Hansen et al. [2022]), yet demand for remote work still exceeding supply world-wide (Aksoy et al. [2022]). These behavioral changes by both employers and workers, aided by information technology platforms, give rise to a key managerial decision in the post-pandemic world: what WFH policies should firms adopt when they are no longer forced to do so?

The prior academic literature in this domain concentrates on an important question informing the managerial choice of offering remote work, that of worker productivity for those randomized to a WFH condition within an organization. These field experiments mostly (though not exclusively, as we detail in the literature review below) find a positive individual productivity effect caused by the remote work condition within the subject organizations.

Another important aspect of the decision to offer remote-eligible work, however, is how the applicant pool responds to such work (Hensvik et al. [2021]). This perspective has received much less attention in the literature, yet is important particularly in tight labor markets, such as the one we study (technology ventures). In addition, the appeal of job flexibility through remote work may manifest differently depending on worker status and circumstance. While substantial contemporaneous popular press discussion points to a cost- or employee-based rationale for WFH policies, such as economizing on office space or employee productivity or retention-based motivations, we investigate a potential *unintended* consequence/benefit to WFH policies, namely the ability to attract gender- and race-diverse applicants. After all, if one barrier to job applicant diversity is flexibility, allowing remote work may be an important job design decision. We therefore pose the following research question related to the labor market: do employers attract more diverse (gender and race) and experienced talent when they designate a job as remote-eligible?

To study this question, we analyze a new data set, from AngelList Talent (rebranded as Wellfound since November 2022), that covers activity from both sides of the startup labor market (employer job postings and applications to those postings), collected before and after the COVID-19 induced shutdowns of March 2020. We believe this data set is unique in that it captures both supply and demand behavior in the labor market at considerable

¹<https://www.census.gov/data/tables/2010/demo/commuting/p70-132.html>.

scale and granularity (prior studies using detailed market data from the full-time market typically analyze one side of the market or the other). Furthermore, because the AngelList Talent platform caters to the growth-oriented, early-stage startup labor market, our sample mitigates undesirable heterogeneity. This labor market also complements the type of work and jobs featured in prior field studies on worker productivity resulting from remote work, which largely (though not exclusively) examine tasks with quite objective performance criteria such as data entry and call center performance.

Our first analysis considers the jobs listed as remote-eligible and the organizations which list them, prior to and extending after the COVID-19 shutdowns. We pay particular attention to portraying descriptive patterns because there has been limited descriptive evidence on these patterns documented across a large number of organizations and jobs in the prior literature. We incorporate the detailed nature of our job listing data, and characterize the organizations and jobs associated with the likelihood a given job is listed as remote-eligible.

This sets up the main analysis we conduct to address the research question of whether organizations attract more experienced and diverse talent if they designate a job as WFH. Because we know (and show in our analysis) that remote job listings are not randomly offered, analyzing job market attraction to remote job listings should ideally endogenize the job design decision to aid inference. Our main empirical strategy is a matching approach, in which we match job postings by week of posting (to account for sharp differences in remote eligibility over time, especially over the course of the COVID-19 pandemic), job title (to recognize that certain types of jobs may be more or less amenable to remote work), and geography (reflecting, among other things, competitive work practices including flexible job design).² An alternative empirical strategy (which we summarize in the text but present in more detail in the appendix) is to identify jobs before the shutdowns which were largely onsite, but because of the shutdowns were forced to be remote conditional on hiring aspirations. We analyze applicant behavior for these jobs to triangulate our main empirical strategy. These approaches allow us to mitigate a common confound in the job design literature, namely that job design and firm characteristics can be co-determined with the decision to offer remote work in ways which are unmeasured and/or unobserved. Within the limitations of these estimation approaches, we find that offering remote work affords organizations more experienced applicants as well as applicants drawn from more diverse gender and (especially) race backgrounds. A final analysis examines the heterogeneity of these patterns across labor markets which themselves differ by diversity and quantifies the intensity of remote preferences for applicants via a compensating wage differential analysis.

²We present result robustness to matching approaches using a host of additional firm-level characteristics in the Appendix.

2 Literature

Two streams of prior work relate to our research question. While both address the shifting organization of work, the first domain takes the employer perspective about considerations when offering remote work (and emphasizes productivity effects). The second domain takes the employee perspective and examines drivers of their preference for/against applying for remote-eligible jobs. We discuss themes within each literature in turn, especially as they relate to our empirical work, which puts together both sides of the market.

2.1 Employers and the decision to offer remote work

The literature on employers’ decision to offer remote work centers on employee productivity. Both [Bloom et al. \[2015\]](#), in the setting of a call center operation in China, as well as [Choudhury et al. \[2021\]](#), in the setting of US Patent and Trademark Office (USPTO) examiners study remote work. The call center study found a 13% increase in employee productivity, mainly due to reduced commute-, break- and sick-time, as well as a quieter work environment (and were 50% less likely to quit as well, though their likelihood of subsequent promotion declined). The patent examiner study found a 4.4% boost in output without need for re-work when allowed to work from anywhere.

With the caveat that the field experimental evidence is drawn from only a handful of organizations, one puzzle which arises is: if offering remote work increases employee productivity, why do all employers not offer it? A possible explanation involves career concerns in which employees may be wary of remote work’s impact on their promotion possibilities, even if they prefer the flexibility. In a call center work setting before and after the COVID-19 shutdown, [Emanuel and Harrington \[2020\]](#) find more productive workers do not want to pool with (latently) less productive ones. They find that the reason is because of the lower promotion rates, consistent with [Bloom et al. \[2015\]](#), which also uses a call center empirical setting. More generally, [Cullen and Perez-Truglia \[2023\]](#) presents quasi-experimental evidence of social interactions, facilitated by face-to-face interactions (employee-manager smoking breaks), leading to career progression.

Another explanation to the puzzle challenges the premise that remote work increases employee productivity in the first place. Using sociometric badges, [Wu et al. \[2008\]](#) find that face-to-face interactions have a positive productivity effect for many tasks. Similarly, using smartphone geolocation data, [Atkin et al. \[2022\]](#) find that face-to-face interactions significantly boost knowledge spillovers (measured via patent citations) within Silicon Valley. In a field experiment of data-entry work in India where performance and error rates can be precisely measured, [Atkin et al. \[2020\]](#) find a positive on-site treatment effect, and evidence that the most capable employees sort into formal office settings. These findings are consis-

tent with recent remarks by JP Morgan Chase’s CEO Jamie Dimon that remote work does not work well “for those who want to hustle” (Benoit [2021]). The negative WFH productivity effect could also arise from behavioral origins, in that workers may not wish to shirk when working remotely, but because of self-control issues (Kaur et al. [2015]), managerial oversight in an on-site capacity can help impose worker self-discipline.

While there is considerable focus in the literature on worker productivity effects of remote work, less discussed is the ability of firms to manage a remote workforce. This relates to the degree of managerial intensity and skills required to manage in the virtual environment. Intertwined with this are job characteristics and the nature of the work demanded; if work is relatively modular, managerial intensity may be lower. Dingel and Neiman [2020], in an analysis of US occupations, find that about 37% of jobs can be done at home and that jobs that fall into this category tend to be higher paying. Notably, the authors find that 100% of “computer and mathematical occupations,” and 87% of “management occupations” which together characterize the majority of jobs in our empirical sample, can be performed remotely, though the productivity of WFH jobs depends on firms’ digital technology adoption (Bai et al. [2021]) among other factors. Unfortunately these same factors are difficult to observe and measure at scale, which presents a host of empirical challenges we aim to overcome (which we will discuss at more length in the data section).

2.2 Employees and the decision to engage in remote work

Here, we discuss employee perspectives in shaping preferences for remote work. Within the context of STEM-related work, women and URM are underrepresented in these jobs relative to their numbers in the overall labor market according to the US Bureau of Labor Statistics.³ The literature has highlighted a multitude of factors including background preparation for such jobs, access to typical hiring channels, disparities in retention rates (especially given extra-work obligations), discrimination in the job market, and self-selection into job applications.⁴ In this section, we discuss three channels through which underrepresented groups (women and URM) in particular may assess remote work: time flexibility, locational flexibility, and face-to-face workplace interactions. While the first two of these emphasize beneficial attributes of remote work, the third category highlights both benefits and costs.

³<https://www.bls.gov/opub/reports/womens-earnings/2020/pdf/home.pdf>;
<https://www.bls.gov/opub/reports/race-and-ethnicity/2020/pdf/home.pdf>

⁴Of course, disparate labor market outcomes can arise from both applicant as well as employer behavior. Using applicant tracking system data (information on who applies, gets a callback, an interview, and an offer), Parasurama et al. [2020] find that underrepresentation of women in information technology (IT) fields tends to be driven by choices of the worker, whereas underrepresentation of under-represented minorities (URM) tend to arise from employer choices.

2.2.1 Time flexibility

While time flexibility is generally valued by the workforce at large⁵, one important theme in the literature is that women, especially those with young children and families (e.g., [Mas and Pallais \[2017\]](#); [Atkin et al. \[2020\]](#); [Barbulescu and Bidwell \[2013\]](#)) disproportionately value time flexibility in their jobs, including the ability to conduct remote work, in part due to societal expectations of family roles. In [Barbulescu and Bidwell \[2013\]](#)’s archival study of female job applicants to managerial positions, expectations of future job demands factor into their choice of whether to apply in the first place for a job. This pattern is not true for male applicants, and therefore the difference may reflect anticipated obligations in professional and family life (and the balance between the two). As the empirical setting is the first jobs of business school graduate candidates, whose average age reflects typical child-rearing years, the job flexibility explanation is especially salient.

Not only do women’s expectations of family duties influence sorting to jobs, when unexpected circumstances arise with regard to child care shortfalls after employment, women tend to disproportionately bear the cost. During the COVID-19 pandemic, in a survey of scientists, [Myers et al. \[2020\]](#) found that female scientists and those with child dependents experienced a substantial decline in time devoted to research. Given these disparities, it may be the case that the flexibility accompanying remote work may be especially appealing to women with young children at home. The preference for flexibility may also be manifested in the lower willingness of women to commute as far to their workplaces as compared to men ([Le Barbanchon et al. \[2021\]](#)).⁶

2.2.2 Locational flexibility

Another context in which women, even those not caring for young children, may value flexible job arrangements such as remote work is the case of dual-earner couples seeking to be spatially co-located (the “two-body problem”). Especially for such couples who are in specialized skill professions as would be the case in the empirical context we examine, the job market may not be equally “thick” for each individual. Indeed, [Benson \[2014\]](#) notes that women tend to be segregated into geographically dispersed occupations (especially so for more specialized science- and engineering-based occupations), and so such geographic

⁵For example, in an experimental setting, [Mas and Pallais \[2017\]](#) estimate that workers on average accept compensating wage differentials (20% of wages to avoid a schedule set by an employer on short notice and 8% for the option to WFH) for work flexibility, though the average is skewed by outlier preferences. Note that these estimates are both for call center workers (a distinct labor market relative to the technology labor market we examine), and are estimated based on on field experimental evidence of one organization (whereas our estimates stem from archival from a wide gamut of organizations).

⁶While average female preferences for a shorter commute relative to men may reflect either individual preferences or household decision constraints (or both), there is also the possibility that decisions regarding value of time and work flexibility can be individual- and even time-of-day-specific [[Chen et al., 2019](#)].

constraints may relate to household relocation decisions. This constraint may be loosened under remote work. In addition, [Sorenson and Dahl \[2016\]](#) note that a contributing factor to the gender wage gap is due to couples often giving greater weight to men’s careers in their location choices. As a result, the gender wage gap may reflect co-location decisions in labor markets which may not be as well-matched for the female as compared to the male partner. With the possibility of remote work, especially if such work is de-stigmatized as may be the case in the post-pandemic environment, dual-earner couples may have expanded possibilities for organizing their work and household lives.

