How does offering remote work affect the diversity of the labor pool? Evidence from technology startups

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

November 3, 2022

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, AngelList Talent, that spans long windows before and after the COVID-19 pandemic-induced shutdowns of March 2020. To overcome 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 job applicants, with larger effects in less diverse geographic areas. When onsite and remote jobs receive similar numbers of applications, remote jobs attract about 1.25 more applications from URM candidates and from women, and applicants have about 1 year more of experience. The average estimated compensated wage differential for remote work is about 7% (for females 17% and for URM candidates 3%) of posted salaries.

Keywords: remote work; WFH; startup labor; technical workers; talent acquisition; diversity and inclusion; equity; COVID-19.

*Correspondence to: dhsu@wharton.upenn.edu & tambe@wharton.upenn.edu. We are grateful for feedback from Matthew Bidwell, Peter Cappelli, Raj Choudhury, Ananya Sen, 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), and the CMU Block Center Future of Work Conference. We particularly thank Kevin Laws at AngelList for data access and Remi Gabillet 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

Startup firms' typical assets are primarily intangible knowledge and human capital-based, which stands in contrast with established firms, which possess a multitude of other assets, not the least of which is an established reputation. Recruiting talent, particularly specialized in technical and managerial areas, is therefore a key organizational challenge, especially for startups.[1] Compounding the issue is the desire of many organizations to diversify (by gender and ethnicity) their recruiting efforts for both equity and productivity reasons (e.g., Yang et al. [2022]).[2]

If one barrier to job applicant diversity is flexibility, allowing remote work may be an important job design decision. 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.[3] A much wider share of the world experienced remote work in 2020-22, with the COVID-19 pandemic due to government-mandated shutdowns, with global remote-eligible job listings in that period rapidly scaling (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 as the world emerges from the pandemic: what WFH policies should firms adopt when they are no longer forced to do so?

There is a wide range of opinion on this question. At one end of the spectrum, some companies such as Twitter and Slack (before their Salesforce acquisition) announced permanent WFH policies. On the other side of the spectrum, firms like Netflix, WeWork, and JP Morgan Chase stated their desire to have employees back in the office as soon as possible. Other organizations, such as Facebook and Google, have taken a hybrid stance, providing some fraction of their workers the flexibility to conduct all of their work remotely and/or combine WFH with on-site work.

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

[1]For example, the Bureau of Labor Statistics projects all occupations in the U.S. will grow 3.7% between 2019-2029, but that science, technology, engineering and math (STEM) jobs will grow 9.5% over the same time interval. Source: https://ncses.nsf.gov/pubs/nsb20212/data. Furthermore, during economic downturns, startups may find recruiting talent even more challenging, as workers may seek the relative safety of incumbent firms (Bernstein et al. [2020]).

[2]The labor supply of women and URM in STEM fields is even more constrained. The 2022 National Science Foundation's Science and Engineering Indicators report, using 2019 figures, indicate that women compose 44% of the bachelor's degree or higher STEM workforce, while Black and Hispanic workers together comprise 15% of the same pool (even while workers with those ethnic backgrounds make up 36% of the 20-34 age group in the U.S.) Source: https://ncses.nsf.gov/pubs/nsb20221/data.

[3]https://www.census.gov/data/tables/2010/demo/commuting/p70-132.html.

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 the applicant pool likely attracted to such work. Specifically, many employers have a renewed goal of striving to make their workforce more inclusive and diverse. 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, 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 scale and granularity (prior studies using detailed labor market data 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, and is centered on job listings posted prior to the COVID-19 shutdowns. We pay particular attention to portraying descriptive patterns because we do not believe such patterns have been previously 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 causal 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). An alternative empirical strategy (which we present 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 literature on job design, 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. Across the estimation approaches, we broadly find that offering remote work affords organizations more experienced applicants, as well as applicants drawn from more diverse gender and race backgrounds. A final analysis examines the heterogeneity of these patterns across labor markets which themselves differ by diversity, and quantifies the intensity of diverse applicant remote preferences via a compensating wage differential analysis.