URM candidates may also value flexibility afforded by remote-eligible job positions. Consider that just five cities (Boston, San Diego, San Francisco, Seattle, and San Jose) accounted for approximately 90% of the growth of technology-sector jobs between 2005 through 2017, according to [Atkinson et al. \[2019\]](#). The affordability (including real estate and cost of living) of these locales has become less accessible to all households, especially for URM demographic groups who have not had the same historic opportunities for wealth accumulation. This can contribute to URM candidates selecting living locations distant from the best employment opportunities, whether that is across- or within-cities. Especially in the within-city context, the seminal work of [Kain \[1968\]](#) suggests the upshot is worse economic outcomes for URM stemming from worse information about the best employment opportunities, a higher likelihood of racial discrimination, and higher direct commuting costs of reaching the desired employment establishment. Collectively, these deleterious effects have become known as the spatial mismatch theory. The general consequence in our empirical setting is a potential mismatch between technology job opportunity locales and geo-spatial residence patterns for URM. Remote work holds the potential for mitigating such mismatches and therefore may elevate job applicant interest from URM candidates.⁷

2.2.3 Limiting face-to-face interactions

Unlike time and locational flexibility which are generally considered beneficial, especially to underrepresented workers, face-to-face interactions may be a double-edged sword. On the one hand, remote work could be valued by diverse job candidates due to *limiting* face-to-face interactions. URM may face higher costs of in-person work interactions relative to others. Racial “microaggressions” have been defined as “brief and commonplace verbal, behavioral, and environmental indignities, whether intentional or unintentional, that communicate hostile, derogatory, or negative racial slights and insults toward people of color” [Sue et al. \[2007\]](#). Studies (e.g., [Harwood et al. \[2012\]](#)) have found that in-person informal work interactions can sometimes result in microaggressions. In the context of a university

⁷For example, the firm Meta in their publicly-available 2022 Diversity Report stated that job candidates accepting their remote job offers were more often from under-represented groups.

environment closing as a result of the COVID-19 pandemic, [Cho and Brassfield \[2022\]](#) in their field work found that remote work brought a sense of relief to some URM both because of the potential for microaggressions on premises, but also because of the threatening conditions in the physical travel to the facility. In a parallel effect for women, it may be the case that one reason remote work may be preferred by some women is because some organizational contexts may be more permissive of sexual misconduct [[Holland et al., 2020](#)].

On the other hand, the literature has also discussed detrimental career progression effects of limited face-to-face workplace interactions (as would be the case with remote work). While workers may prefer the flexibility of remote work in the short term, the literature has also documented that such a work arrangement may be detrimental to the employee from the standpoint of promotion prospects. For example, in the [Bloom et al. \[2015\]](#) field experiment, eligible workers who opted for remote work were about 8.5% less likely to be promoted over the next 20 months. Possible reasons for this include less “soft” information about an employee which might otherwise be garnered through in-person informal social interactions (including the possibility of being “out of sight, out of mind” for promotion consideration). In addition, remote employees might face higher costs of on-the-job training (including potentially fewer opportunities to develop interpersonal skills which might be especially important in being promoted to managerial positions).

Concerns regarding career progression may be particularly important for those with URM status, as such individuals may be less equipped to utilize the external job market if their internal career progression in their current employer is hampered. Since being referred to a job is an important method of hiring, and because racial minorities are both less entrenched and represented in organizations, there is less opportunity for co-ethnic referrals as a channel for such racial minorities ([Fernandez and Fernandez-Mateo \[2006\]](#)). More generally, URM individuals tend to react to their job market headwinds *not* by self-selecting into segments of the job market to avoid discrimination, but by engaging in wider job searches ([Pager and Pedulla \[2015\]](#)). In recognition of the difficulty of engaging in the external job market, it is not clear how URM might value remote work on net, as they may encounter a range of frictions in the external job market should their internal career progression stall.

For women, the potential career progression costs of remote work may instead be more related to lower quality of feedback on their work product. In the empirical context of female software engineers in a Fortune 500 firm, [Emanuel et al. \[2022\]](#) find that onsite work, particularly for females, is associated with greater feedback on their work product and lower incidence of separation from the firm.⁸

⁸Under the assumption that on-the-job social interactions with managers may be more difficult under remote work relative to on-site, face-to-face interactions, the study by [Cullen and Perez-Truglia \[2023\]](#) documenting a gender pay gap resulting from lower rates of promotion is salient. Using exogenously-given managerial rotation in a large commercial bank in Southeast Asia, the authors find that employee-management

The net effect of how workers of various types value remote work is therefore ambiguous, with flexibility (along the time and locational dimensions discussed) being potentially traded off against career progression. We therefore consider the question of whether the introduction of remote-eligible jobs will increase the number of gender and ethnically diverse applicants to be largely an empirical one.

3 Data

There are two empirical difficulties in examining how making jobs remote-eligible is likely to impact the composition of the applicant pool. First, there are a number of unobserved and unmeasured factors on both sides of the marketplace which make a typical observational study difficult. Second, the methodology of many of the recent studies in this literature relies on field experiments, which enhance causal interpretation, but have the drawback of being unable to generalize across field sites (organizations) and are not suited to studying labor *markets*. In this section, we describe the data, variables, and empirical strategy we use to address these dual challenges.

3.1 Overview

We analyze data from AngelList Talent, a leading online platform for startup labor market activity for entrepreneurial ventures. This job marketplace is part of a larger AngelList platform catering to the startup ecosystem (AngelList Venture, for example, is a financial capital marketplace). Although not exclusive to technical occupations, most of the activity on AngelList Talent is for positions in technical or technical-adjacent positions, such as machine learning engineers, data scientists, or product managers, together with associated management and sales/marketing positions.⁹ Importantly for our study, prospective employers on the platform are asked to explicitly designate whether the job being posted is “remote”, which is a relatively unique feature among labor market platforms.

AngelList Talent has since its inception in 2011 attracted some 62,000 companies to post over 215,000 jobs. The site has had approximately 3.6M unique job candidates upload profiles to their platform, and has connected about 1M applicants to jobs. We do have further information about the applicants, including their applications to job postings from within the platform, as well as limited information through the platform about activity further down the hiring funnel (such as mutual employer interest). However, we do not

socializing (in this case via smoking breaks) importantly predicts promotion and wage growth. Since female smokers are relatively rare at the focal firm, the authors mainly study male-male employer-manager interactions.

⁹The COVID-19 induced recession of 2020 affected workers unequally, with demand for technology workers, such as those working for venture capital-backed startups, remaining robust (e.g., [Gompers et al. \[2022\]](#)), while the economic situation for all small businesses on average has been characterized as much more precarious and financially fragile (e.g., [Bartik et al. \[2020\]](#)).

have information about hiring offers or applicant mobility events (such activities take place outside of the platform).

While the AngelList Talent data span historical job postings since the platform’s inception since 2011, we concentrate our attention on job postings and applicant behavior, aggregated to the weekly level, for the two years before and after the COVID-19 shutdowns (which we take to be March 15, 2020, the first workday after the date on which the US declared a federal emergency). From this platform, we have job listing data on all of the job postings employers made to the platform including a description of the firm, open position titles, remote work designation, and compensation details (wages and equity). From the applicant data, we have information on demographics, some information on prior experience, and job listings that they applied to (including details of timing). To engage on the platform, job candidates fill in their profile information (such as race and gender, job preferences, and resume data, including educational history, skills, and prior employers).

The gender and race field we analyze are self-reported by website participants, who can choose whether or not to disclose these fields. In this data set, about 42% of applications are from applicants who report their race, and about 70% are from applicants who choose to report their gender. The choice of whether to disclose race or gender may be correlated with remote work preferences. In Appendix D, we evaluate the effects of this missing demographic data on our analysis. This appendix reports detailed distributions of applicants by gender and race and then compares them with the distributions produced when using applicant names to infer race and gender for a large sub-sample of applicants.¹⁰ We also report the results of an analysis that recreates our main analysis on a narrower time window during which we have both self-reported and name-inferred race and gender information, and we find that the point estimates are similar when using either source of demographic data.

On this labor platform, applicants are able to search for positions (and can filter their searches by role, compensation, location, skills required, startup size, etc.) and can apply for a position. Examining both the labor demand and supply sides of this market together allows us to: (1) characterize organizational and job characteristics associated with remote-eligible job listings, and (2) analyze how a shift in the characteristics of job listings (remote-eligible) induced by the pandemic is met by a change in applicant composition.

3.2 Data description and variables

In this section, we provide a high-level description of (a) job listings and (b) applicant characteristics. Before doing so, however, we characterize the types of organizations listing on

¹⁰We obtained applicant names for a large sub-set of these workers for this comparison. However, because we only have names for a subset of workers, we prioritize the self-reported race and gender information for our main analysis.

the AngelList Talent platform (not just in our sample window). In total, of the approximately 214,000 organizations posting jobs on the platform, about 63% of the organizations have an employee headcount of less than 501 people, with 19% of organizations on the platform belonging to the larger than 50 but less than 201 employee category. The remaining 37% of the overall sample of organizations are larger than 500 employees, with the largest segment (almost 22%) within that being the over 1,000 but less than 5,001 category.¹¹

AngelList Talent’s platform covers job listings from around the world (though for our sampling purposes, we exclude international locations in an effort to stem undesirable heterogeneity), with the top two locations being San Francisco and New York City which together comprise well over half of the total job listings on the platform since 2011. This may not be surprising given both the locus of venture capital backed startup activity and potential home bias in AngelList’s headquarters (which is located in San Francisco).

We examine two areas empirically, with the first area laying the groundwork for the main analysis tied to our research question (listed as the second area): (1) factors shaping the likelihood a job is listed as remote-eligible (before and after the COVID-related shutdowns in the US, March 15, 2020), and (2) how WFH job status influences job applications from those with gender and racial diversity, and from different experience levels.

3.2.1 Description of remote listings

In our job listings dataset, the top four job titles are: software engineer (14,750), senior software engineer (13,758), product manager (10,147), and sales development representative (7,168). Figure 1 shows the fraction of all posted full-time jobs designated as being remote, where each x-axis tick corresponds to a week. The illustration spans the period beginning four years prior to the pandemic closures and another two years after the closures. There is a gradual rise in the fraction of jobs listed as remote in the years leading up to the pandemic, with a sharp jump corresponding to the period in which companies were forced to go fully remote, in the latter part of 2020. We can also see that, after a brief dip in the fraction of remote jobs (corresponding to the period in which COVID rates had begun to subside but the Delta variant was just beginning to spread), the incidence of jobs listed as remote continued to rise. By the end of our sample, almost 80% of jobs were listed as remote.

Jobs listed as remote-eligible are heterogeneous by firm size, though. Figure 2 is a histogram of listings by firm size before and after the pandemic (as well as the differences displayed in the final panel), where the fraction of remote jobs in each category is shaded differently. The figure indicates that the distribution of listings posted on this platform

¹¹Organizations listing on the AngelList platform since its inception are typically relatively early stage, with the top two highest attained funding stages (at the time of listing) as Series A, with about 24% of the sample, and seed stage (16%).

shifts towards smaller firms after the pandemic. The shift in fraction of jobs listed as remote is particularly pronounced, with listings from smaller firms (less than 200 employees) comprising the most notable share of all remote listings.

Figure 3 is a heat map that characterizes – before and after the pandemic-induced shutdowns – the degree of remote work in listings for jobs in different occupations and salary brackets. Before the pandemic, we see that about 20% of jobs are listed as remote. The lowest incidence of remote work is in jobs in the highest wage bracket, over \$500,000, and jobs with wages less than \$50,000 had the highest incidence of remote work. Among these, it appears that management positions have consistently lower degrees of remote work, across all wage brackets. Among all listings, positions for management and marketing in the lower wage brackets had the lowest incidence of remote work eligibility.

The heat map changes significantly in the post-closure period. As expected, all cells are lighter, indicating a much higher incidence of remote work across all job categories and wage brackets. Management positions are listed with a higher incidence of remote eligibility than engineering or operations jobs across most wage categories. Again, jobs in the highest wage bracket are the least likely to be listed as remote, although in the lower wage brackets, there is less of a clear pattern in which occupation-wage combinations are most likely to support remote work.

We now embed factors associated with remote-eligible jobs in a multivariate framework. In Table 1, we present correlations between organizational and job characteristics and whether a job is listed as remote or onsite. The first column restricts the sample to listings from 2018-2019 which, of course, predates any pandemic-induced forced closures. During this period, we can see that smaller firms are more likely to list jobs as remote (firms with 10 or fewer employees constitute the reference category¹²), although this pattern reverses for the largest firms which are most likely to have the organizational infrastructure in place to support remote work. Consistent with the descriptive patterns, product and operations roles are less likely to be remote (‘designer’ is the reference job role category). Finally, salary and equity move in opposite directions – higher wage jobs are less likely to be remote. Higher equity positions are *more* likely to be remote, perhaps because younger, smaller firms which tend to offer remote jobs and tend not to have large physical footprints are also likely to grant more equity for early stage employees.

The second column shifts the sample to a post-closure period. In the later sample period, many of the coefficient estimates on organizational and job characteristics are no longer significant, as might be expected because due to the forced workplace closures, jobs being

¹²Based on anecdotes from the popular press and conversations with the data team at AngelList Talent, these very small companies have no physical footprint (they are least likely to have the financial means to spend on renting a dedicated office), so it is natural that they post remote jobs before they establish a physical business location.

listed as remote are less likely to be selected on characteristics.

3.2.2 Empirical design examining the consequences of remote listing

In our focal analysis, we examine how a host of job applicant characteristics relate to (and are induced by) the managerial decision to make jobs remote-eligible. In this analysis, our outcome variables center on applicant reactions to job postings, and span the following: (1) *applicant experience* is the number of years of work experience in the current job role among applicants to a job; (2) *number of female applicants* is the average number of applications sent by female applicants to job postings; and (3) *number of URM applicants* is the average number of applicants to a given job which are likely to belong to an underrepresented minority group. Table 2 contains the summary statistics for these variables.

The explanatory variables are *post-closure*, an indicator for after the COVID-induced shutdowns; and *remote* is a dummy variable for whether a job is listed as remote-eligible. An important control variable we include in the empirical specifications is *number of applications*, which is the count of the total number of applications sent to a job listing. Most of our specifications include fixed-effects for weeks since posting the listing and job type and Appendix B reports additional tests that add fixed-effects for primary city and/or employer.

To mitigate the issues around endogenous remote-eligible job listings, our empirical strategy matches job listings by detailed observable characteristics. In this coarsened exact matching (CEM) procedure, we balance the “treated” (remote) observations with the “control” observations on week of job listing, job title, salary and equity. This matching approach (details are contained in Appendix A and Appendix B, which also present results of a more extensive matching approach with firm characteristics, which when included, yields qualitatively similar estimates but drops the sample size by over 90% - and so we offer those results as robustness tests) helps improve inference attributed to the *remote* job effect on applicant characteristics.

4 Results

4.1 Talent attracted to remote job listings

We begin our analysis in this section with descriptive patterns of the applicants pre- and post-shutdowns. Figure 4 depicts the total applications sent to jobs in the two years before and two years after the March 2020 shutdowns, also dis-aggregated by applications to onsite and remote jobs. The pattern is dramatic. Figure 4 demonstrates that before the shutdowns, the fraction of applications sent to remote job listings as a share of all applications is essentially flat and stable. Applications sent to onsite jobs appear to be slowly declining as an overall proportion. After the shutdowns, the corresponding percentage for onsite

applications collapses, and the vast majority of applications are those that are sent to remote positions (which also reflects job postings increasingly listed as remote over time).¹³

In Figure 5, we plot applicants’ industry experience and diversity (number of female applicants and number of URM applicants) over time (again, the dashed vertical lines denote the temporal closure boundary). Throughout the figure, we difference (remote minus onsite) the fraction of applications which are from female applicants (top panel), URM candidates (middle panel), and by experience (bottom panel). While females throughout the time horizon in general seem to apply more to onsite job listings, the opposite is true for URM candidates and those with more experience. For females, the number of applicants to job listings appears to be slightly higher for on-site as compared to remote jobs before the closures. That pattern is somewhat amplified after the closures, with a higher fraction of female applicants to onsite jobs. For URM-status applicants, the trend is more pronounced. There is a higher proportion of URM applicants to remote job listings before the closure, but this difference clearly widens after the closures. The pre-post pattern on the difference in applicant experience appears consistent. As is true for all patterns shown in this table, no empirical strategy has yet been employed, and so should be interpreted as purely descriptive.

Our next step is to examine these results within a more complete OLS regression framework in which we do not yet tackle the issue of endogenous remote job listings. In Figure 6, we examine the correlates of applicant characteristics received to job listings, where the sample can include a window of time that includes the pandemic closures. The key specification we estimate is $y_i = \text{remote}_i + \log(\text{salary})_i + \log(\text{equity})_i + \log(\text{numapps})_i + \gamma_i + \epsilon_i$ where y_i is logged women applications, logged URM applications, or logged years experience of the applicant pool for each job posting i and γ_i is a vector of job listing fixed effects that includes job title and week of posting.

Each row in each panel corresponds to a separate analysis, and shows the point estimate and 95% confidence interval for the remote coefficient (together with the sample size). Standard errors are clustered by weeks since posting and employer. The top point estimate in each of the three panels in the figure is from the OLS regression shown above and indicates a positive correlation between remote listings and the logged number of female applicants, URM applicants, and average experience, respectively.

The second through fourth row of each panel are coefficient estimates when using matched samples (as an empirical strategy of addressing endogenous remote listings). The second row uses the same time window as the first, but constructs a sample that includes

¹³The secular trend of total applications and applications per job post has been trending down on the platform, even before the pandemic shutdowns. The pattern may reflect relatively high platform engagement in the 2018 time frame in which certain technology domains, such as “web3”, were peaking. In addition, it appears that the onset of the pandemic dampened employer job recruiting behavior, perhaps due to challenges of onboarding new hires. Because we are matching on week of job posting in our CEM approach, this secular trend may be less important in our result interpretation.

remote listings and onsite listings that are matched on weeks since posting, job role, wages, and equity, where job role is an exact match, and wages, equity, and weeks since posting are allowed to be coarsened matches.¹⁴ Then the specification described above is estimated on the restricted matched sample. The second row in each panel indicates that constructing a matched sample does not significantly alter the point estimate on the remote coefficient compared to the (unmatched) OLS estimates from the first row in each panel.

The implied economic effects of the estimates from the second row, in which matching takes place over the entire time window (2018-2022), are significant. A discrete change in job posting to *remote* status (holding all else constant) is associated with approximately 15% increase applicants who are female, 33% increase in applicants with URM status, and 17% increase in applicant experience.

The third and fourth rows in each panel construct matched samples from the 2019 data and then the 2021-2022 data, respectively. These time periods were chosen because 2019 predates any of the changes associated with the pandemic, and 2021-2022 excludes the 2020 adjustment period in which employers and applicants were uncertain about future policies.

We can see that the diversity of the applicants (fraction of female and URM applicants) is higher, in that there is greater representation of female and URM applicants for remote listings, in the 2018-19 pre-period. The final row in each panel shows that in the post-closure period, when remote listings become more prevalent, these diversity differences are attenuated (there is less separation of applicants into jobs by demographic). However, as before, the estimated economic effect of remote work in attracting URM applicants is greater compared to female applicants. The effects of remote listing on applicant experience are similar before and after the closures, even when using the matched sample.

The estimates from these regressions on both matched and unmatched samples support the interpretation that designating a job opening as remote-eligible is associated with a labor pool that has more female, URM, and experienced candidates. The magnitudes of the point estimates from these models are consistent with the interpretation that remote listings receive 1-2 more applications from URM or female candidates, after conditioning on the total number of applications received by the posting.

The specifications that have been used to this point use fixed-effects for broad job roles, which come from a taxonomy specified by the job provider. Figure 7 substitutes fixed-effects for fine-grained job titles included in the job listings. Because of the long tail of job listings that appears in this sample, we limit the sample to listings that fall into the 1,000 most common job titles reported in these postings. We replicate our baseline matched sample specification using fixed effects for week of posting and job title (where we use title now instead of job category) along with salary, equity, and number of applications to the listing.

¹⁴Details about the matching and the covariate balance within the matched sample are in Appendix A.

This substitution does not impact the sign or significance of our *remote* coefficient although the magnitude of the estimates is smaller than in Figure 6, which suggests that some, but not all, of the demographic differences in remote jobs can be attributed to job title heterogeneity. When holding job title fixed, we still observe a positive association between *remote* posting and the fraction of URM and women applicants, as well as applicant experience.

Appendix B reports further results from specifications using a more complete panel of fixed-effects, including location and employer, at the expense of restricting the sample size. In general, the results from these specifications show that within the restricted sample, these effects are robust to accounting for time-invariant firm and employer characteristics, such as firm size, industry, or the geographic location of the job.

Rather than estimating a single effect as in Figure 6, we can alternatively estimate a distributed lag model, which has the benefit of tracing out the evolution of changes in applicant behavior, at the expense of having many more parameters to estimate. To illustrate how applicant behavior changes in the two years before and after the March 2020 closures, consider Figure 8, which again uses the matched sample. This figure plots the estimated coefficient on the interaction term between *remote-week* in which the outcome variables are the (logged) fraction of female applications, fraction of URM applications, and applicant experience. The control variables in the OLS specification include (logged) salary, equity, number of applicants, and job-role fixed effects. The vertical dashed line designates the closure week, and Figure 8 illustrates the week-by-week patterns (with the point estimate designated with a dot, and the vertical black bar displaying the 95% confidence interval). This figure provides a visual representation of the temporal evolution of applicant response to remote-eligible jobs. The pattern seems most pronounced for URM candidates.

A separate, non-matching based empirical approach focuses our analysis on jobs that have onsite characteristics, but were forced into the remote condition due to changes in the pandemic business environment. This allows us to estimate applicant interest for jobs which *would have been on-site but for the pandemic closures* - and so are instead remote. In this manner, we estimate the change in applicant characteristics (diversity and experience) stemming from the exogenous switch to remote work in a way that is separate from our main matching approach and so allows us to triangulate the results. We briefly describe the set-up and results here, but present details in Appendix C. In the first step, we specify a logistical regression model on job listings posted in 2019 to predict onsite status. The explanatory variables are job characteristics. This allows us to conduct the analysis of interest, estimating female, URM, and experience effects on job listings which we predicted would have been onsite pre-pandemic but post-pandemic, are largely remote. The results are in Figure 9, and are consistent with the results from the earlier matching approach. In summary, the *remote* coefficient is highly statistically significant and the estimates suggest

that such job status is associated with a 10%, 21%, and 18% increase in female applicants, URM applicants, and applicant experience, respectively.

4.2 Remote effects by geographic location

We also explore remote-diversity effects by local labor market pool diversity in Figure 10. To prepare the data, we sort city labor markets by local applicant diversity, proxied by the composition of applicants to onsite jobs.¹⁵ For the URM case, the lowest group, quintile 1, represents the least URM diverse markets, while quintile 5 is the most URM diverse markets. We take the same approach when arraying labor markets into gender diversity and experience quintiles. The figures plot the regression coefficients and 95% confidence intervals for each quintile separately using the OLS estimating equation associated with Figure 6. The pattern for URM candidates is clear: the estimated effect of listing a job as remote-eligible in geographies which have the least diverse local markets is highest, whereas the locales with the most diverse local markets experience much smaller benefits in terms of applicant diversity. We find muted effects for gender, and none for experience, which may be as expected because the literature has principally emphasized geography and migration costs in the context of ethnicity.

Figure 11 shows which “donor” cities are most responsible for the increased URM representation for remote listings in the lowest quintile of the URM panel in Figure 10. We restrict the data to only include applicants to job listings posted in the lowest quintile markets. We then estimate whether an applicant is likely to be URM based on job role and city of applicant. Figure 11 plots the magnitudes of the eight largest city fixed-effects. Miami, Rochester, and Atlanta have the three largest coefficients, and are all cities that are well known for the sizes of their URM communities.

4.3 Compensating differentials for remote work

The job search preferences reported by workers along with the wage fields in these job postings data allow us to compute preferences for remote work as well as the revealed compensating wage differential (if any) that workers associate with remote jobs as compared to onsite positions.

The results of these analyses are shown in Table 3. First, column (1) shows the results of workers’ self-reported preferences for remote work. When workers join the platform, they submit profile information about themselves, including the types of jobs they would prefer, their preferences for remote work, and as discussed above, their demographic information.

¹⁵To ensure that this method of sorting cities did not generate results mechanically, we also employed an alternative sorting approach based on Census data on cities’ racial composition. This approach produced very similar results.

Column (1) analyzes the relationships between demographics and preferences for remote work. A logit specification that tests which workers prefer remote work (workers in this category had to specify that they would exclusively consider remote positions) shows that women and URM candidates both have a preference for remote work, conditional on job role and experience levels.

Column (2) estimates compensating wage differentials based on the application choices of workers. If workers have different wage thresholds at which they are willing to consider remote and onsite work, we might observe this in the jobs to which they apply. In column (2), the unit of observation is a job application, and we estimate how much lower posted wage levels tend to be when the targeted job is remote. Our compensating differential estimates for remote work indicate posted wages that are about 7.5% lower when remote jobs are targeted.¹⁶ The muted effects on the interaction terms between remote and URM/gender suggest that conditional on jobs being remote, shifting the wage level does not have a large impact on the composition of the applicant pool. These estimates are due to differences in preferences both within and across applicants for remote jobs. In column (3), we add applicant fixed-effects to investigate how much workers are willing to give up to work remotely after accounting for heterogeneity in applicant preferences, and find that the size of the compensating differential falls by about half.

We can also leverage the stated preferences of applicants for remote work in our analysis. Figure 12 compares the remote coefficient in column (4) of Table 3 when the sample is restricted to groups with the following stated preferences: ‘Remote only’, ‘Remote preferred’, ‘Onsite or remote’, or ‘Onsite preferred’. The compensating differential is largest for applicants with preferences for remote work, and falls to zero for applicants who prefer onsite work. Appendix E conducts further analyses with preferences and applicant level behavior and shows that the stated preferences of remote work by female and URM applicants are consistent with their application behavior and that a shift in applications to remote work was driven by within-user changes in preferences. This appendix also reports the results of an analysis that evaluates, at an application level, whether remote designation on a job predicts applicant demographics and how this shifts in the pre- and post-periods.

5 Discussion

While the basic technological infrastructure ingredients enabling remote work may have been in place prior to the COVID-19 pandemic, the added behavioral shift (on the part of both employees and employers) facilitated by the prolonged work-from-home environment

¹⁶These estimates are in line with those reported by [Mas and Pallais \[2017\]](#) who find a compensating differential of 7% within their call center empirical context. That study makes use of a wider set of control variables for worker demographic characteristics relative to what is available in our setting. This also motivates our inclusion of worker fixed effects in the final specification of the table.

gave more individuals direct WFH experience. As the US emerges from the pandemic, a salient managerial decision is job design, including the scale and scope of remote work. This paper is a window into what might happen in the startup labor market, particularly in terms of applicant diversity and experience, should employers elect to make jobs remote-eligible.

Given the competitiveness and importance of the talent pool, especially for startups (whose assets are typically disproportionately intangible and knowledge-based), our findings on women and especially URM job applications spurred (15% for the former group, 33% for the latter) by the job design decision of making a job remote-eligible are notable. Our compensating wage differential analysis quantifies this preference (7% of wages for women, 9% of wages for URM). Given the disproportionately technical and managerial nature of jobs typically listed on the platform from which we draw data, however, the generalizability of the results to other types of jobs remains to be seen.

An area for result interpretation in our study is the extent to which potential employers in our data decide to wait out the COVID shutdowns and forestall their hiring activities.¹⁷ While we track job postings and applicant behavior over two years after the March 2020 quarantine mandate, the full, long-term dynamic equilibrium of labor market adjustments on both sides in the aftermath of the 2020-2022 pandemic may be yet to be revealed.

With this caveat, how do the results link to our theoretical discussion of the potential channels by which applicants might differentially value remote work? We identified three ways in which diverse talent might value remote work: time flexibility, locational flexibility, and limiting face-to-face workplace interactions (the benefits stem from limiting potentially hostile workplace reasons, but at the possible cost of career advancement).

While future research would ideally sharply test the salience of these mechanisms to diverse worker groups, this section discusses inferences on these mechanisms given our empirical results. Our first observation is that from a theoretical standpoint, certain groups of women benefit from remote-work-induced time flexibility (especially those with young families), locational flexibility (dual-career couples), and limiting face-to-face interactions (especially those at risk of workplace misconduct). For URM, the literature suggests remote-work benefits accrue mainly due to locational flexibility (URM are less likely to reside in the same location as high opportunity employers) and limiting face-to-face interactions (microaggressions). Our empirical finding that URM applicant interest is much more responsive to remote-eligible jobs relative to women suggests that the benefits to URM, while numerically fewer in categories relative to women, may apply more generally and deeply. Perhaps

¹⁷If such behavior were randomly distributed among hiring firms net of the controls (recall that our specifications include fixed effects for location, time, and organization), this would not present a problem. However, if such firms' job posting behavior is systematically related to their beliefs about their ability to successfully manage remote work and/or the job skills they are recruiting for (in ways which we do not measure and control), our results may be biased.

this is due to the remote-work benefits accruing only to women in particular situations described above. In addition, our geographic empirical evidence that URM in the least diverse cities (where URM are the minority) value remote-work much more than URM in the most diverse cities suggests that spatial mismatches between URM worker location and employment locations are an important driver of deeper interest in remote work. As a result, future research may wish to delve into this aspect more deeply (such spatial mismatches have not generally received much attention, especially in the startup technical labor market literature).

A second inference based on the empirical results centers on the labor market costs of remote work with regard to career advancement across URM and women. Our estimates of URM being more responsive in their application interest to remote work relative to women may suggest that URM incur relatively lower costs in engaging in this form of alternative work arrangement relative to women. The extant empirical evidence suggests that women suffer from feedback on their work product when they are not working on-site, while the main harm to URM in a parallel situation may be in the form of limited external labor market access (should their internal career stall). The cost to women seems more general relative to the circumstance for some URM individuals in this regard. Again, our empirical results highlighting the heterogeneity of URM applicant reaction according to city diversity are consistent with a potential conditional cost for this group. This suggests that more generally, future work on remote and alternative work arrangements should examine differences across women and URM as well as heterogeneity within them.

With respect to additional avenues for future research, given our finding that diverse talent prefer remote jobs (all else equal) together with the finding from the literature that onsite work is associated with better career progression outcomes, follow-on work may examine the extent to which workplace policies and interventions, such as engineering social interactions for remote and hybrid workers (e.g., [Choudhury et al. \[2020\]](#)), may mitigate damaging career promotion prospects.

While we do not observe productivity or other worker-level output measures in this study, future work in this domain would ideally compare the quality of talent recruited as well as the productivity and creativity of employees under different alternative work arrangements. Doing so would provide greater insight into the potential broader impact of the shifting applicant behavioral patterns we document, as well as expose potential managerial policy interventions. For example, an outcome of enhanced worker productivity from remote work followed by poor promotion prospects implies different policy choices as compared to the same promotion outcome, but with reduced worker productivity.

Another fruitful domain for future work would be to explore the boundary conditions of remote work and the future of organizing and managing work, especially as related to

competitive (startup) labor markets. As suggested by our compensating wage differential analysis (which reinforces prior findings such as [Mas and Pallais \[2017\]](#) and [Chen et al. \[2019\]](#)), since individuals value remote work in financial terms, it may be the case that employers may be more likely to offer such job flexibility in competitive labor markets. Relatedly, we suspect that the concept of “jobs which can be effectively performed remotely” is not exogenously-given and fixed. However, it may take managerial and process investments for organizations to learn to effectively manage in such environments. For example, [Srikanth and Puranam \[2011\]](#) find that coordinating tacit knowledge even among geographically-distributed work teams is possible.¹⁸ Other areas to explore include organizing interdependent (as compared to modular) remote team work, the degree to which newly-hired versus existing employees are eligible for certain configurations of remote work, and more generally how to reduce conflict in geographically distributed teams [[Hinds and Bailey, 2003](#)].

To conclude, it is likely that the entire category of “alternative” work and job design is going to be the subject of much experimentation going forward, probably in ways much more subtle than that which much of the world experienced during the COVID-19 shutdowns. Hybrid work is one arena already being studied (e.g., [Bloom et al. \[2022\]](#), [Lewandowski et al. \[2022\]](#)), but we believe there are a myriad of alternatives. As a result, we hope that the work presented here is a window to understanding remote work and the startup labor market, particularly with regard to gender and racial diversity, but recognize there is much work to be done to understand this aspect of the future of organizing work.

References

- Aksoy, C. G., Barrero, J. M., Bloom, N., Davis, S. J., Dolls, M., and Zarate, P. (2022). Working from home around the world. Technical report, National Bureau of Economic Research.
- Atkin, D., Chen, M. K., and Popov, A. (2022). The returns to face-to-face interactions: Knowledge spillovers in Silicon Valley. Technical report, National Bureau of Economic Research.
- Atkin, D., Schoar, A., and Shinde, S. (2020). Worker sorting, work discipline and development. Technical report, MIT Sloan working paper.
- Atkinson, R., Muro, M., and Whiton, J. (2019). The case for growth centers. *Brookings Institution: Washington, DC, USA*.
- Bai, J. J., Brynjolfsson, E., Jin, W., Steffen, S., and Wan, C. (2021). Digital resilience: How work-from-home feasibility affects firm performance. Technical report, National Bureau of Economic Research.

¹⁸While we only consider individual demographics within the context of race and gender based diversity, there are other areas of diversity to consider, such as tenure and functional diversity. Elevating the level of analysis to teams along the expanded set of diversity dimensions and considering both behavior and performance also merits future study as they relate to remote teams.

- Barbulescu, R. and Bidwell, M. (2013). Do women choose different jobs from men? mechanisms of application segregation in the market for managerial workers. *Organization Science*, 24(3):737–756.
- Bartik, A. W., Bertrand, M., Cullen, Z., Glaeser, E. L., Luca, M., and Stanton, C. (2020). The impact of COVID-19 on small business outcomes and expectations. *Proceedings of the National Academy of Sciences*, 117(30):17656–17666.
- Benoit, D. (2021). Jamie Dimon on booming economy and finally getting off zoom. *Wall Street Journal*.
- Benson, A. (2014). Rethinking the two-body problem: The segregation of women into geographically dispersed occupations. *Demography*, 51(5):1619–1639.
- Bloom, N., Han, R., and Liang, J. (2022). How hybrid working from home works out. Technical report, National Bureau of Economic Research.
- Bloom, N., Liang, J., Roberts, J., and Ying, Z. J. (2015). Does working from home work? Evidence from a Chinese experiment. *The Quarterly Journal of Economics*, 130(1):165–218.
- Brynjolfsson, E., Horton, J. J., Ozimek, A., Rock, D., Sharma, G., and TuYe, H.-Y. (2020). COVID-19 and remote work: An early look at US data. Technical report, National Bureau of Economic Research.
- Chen, M. K., Rossi, P. E., Chevalier, J. A., and Oehlsen, E. (2019). The value of flexible work: Evidence from Uber drivers. *Journal of Political Economy*, 127(6):2735–2794.
- Cho, K. S. and Brassfield, L. (2022). An afterthought: Staff of color and campus wellness within higher education responses to COVID-19. *American Behavioral Scientist*, page 00027642221118254.
- Choudhury, P., Foroughi, C., and Larson, B. (2021). Work-from-anywhere: The productivity effects of geographic flexibility. *Strategic Management Journal*, 42(4):655–683.
- Choudhury, P., Lane, J. N., and Bojinov, I. (2020). Virtual water coolers: A field experiment on the role of virtual interactions on organizational newcomer performance. Technical report, Harvard mimeo.
- Cullen, Z. B. and Perez-Truglia, R. (2023). The old boys’ club: Schmoozing and the gender gap. *American Economic Review*, 113(7):1703–1740.
- Dingel, J. I. and Neiman, B. (2020). How many jobs can be done at home? *Journal of Public Economics*, 189:104235.
- Emanuel, N. and Harrington, E. (2020). ‘Working’ remotely? selection, treatment and the market for remote work”. Technical report, Harvard mimeo.
- Emanuel, N., Harrington, E., and Pallais, A. (2022). The power of proximity: Office interactions affect online feedback and quits, especially for women and young workers. Technical report, Harvard mimeo.
- Fernandez, R. M. and Fernandez-Mateo, I. (2006). Networks, race, and hiring. *American Sociological Review*, 71(1):42–71.
- Gompers, P. A., Kaplan, S. N., and Mukharlyamov, V. (2022). Private equity and COVID-19. *Journal of Financial Intermediation*, page 100968.
- Hansen, S., Lambert, P. J., Bloom, N., Davis, S. J., Sadun, R., and Taska, B. (2022). Remote work across jobs, companies, and countries. Technical report, working paper, July.

- Harwood, S. A., Hunt, M. B., Mendenhall, R., and Lewis, J. A. (2012). Racial microaggressions in the residence halls: Experiences of students of color at a predominantly white university. *Journal of Diversity in Higher Education*, 5(3):159.
- Hensvik, L., Le Barbanchon, T., and Rathelot, R. (2021). Job search during the COVID-19 crisis. *Journal of Public Economics*, 194:104349.
- Hinds, P. J. and Bailey, D. E. (2003). Out of sight, out of sync: Understanding conflict in distributed teams. *Organization Science*, 14(6):615–632.
- Holland, K. J., Cortina, L. M., Magley, V. J., Baker, A. L., and Benya, F. F. (2020). Don’t let Covid-19 disrupt campus climate surveys of sexual harassment.
- Kain, J. F. (1968). Housing segregation, negro employment, and metropolitan decentralization. *The Quarterly Journal of Economics*, 82(2):175–197.
- Kaur, S., Kremer, M., and Mullainathan, S. (2015). Self-control at work. *Journal of Political Economy*, 123(6):1227–1277.
- Le Barbanchon, T., Rathelot, R., and Roulet, A. (2021). Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics*, 136(1):381–426.
- Lewandowski, P., Lipowska, K., Smoter, M., et al. (2022). Mismatch in preferences for working from home-evidence from discrete choice experiments. Technical report, Instytut Badan Strukturalnych.
- Mas, A. and Pallais, A. (2017). Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–59.
- Myers, K. R., Tham, W. Y., Yin, Y., Cohodes, N., Thursby, J. G., Thursby, M. C., Schiffer, P., Walsh, J. T., Lakhani, K. R., and Wang, D. (2020). Unequal effects of the COVID-19 pandemic on scientists. *Nature Human Behaviour*, 4(9):880–883.
- Pager, D. and Pedulla, D. S. (2015). Race, self-selection, and the job search process. *American Journal of Sociology*, 120(4):1005–1054.
- Parasurama, P., Ghose, A., and Ipeirotis, P. G. (2020). Determinants of occupational segregation across race and gender: Evidence from sourcing, screening, and hiring in IT firms. *NYU Stern School of Business*.
- Sorenson, O. and Dahl, M. S. (2016). Geography, joint choices, and the reproduction of gender inequality. *American Sociological Review*, 81(5):900–920.
- Srikanth, K. and Puranam, P. (2011). Integrating distributed work: comparing task design, communication, and tacit coordination mechanisms. *Strategic Management Journal*, 32(8):849–875.
- Sue, D. W., Capodilupo, C. M., Torino, G. C., Bucceri, J. M., Holder, A., Nadal, K. L., and Esquilin, M. (2007). Racial microaggressions in everyday life: Implications for clinical practice. *American psychologist*, 62(4):271.
- Wu, L., Waber, B. N., Aral, S., Brynjolfsson, E., and Pentland, A. (2008). Mining face-to-face interaction networks using sociometric badges: Predicting productivity in an IT configuration task. *Available at SSRN 1130251*.

Figure 1: Time series of jobs listed as remote




Figure Notes: This chart illustrates the fraction of full-time, US-based job listings in our sample that are tagged as remote-eligible. The x-axis denotes the number of weeks since the March 15, 2020 mandated COVID closures with the orange line indicating the date of closure. The y-axis is the 15 day rolling mean of the fraction of listings that are remote-eligible.

Figure 2: Employer size and remote-eligible jobs

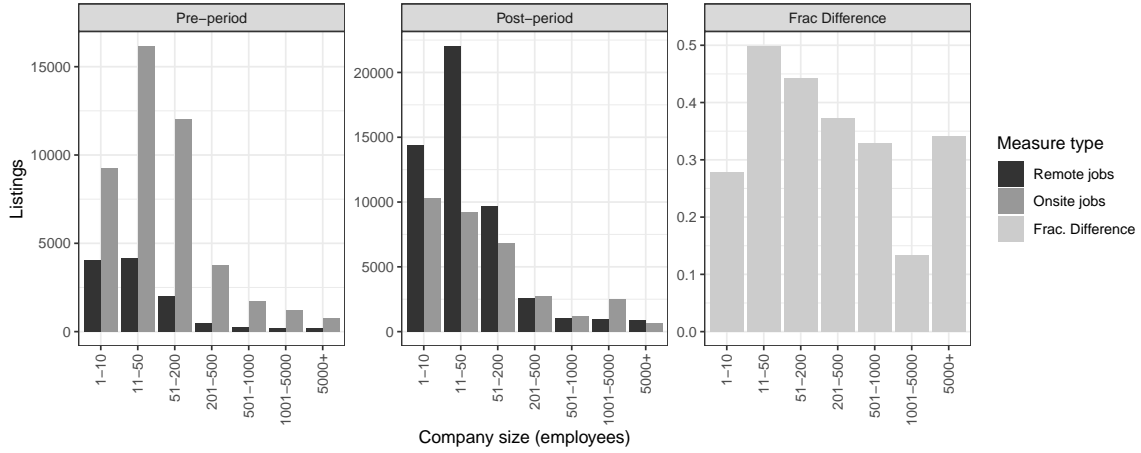


Figure notes: These charts report changes in the distribution of full-time, US job listings by company size before and after the pandemic-induced shutdowns in March 2020. The pre-period in panel 1 includes dates from Jan 1, 2018 through March 15, 2020. The post-period in panel 2 are dates after March 15, 2020. The gray bars designate the number of onsite listings, and the black bars designate the number of remote-eligible listings. Panel 3 reports the change in the fraction of overall listings accounted for by remote listings from the pre-period to the post-period.

Figure 3: Heat map of remote listings by job role and wage




Figure Notes: These heat maps illustrate the fraction of full-time, US-based job listings that are designated as remote. Each cell illustrates the fraction of job listings that are remote by wage bracket and job role. The top panel includes job listings posted before or on March 15, 2020, and the bottom panel is job listings posted after that date.

Table 1: Logit predictions of remote designation for job listings

Sample	2018-2019		2021-2022	
	Coef	Std. Err.	Coef	Std. Err.
Model:	(1)		(2)	
<i>Variables</i>				
Employees (11-50)	-0.425***	(0.054)	0.062	(0.078)
Employees (51-200)	-0.824***	(0.081)	-0.312*	(0.185)
Employees (201-500)	-0.823***	(0.110)	-0.489	(0.298)
Employees (501-1000)	-0.339**	(0.170)	0.438	(0.484)
Employees (1001-5000)	-0.843***	(0.166)	-0.268	(0.525)
Employees (5000+)	0.412**	(0.170)	0.002	(0.367)
Engineering role	0.034	(0.070)	-0.079	(0.102)
Management role	-0.071	(0.107)	-0.277*	(0.155)
Marketing role	-0.124*	(0.074)	-0.062	(0.111)
Operations role	-0.381***	(0.072)	-0.871***	(0.113)
Product role	-0.532***	(0.103)	-0.084	(0.127)
Sales role	0.037	(0.077)	-0.144	(0.118)
Log(Salary)	-0.300***	(0.016)	-0.129***	(0.031)
Log(Equity)	0.487***	(0.032)	0.011	(0.059)
Mean values (Remote)	0.2201		0.8101	
Observations	36,900		14,705	
Squared Correlation	0.068		0.017	

Table notes: The models in all columns are logit regressions with each observation corresponding to a job posting and the dependent variable indicating whether or not the job is designated as remote. *Post* indicates whether the job is posted after March 15th, 2020. The omitted category for firm size is 1-10 employees, and for job role, it is Designer. The first column is restricted to listings posted in 2019. The second column is restricted to listings posted from 2021 onwards. Standard errors are clustered on employer.

Table 2: Characteristics of job listings

Var	All	Onsite	Remote
Number of apps	50.72 (138.219)	31.247 (87.672)	80.448 (187.277)
Number of female apps	10.205 (29.915)	6.247 (17.16)	15.807 (41.123)
Number of URM apps	7.584 (22.441)	3.827 (9.955)	12.548 (31.539)
Mean experience in years	5.611 (2.398)	5.469 (2.571)	5.811 (2.114)
Salary (thousands of USD)	120.71 (55.437)	118.563 (52.111)	123.328 (59.135)
Equity	1.898 (5.519)	1.306 (4.603)	2.62 (6.388)
Observations	465,865	184,387	281,478

Table notes: This table reports means and standard deviations for key variables used in the analysis sample. Odd (even) rows report means (standard deviations). The first (second) row is applications submitted to job listings. The third (fourth) row is number of female applicants to each listing, for those applications where gender is reported. The fifth (sixth) row is number of applicants that are submitted by URM candidates, from those applications where users report race (42% of the sample of job applications). The seventh (eighth) row is years of experience of applicants for job listings that received at least one application. The ninth (tenth) row is salary and the eleventh (twelfth) row is equity figures (reported in percentage points).

Figure 4: Applications sent to remote and onsite job listings

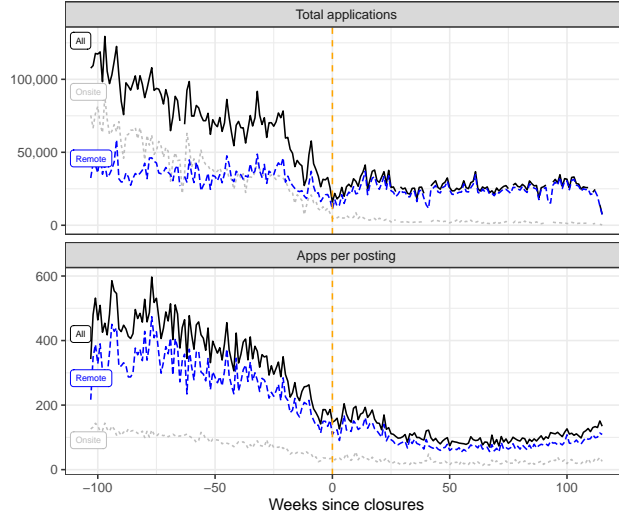


Figure Notes: This figure illustrates changes in numbers of applications to US, full-time job listings in our sample over the time window spanning two years before the pandemic-induced closures and about two years after the closures. The solid line is total number of applications to all job openings in this time window. The dotted gray line is applications to onsite applications only. The dotted blue line is applications to remote positions only. The dashed vertical line at the "0" indicator is the week corresponding to March 15, 2020. The top facet illustrates all applications and the bottom facet illustrates applications per posting.

Figure 5: Differences in characteristics of applicant pools to remote and onsite jobs




Figure Notes: These trend lines illustrate the time-series of the differences in the fraction of applicant types to remote and onsite work. In each panel, the sample of jobs is limited to full-time, US listings. The top panel plots the difference in the fraction of female applicants to remote and onsite listings over the sample period. Most of the series is below zero, suggesting more female applicants to onsite relative to remote jobs. The second panel plots the difference in the fraction of URM applicants to remote and onsite listings during the sample period, where the application sample is limited to cases where applicants report race. The third panel is the difference in industry experience of applicants to remote and onsite listings. In all panels, the y-axis is the difference between remote and onsite jobs. In all panels, the x-axis represents weeks since closures and the orange line is March 15, 2020, the date of the pandemic-induced closures.

Figure 6: OLS and matched sample estimates of remote coefficient




Figure Notes: These figures illustrate the results of OLS and matching tests. The sample is limited to US full-time job listings from 2018 to 2022. The reported point estimate is the coefficient estimate on the *remote* indicator in the specification: $y_i = \text{remote}_i + \log(\text{salary})_i + \log(\text{equity})_i + \log(\text{apps})_i + \gamma_i + \epsilon_i$. Here, y is the logged measure shown in the panel header (female apps, URM apps, average experience of applicants), γ is a vector of fixed effects including week of posting and job role, and i indexes job listings. In each panel, the first row is the OLS estimate on the full sample. The last three rows are estimates from samples matched on salary, equity, job role, and week of posting with *remote* indicated as the treatment. These last three rows show estimates produced from matched samples drawn from the whole sample, only 2018 and 2019, and only 2021 onwards, respectively. The error bars on each point estimate indicate a 95% confidence interval and standard errors are clustered on employer and week of posting.

Figure 7: Finer grained job title controls on matched sample




Figure notes: This chart illustrates the estimate on the remote coefficient using the same specification as in Figure 6, but instead of broad job categories, we substitute fixed-effects for specific job titles. The sample is limited to postings where the posted job title is one of the 1,000 most frequently appearing job titles.

Figure 8: Distributed lag model estimates of remote effects

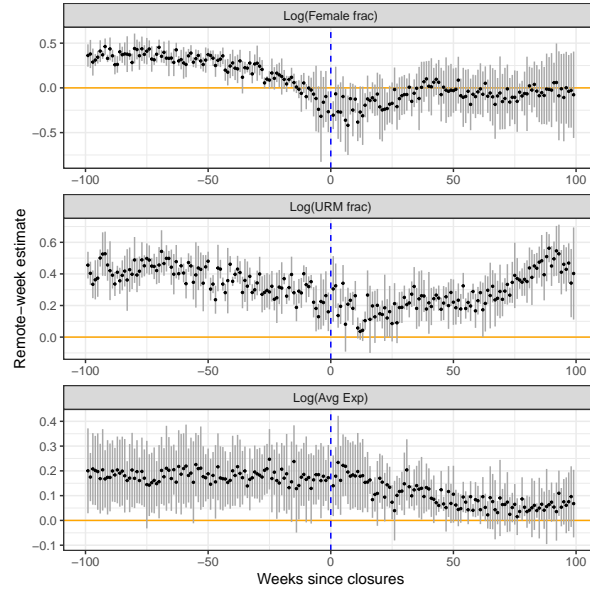


Figure Notes: This figure illustrates the coefficient estimate on $remote \times week$ from the following specification: $y_i = remote_i \times week_i + \log(salary)_i + \log(equity)_i + \log(apps)_i + \gamma_i + \epsilon_i$ where $week_i$ is the number of weeks since the pandemic closures and γ_i includes job-role fixed effects and i indexes the job posting. The sample is the same used in Figure 6. The lines in grey indicate 95% confidence intervals. y_i is the logged fraction of female applications, the logged fraction of URM applications, and the logged average experience of applicants in the first, second, and third panels, respectively. The time period in each panel is limited to 100 weeks before and 100 weeks after the pandemic-induced closures.

Figure 9: Estimates using listings that would have been onsite pre-pandemic




Figure notes: This figure estimates the model $y_i = Remote_i + \log(Salary) + \log(Equity) + \log(Apps) + \gamma_i + \epsilon_i$ but it is restricted to a sample of listings that were posted after the pandemic and would have been predicted to be onsite, based on job characteristics and activity in the pre-period. Further details of how this comparison was constructed can be found in Appendix C. Standard errors are clustered on employer and week of posting.

Figure 10: Remote effect by quantile of local market diversity




Figure Notes: This chart indicates the *remote* coefficient estimate from the specification used in Figure 6 but with the sample divided into five quantiles according to levels of diversity in the local market, proxied by the demographic content of applications to onsite jobs. For this analysis, we only retain markets that received at least 40 onsite applications. For the URM analysis, the lowest quantile (1) includes the least diverse URM markets (Boulder, Seattle, Menlo Park, Redwood City, Somerville, Berkeley, Pittsburgh, Minneapolis, Sunnyvale, Bellevue, Cincinnati, Madison) and the highest quantile (5) includes the most diverse markets (New York, Dallas, Orlando, Atlanta, Miami, San Antonio, Las Vegas, Houston, Indianapolis, Baltimore, Tampa, and West Hollywood)

Figure 11: Top URM donor cities for least diverse local markets




Figure Notes: This figure plots the eight largest city fixed-effects from the regression $urm_i = jobtype_i + city_i + \epsilon_i$ on a sample of all applications to remote jobs posted in the least diverse cities (Quantile 1) from Figure 10. The sample is also restricted to applications that originate from a city that is different than the one in which the job is being posted, and to originating cities that produce at least 500 job applications in this sample.

Table 3: Remote work preferences and compensating differentials by demographic

Dep. Var	Remote	Log(Salary)	
Model:	(1)	(2)	(3)
	Logit	OLS	OLS
<i>Variables</i>			
URM	0.162*** (0.045)	-0.036*** (0.007)	
Remote \times URM		-0.016*** (0.004)	
Female	0.571*** (0.131)	-0.043*** (0.006)	
Remote \times Female		-0.010 (0.008)	
Post-closure	0.248** (0.107)		
Log(Years experience)	0.401** (0.169)		
Remote		-0.075*** (0.019)	-0.041*** (0.001)
<i>Fixed-effects</i>			
Year posted		Yes	Yes
Applicant			Yes
Job role	Yes	Yes	
<i>Fit statistics</i>			
R ²		0.167	0.365

Table notes: This table analyzes preferences for remote work based on data submitted by job-seekers on this platform. The first column is a logit model using stated preferences for remote work as the dependent variable. The second column estimates how the salaries in postings that applicants apply to differ for remote jobs. The third column adds applicant fixed-effects. The number of observations is 50,594 for column (1), is 3,328,354 for column (2), and is 3,527,276 for column (3).

Figure 12: Remote effect by stated preferences for remote work




Figure Notes: This chart indicates the *remote* coefficient estimate from the specification used in column (3) of Table 3 but with the sample stratified by the stated remote preferences of workers.

Online appendices for “Remote work and job applicant diversity: Evidence from technology startups”

David Hsu and Prasanna Tambe
Wharton School, University of Pennsylvania

Appendix A Evaluation of match quality

This section presents results that assess balance in our matched sample. Figure A.1 shows the standardized mean differences of the different matching covariates between the treatment (remote) and matched onsite sample, and between the treatment (remote) and donor samples (all onsite). In the figure, the filled circles indicate mean differences for the matched (adjusted) sample, and the hollow circles indicate differences in the unadjusted sample. We can see that the matching process makes a substantial difference for covariate balance.

Figure A.1: Covariate balance for the matched and unmatched samples




Figure notes: This figure illustrates standardized mean differences between the matched remote and onsite listings sample (Adjusted, filled circles), and for the raw sample (Unadjusted, hollow circles).

Figure A.2 illustrates the distributions of the samples for salary, equity, and numbers of applications before and after the matching process. Again, from the density distributions, we can see that the density plots for salary, equity, and number of applications appear similar for the matched sample.

Figure A.2: Distributions of salary, equity, and number of applications

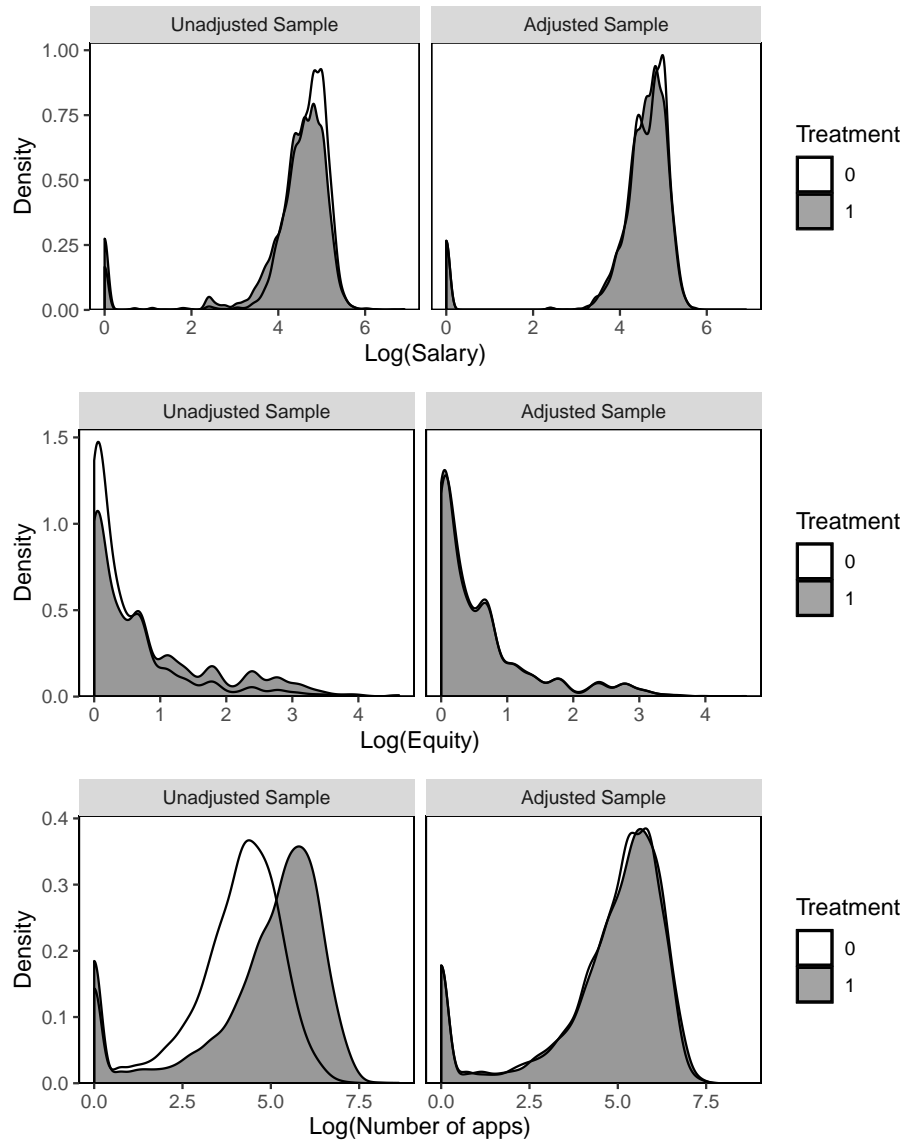


Figure notes: This figure illustrates the distribution of salary, equity, and number of applications for the remote sample, the matched onsite sample, and the complete onsite donor sample.

Appendix B Robustness tests for main specification

In this section, we present the results of additional robustness tests for our main specification. Our baseline regression, presented in the main body of the paper, includes fixed-effects for job type and for week of posting, while including controls for salary, equity, and the total number of applications received by the job listing. In the following sections, we add fixed-effects for city and employer and we vary the matching variables.

B.1 Additional fixed-effects

First, we include fixed-effects for location (city) and employer. We limit the sample to the top 1,000 cities and top 1,000 employers, in terms of the number of applications from each of these categories. The first row of each panel in Figure B.1 replicates the baseline specification on the limited sample restricted to the top 1,000 employers and cities. The second row of each panel then adds fixed-effects for the city in which the job is located according to the job listing. The third row also adds employer, such that the specification in the final row includes fixed-effects for job type, week of posting, city of job, and employer and it includes controls for salary, equity, and number of total applications sent to the posting. In all specifications, the coefficient estimate on the *remote* variable is stable, indicating that the correlations we observed between the remote job designation and the application demographics using our baseline specification are not likely to be attributable to time-invariant factors associated with the location of the job or the employer that is offering the job. The last of these specifications is particularly restrictive, and rules out factors such as company size, city amenities, or industry effects that could otherwise influence application patterns.

B.2 Additional employer-level matching variables

We also conduct our matching on a more extensive set of variables that include firm-level characteristics. Figure B.2 illustrates the results of performing the matching analysis shown in the second row of Figure 6 which uses data from 2018 to 2022. First, we match on firm-level characteristics provided by AngelList which in addition to job title, salary, and equity, include how active the firm is on the site in terms of number of jobs posted, how active their recruiters are, and the size of the company. Second, we match the startups in this sample on a variety of characteristics acquired from the Pitchbook database, which include industry, the year of founding, the number of employees, the financing status of the company, the firm’s ownership status, and the total money raised. The sample size for this analysis is smaller because it is limited to the number of startups that can be matched with Pitchbook. We can see in Figure B.2 that for the URM and experience related dependent variables, the coefficient estimates are similar in direction and magnitude to the results from our main, large sample specification.

Figure B.1: Varying fixed effects on the matched sample

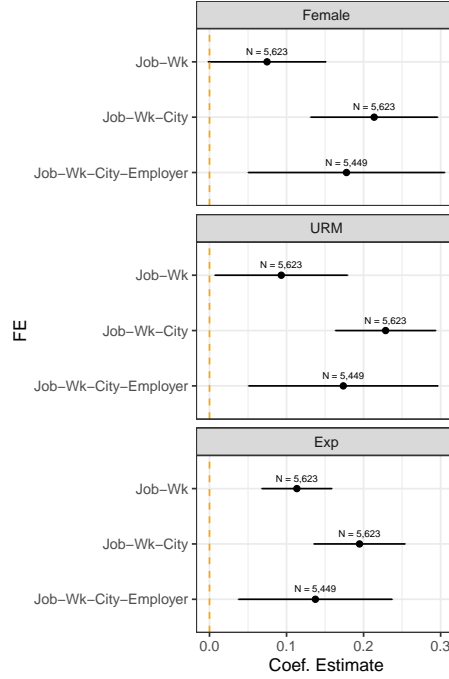


Figure notes: These charts show the remote coefficient when using the specification from Figure 6, except that we add additional fixed-effects. We limit the sample to postings located in the top 1,000 cities and top 1,000 employers in terms of number of listings posted in the city or by the employer, respectively. The estimate in the first row uses the specification from Figure 6 on the restricted sample. The second row adds fixed-effects for the city of the posting and the third row adds fixed-effects for employer.

Figure B.2: Additional firm-level variables for matching




Figure notes: This chart illustrates the estimate on the remote coefficient using the same specification as the second row in each of the panels in Figure 6 (i.e., matched estimator using data from 2018-2022), but with additional firm-level variables used during matching. The top panel uses firm-level variables from AngelList. The bottom panel uses firm-level variables from Pitchbook, matched to the AngelList employers.

Appendix C Alternative approach based on forced closures

A correlational analysis of remote designation on applicant characteristics carries the implicit assumption that remote listings within the matched sample are exogenously determined. However, a significant confound to this interpretation is that remote listings are shaped by factors associated with the firm and the job position itself (in ways which are unmeasured and/or co-determined).¹⁹ In the main body of the paper, we briefly report results from an approach that uses the shutdowns to estimate the effects on listings that would have been onsite, had it not been for the pandemic (shown in Figure 9 of the main paper). Here, we discuss this approach in detail.

The unanticipated COVID-19 shutdowns afford us an empirical strategy to limit the sample to a set of jobs that was plausibly forced into the treatment due to the pandemic-induced shock. To do so, we identify “treated” listings that are made remote by employers after the work closures, but would have otherwise likely been listed as on-site based on employer and job characteristics. We build a model classifying jobs as remote or on-site given the characteristics of the job. We use it to identify jobs in the post-closure period that – had business conditions not shifted towards remote work – would have likely been on-site.

To predict jobs that would have been onsite based on job and employer characteristics, we train a logistic regression model on job listings posted in 2019. Table C.1 reports classification accuracy when identifying whether a listing is likely to be listed as on-site or remote-eligible, given job features including: job titles, week of year, location, and employer. When applying this prediction model to an out-of-sample baseline–pre-closure sample from 2018—about 90% of the jobs it classifies as onsite were actually onsite. This number shifts significantly when applied to post-closure data (2021), when many jobs that might have normally been onsite had to be made remote. In the 2021 sample, 76% of jobs that were predicted to be onsite were taken remote.

Table C.2 limits the analysis to a matched sample of listings from the post-closure period that would have been predicted as onsite. The specification used here is the same as the one used in Figure 6. This approach, then, separates the effects of (partially unobserved) factors that might differ in remote and onsite jobs from the effects of designating a job with traditionally onsite characteristics as remote.

Our estimates indicate that on the limited matched sample of listings, the estimated coefficient on remote work remains similar, suggesting a positive effect on the fraction of female applicants, the fraction of URM applicants, and the average experience of applicants. In addition to the matched sample analysis in the main body of the paper, this analysis provides further evidence that holding fixed the characteristics of an otherwise onsite job, designating it as remote increases the diversity and experience of the application pool.

¹⁹To fix ideas, consider a simple example of the top five job titles for pre-seed companies in the data: software engineer, account executive, product manager, senior software engineer, and sales development representative. By contrast, for firms which recruit on the platform but which are post-IPO, the corresponding job titles are: enterprise account executive, senior software engineer, solutions consultant, customer success manager, and account executive. This suggests that there is considerable across-firm variation in offering remote work (likely related to unobserved managerial capabilities in managing a remote work force), and furthermore, that there is a dispersion of jobs (and associated skills) which are more or less amenable to the remote work environment - but the data may be generated by unobserved processes and/or may be interdependent with unmeasured variables.

Table C.1: Prediction accuracy for jobs predicted to be onsite

Year	Actually onsite	Actually remote
2018	89.57%	10.43%
2021	23.82%	76.18%

Table notes: This table reports statistics for prediction accuracy from a logit model that uses 2019 job listings to train a model on whether a job is likely to be posted as onsite or remote, based on job characteristics including role, compensation, and other factors. This model is then used to classify listings in the post-pandemic period as whether they would have been likely to be onsite or remote, in the absence of the pandemic-enabled job shifts. The top row reports the prediction accuracy on out-of-sample (2018) results pre-pandemic and the second row is the prediction accuracy for 2021, the first fully post-pandemic year. This table is meant to provide supportive evidence for Table C.2, which analyzes outcomes for remote post-pandemic listings that would have been onsite had the pandemic not forced job changes.

Table C.2: Estimates restricted to listings that would have been onsite, pre-pandemic

Dep. Variable Model:	Log(Female) (1)	Log(URM) (2)	Log(Exp) (3)
<i>Variables</i>			
Remote	0.101*** (0.031)	0.209*** (0.027)	0.184*** (0.016)
Log(Salary)	-0.216*** (0.034)	-0.166*** (0.038)	0.372*** (0.034)
Log(Equity)	-0.094*** (0.026)	-0.059** (0.027)	0.073*** (0.011)
Log(Num apps)	0.775*** (0.013)	0.744*** (0.011)	-0.079*** (0.006)
<i>Fixed-effects</i>			
Week of posting	Yes	Yes	Yes
Job role	Yes	Yes	Yes
<i>Fit statistics</i>			
Observations	2,407	2,407	2,295
R ²	0.82509	0.83897	0.44994
Within R ²	0.78054	0.82019	0.33223

Clustered (weeks.since) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table notes: This table estimates the model $y_i = \text{Remote}_i + \text{Log}(\text{Salary}) + \text{Log}(\text{Equity}) + \text{Log}(\text{Apps}) + \gamma_i + \epsilon_i$ but restricted to the sample of listings that were posted after the pandemic and would have been predicted to be onsite, based on 2019 activity.

Appendix D Analysis of missing demographic data

D.1 Demographic measurement and bias

The gender and race field included in the main paper are self-reported by website participants, who can choose whether or not to disclose these fields. In this data set, about 58% of applications are from website participants who choose not to report their race, and about 30% are from applicants who choose not to report their gender. Because this choice is not random, factors that affect one’s willingness to disclose race or gender may be correlated with preferences for remote work. For instance, we find that remote positions receive more applications from women and URM candidates, but women or URM workers who do not report demographic information may prefer onsite work more than their counterparts who do report this information. The following sections probe the sensitivity of our findings to measurement error caused by job seekers selectively not reporting race or gender.

D.2 Demographics inferred from applicant names

To evaluate the degree of missing demographic data problems, we can use other indicators of applicant demographics to assess whether applicants of different race or gender are over or under represented in the group that self-reports these attributes. We can conduct this analysis because for a sub-sample of the data, we have data on applicant names, which can be used to infer gender and race.²⁰ We infer gender and race using the *predict_race* package in R. This package uses Census and Social Security Administration files to generate probabilities of how likely a name is to belong to a gender or race group. Our approach is to sum these probabilities across all applicants to generate the total distribution. For instance, if a name is 80% likely to be female and 20% likely to be male, then the total counts for females and males used to compute the aggregate distribution are incremented by .8 and .2, respectively. This approach yields an estimate of the distribution of race and gender in the sample as predicted solely by name.

Algorithmic inference of demographic attributes from names has become common in the academic literature. Of course, inferred demographic characteristics have their own sources of error. Because they are estimates based on Census data on name-gender-race frequencies, names common in some socioeconomic classes may be less likely to be correctly assigned using algorithmic techniques, which could be correlated with errors in the self-reported data if applicants from these groups are also less willing to provide their demographic information. Nonetheless, this comparison provides a way to calibrate the extent of the reporting error that may exist in the self-reported demographic data.

D.2.1 Comparing self-reported gender with gender inferred from names

Table D.1 reports the distributions of job applicants by i) self-reported gender and ii) gender inferred from applicant name. Column (1) of this table reports the categories from which a site participant can choose when self-reporting. Column (2) reports total counts of participants who self-report who choose each category and column (3) converts these

²⁰This sample is for too narrow a window, only a few months, to replicate the main analysis this way, but we can use it for this demographic calibration exercise.

counts into fractions. We can see that a significant fraction of the applicants – about 64% – do not fill out this field. In column (4), we re-compute this distribution only using data on applicants that do report their gender. We can see that the distribution is imbalanced in favor of men, with over 50% being men, 35% being women, a small number choosing the non-binary designation, and about 13% choosing not to disclose (but did not leave the field blank). The comparable distribution using name inference is shown in column (5). About 13% of applicant names cannot be matched. Among those that can, about 60% are identified as men, about 28% are identified as women, and a very small number as non-binary. In sum, the self-reported demographic data suggests the sample is less imbalanced towards men than might be suggested when using names, but the two distributions are not far enough apart to create concern that the self-reported data significantly misrepresent the site users.

Table D.1: Comparison of reported gender attribution with name-imputed gender

Gender	Count reported	Frac. total	Frac. reported	Frac. inferred
Male	600,016	0.187	0.517	0.595
Female	407,581	0.127	0.351	0.278
Non-binary	3,862	0.001	0.003	0
Other	423	0	0	NA
Not specified	148,335	0.046	0.128	NA
Not provided	2,047,392	0.638	NA	NA
Not matched	NA	NA	NA	0.126

Table notes: This data reports the distribution of gender in the self-reported demographic attribute data. We report information on the demographics of all site registrants, so the numbers on missing gender or race differ from those in the main text which focus on the degree of missing demographic data for applicants to jobs rather than all registered site users. Column (2) is counts in each gender category. Column (3) is the fractional distribution of these counts. Column (4) is the fractional distribution of counts only in categories where gender information was provided by job applicants. Column (5) is the fractional distribution when using names to infer gender.

D.2.2 Comparison of reported race attribution with name-imputed race

Table D.2 compares the distribution of self-reported race with the distribution of race imputed from applicant names. To conduct this comparison, we aggregate the self-reported data for applicants into one of the four broader race categories available in the software that assigns names to races – White, Asian, Hispanic, and Black.²¹ A key observation from the figures in this table is that when compared to the distribution of race categories based on name, Asian applicants appear much less-likely to self-report their race. Partly for this reason, the relative shares of workers in other categories is higher. In terms of our URM classification based on self-reported race used in the main analysis, this means that the fraction of workers labeled URM will be higher in our sample than the true base rate. In the next section, we discuss how this measurement error is likely to impact our estimates.

²¹These categories likely reflect the treatment of race as recommend by the US Census Office of Management and Budget. For more information, see <https://www.census.gov/topics/population/race/about.html>.

Table D.2: Comparison of reported race attribution with name-imputed race

Race	Count reported	Frac. reported	Frac. inferred
White	436,838	0.605	0.539
Asian	93,454	0.129	0.283
Hispanic	99,176	0.137	0.092
Black	92,577	0.128	0.086

Table notes: This table compares workers who self-report their race with the race imputed from their names by race prediction software. For this comparison, we omit applicants who self-report being in the Middle Eastern or Native American categories because these categories do not map to output in the race prediction software. We also omit from this list a single observation labeled as “unmapped”. In the self-reported data, we aggregate applicants in the South asian, East asian, South east asian, and Native Hawaiian or other Pacific Islander categories into a single “Asian” category to match the output of the race prediction software.

D.3 Sensitivity tests of key analyses using applicant names instead of reported demographic data

We can also compare, using the sample with names, how the estimates from our main analysis compares with estimates produced when using gender and race inferred from the applicant name data. That is, we replicate our main analysis but use applicant names to compute female and URM applicants to job listings, rather than self-reported demographic data.²² Figure D.1 depicts the results from the model used in Figure 6 on the restricted window sample, using both self-reported demographics and name-inferred demographics. For all three dependent variables, the estimated coefficient is not significantly different using the two different data sources. Although these analyses cannot conclusively address missing data problems in the self-reported demographic data, they suggest that the measurement error introduced by missing data is unlikely to substantively change the nature of our conclusions, at least when compared with demographic data inferred using other methods.

Figure D.1: Replication of main OLS estimates with name-inferred attributes




Figure notes: This analysis replicates our main estimates of application pool demographics on whether a job is remote using a narrower time window with self-reported demographic data alongside name-inferred demographic data.

²²We do not attempt to draw any strong inferences from these estimates because the data for this analysis spans a few months around the start of pandemic, when the labor market was very much in flux.

Appendix E Heterogeneity in user preferences

This appendix further leverages data on user preferences to corroborate findings in the main section of the paper.

E.1 Stated preferences for remote work

First, we analyze data on workers’ stated remote preferences. Beyond observing their application patterns, workers on this platform state whether they have preferences for remote or onsite work when they first submit their information to this site. Table E.1 compares the remote preferences of male and female workers in our sample. We can see a notable difference, wherein women in aggregate have much stronger stated preferences for work that is wholly or preferred remote. The same is true for URM candidates. These preferences are consistent with those that might be inferred from the findings on application patterns in the main analysis.

Table E.1: Remote preferences by gender

Preference	Female	Male
Remote only	0.263	0.159
Remote preferred	0.401	0.369
Onsite or remote	0.321	0.443
Onsite preferred	0.015	0.029

Table notes: This table reports how remote preferences differ between male and female workers who report gender and remote work preferences.

Table E.2: Remote preferences by race

Preference	URM	Other
Remote only	0.230	0.190
Remote preferred	0.397	0.377
Onsite or remote	0.359	0.407
Onsite preferred	0.014	0.027

Table notes: This table reports how remote preferences differ between URM workers and workers in other categories for workers who report race and remote work preferences.

E.2 Applications by job-type with and without user fixed-effects

Next, we examine shifts in patterns of applications to remote work to assess the extent to which differences within the sample can be attributed to within vs across-user changes in preferences. Figure E.1 shows the changing tendency of an application to be submitted to a remote position from week-to-week during our sample period (the model being estimated is shown in the Figure notes). The solid dots represent the unconditional value of the estimate of each week on whether or not an application is sent to a remote position. Some of this change may be due to applicant preferences changing, and some may be due to workers with remote preferences becoming relatively more active in terms of job search. The hollow dots illustrate the coefficient values of the individual weeks after including user fixed effects.

We can see that after adding user fixed-effects, there is still steady growth towards remote application. This indicates that the shift within users accounts for most of this change.

Figure E.1: Applications to remote, with and without user FE

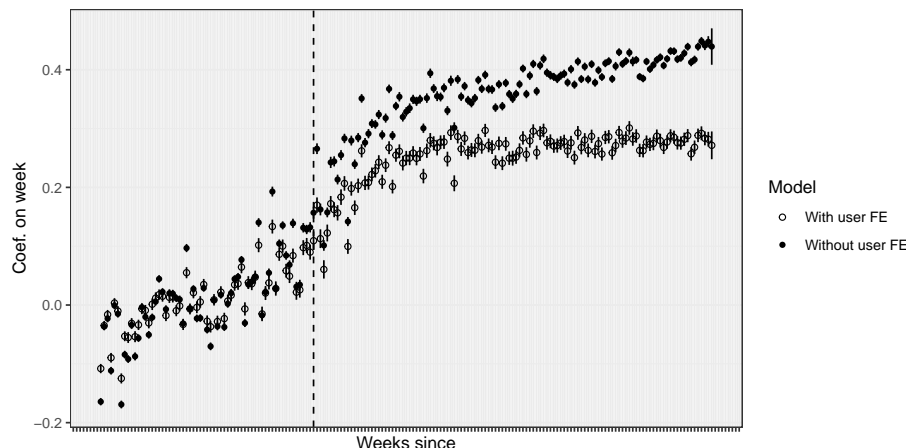


Figure notes: This figure illustrates the shift in applications to remote positions. It illustrates the coefficients from $remote_i = week_i + \epsilon_i$ where i indexes the application. These estimates are reported with and without user fixed-effects.

E.3 Analysis of applications to remote positions

Table E.3 presents an analysis of the relationships between job characteristics and whether an application is being received from a female or URM applicant. Job characteristics include whether the job is remote, firm size, occupational role, salary, and equity compensation. In column (1), we see a positive correlation between remote work and female, i.e. remote work is predictive of an applicant being female, after conditioning on other attributes. The smallest firm size category is omitted, so the coefficients indicate an inverted relationship between size and female. Firms that are small and large are less likely to receive applications from female candidates. Higher salary and higher equity positions receive more applications from men. There is a positive relationship between remote and female, although this relationship fades in the post-period, as reported in column (2).

URM applicants are analyzed in column (3). Both very small and very large firms are less likely to be predictive of an application being from URM candidates. Higher salary and higher equity positions are also less likely to predict applications from URM candidates. One difference that we can see between these two columns is that conditional on other attributes, *remote* is not predictive of receiving an application from a URM candidate.

When interpreting these results, please keep in mind that unlike our main job-posting centric analysis, we have not pre-processed the data to match applicants by characteristics as we do in the job-posting analysis (see Appendix A). In addition, the applicant-centric analysis is not as amenable to empirical specifications which include detailed fixed effects, such as those in our main results (including robustness tests in Appendix B). More generally, there is no strong identification strategy in this applicant-centric analysis such as our main matching specification or from the forced closures (Appendix C). The conclusion is that directionally, the results appear to be broadly consistent.

Table E.3: Analyses of application patterns

DV	Female		URM	
Years	2018-2019	2021-2022	2018-2019	2021-2022
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
Remote	0.064*** (0.012)	−0.028 (0.028)	−0.042*** (0.009)	0.105*** (0.019)
Employees (11-50)	0.043*** (0.015)	0.114*** (0.023)	0.043*** (0.012)	0.057*** (0.015)
Employees (51-200)	0.078*** (0.019)	0.235*** (0.037)	0.053*** (0.014)	0.107*** (0.021)
Employees (201-500)	0.097** (0.041)	0.209*** (0.068)	0.071*** (0.022)	0.046 (0.040)
Employees (501-1000)	0.051 (0.046)	0.342*** (0.108)	0.063** (0.031)	0.158 (0.098)
Employees (1001-5000)	0.116 (0.074)	0.810*** (0.215)	0.105** (0.049)	0.294* (0.153)
Employees (5000+)	0.058 (0.054)	1.017*** (0.133)	−0.002 (0.048)	0.332*** (0.077)
Engineering role	−0.749*** (0.018)	−1.648*** (0.028)	0.114*** (0.014)	−0.436*** (0.018)
Management role	−0.501*** (0.048)	−0.712*** (0.073)	−0.080*** (0.023)	−0.434*** (0.040)
Marketing role	0.274*** (0.019)	0.098*** (0.032)	0.027* (0.014)	−0.138*** (0.019)
Operations role	0.440*** (0.021)	0.440*** (0.034)	0.363*** (0.016)	0.149*** (0.026)
Product role	−0.257*** (0.022)	−0.428*** (0.034)	−0.033** (0.016)	−0.408*** (0.029)
Sales role	−0.182*** (0.020)	−0.325*** (0.037)	−0.013 (0.016)	0.014 (0.019)
Log(Salary)	−0.310*** (0.017)	−0.513*** (0.025)	−0.248*** (0.013)	−0.465*** (0.018)
Log(Equity)	0.006 (0.014)	−0.265*** (0.031)	−0.029*** (0.009)	−0.114*** (0.019)
<i>Fixed-effects</i>				
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	903,159	756,456	919,960	768,031
Squared Correlation	0.048	0.125	0.006	0.022

Table notes: This table estimates correlations between firm and job listing attributes and whether an application is received from a female or URM candidate. Each row in these data is a job listing-application pair for which race and gender information are available for the applicant. Columns (1) and (2) use female applicant as the outcome variable. Column (1) is for the period 2018-2019 and column (2) is for the period 2021-2022. Columns (3) and (4) use URM as the dependent variable. Column (3) is for the period 2018-2019 and column (4) is for the period 2021-2022. Standard errors are clustered on startup.