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. 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

With the rise of a service-based (knowledge) economy, employers are rethinking the role of in-office work, which has implications for the possibility of decoupling the physical location of organizational offices or headquarters and where individuals conduct their work. Indeed, some employers are reconceptualizing the role and physical footprint of corporate offices and workspaces. As previously noted, as the world re-emerges from COVID-19 lockdowns in 2021, the managerial choice of whether and to what degree remote work will be allowed is likely to be a significant one shaping the future of work and job design.

Two recent studies are notable for offering direct empirical evidence on employee productivity and other impacts of remote work. 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 transitioning from being allowed to "work from home" (WFH) to "work from anywhere" (WFA), document increased employee productivity. In the call center setting, the first study found a 13% increase in employee productivity, mainly due to reduced commute-, break- and sick-time, as well as a quieter work environment. The WFH employees were randomly selected from the organization's employees who indicated they would be willing to shift their work in this manner accordingly. The randomized WFH employees 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 in the WFH to WFA transition (such employee status depends in part on employee seniority and request for such status). In addition, this study finds that geographically clustered WFA workers within the same technological unit experience higher productivity (though this effect does not hold for those in different units), suggesting some localization of potential peer effects. While not formally tested, this study also provides qualitative/anecdotal evidence for enhanced employee allegiance to the work organization stemming from the WFA policy.

Related to the theme of productivity, there is almost no literature we could identify directly linking employers' decision to offer remote work and the ability to attract diverse and experienced workers, especially in the setting of more specialized labor markets (technical and managerial roles) as we will consider here. One recent paper using archival data is tangentially related: Yang et al. [2022] report that biomedical research teams producing peer-reviewed scientific publications with mixed gender participants outperformed single-gender research teams in novelty and impact of their published articles (while mixed gender teams are less frequent than would be expected by chance). While future work would ideally examine both possible mechanisms and assess causal links in this domain, the study suggests that organizational policies (including job design) to attract more diverse teams and employees may be in their own interests.

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, and employers benefit via enhanced employee productivity. 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 (resulting in lower promotion rates, which they also find, consistent with Bloom et al. [2015]). Barrero et al. [2020] argue that the COVID-19 shutdowns and widespread experience with remote work has a potential destigmatizing effect and a reason to expect remote work persistence after the COVID-19 pandemic lifts. Beyond mere stigma, in a field experiment of on-the-job market consequences of summer interns' virtual social interaction with top managers (in the summer of 2020), Choudhury et al. [2020] find that randomization into informal social interactions ("water cooler" interactions) with senior managers improves prospects of being offered a subsequent full-time position at the firm. These results suggest both that social interactions in the workplace are important for career advancement, and that it is possible, even in a remote context, to design such interactions (though they are perhaps less likely to occur by happenstance in the virtual as compared to on-site setting).

Another explanation of 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 (in the studies finding a positive WFH treatment effect, the pool from which randomization for the treatment is drawn reflects a prior employee opt-in to the possibility of the treatment condition). The possibility of unobserved employee sorting is also strongly implied by the Emanuel and Harrington [2020] study as well (in the post COVID-19 shutdown period, call center workers hired into remote jobs were 18% less productive than those hired into on-site jobs). These findings are consistent 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] and WeWork's CEO Sundeep Mathrani who commented: "those who are uberly engaged with the company would want to go to the office at least two-thirds of the time, at least" (Dill [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 WFH and remote workforce. This relates to the degree of managerial intensity and skills required to manage in the virtual environment. The managerial infrastructure of the organization may therefore be an important driver of the likelihood of offering remote work. Intertwined with this is job characteristics and the nature of the work demanded. In particular, for some types of work which are highly interdependent among individuals or work teams, the degree of managerial intensity and coordination may be higher; conversely, 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]). This discussion is similar in spirit to an earlier literature predicting the extent to which the output of jobs can be traded at a distance (outsourced or offshored) [Jensen et al., 2005; Blinder et al., 2009].

Not only does the degree of managerial oversight depend on the nature of the work and job, it may depend on employee characteristics such as tenure at the organization in an onsite capacity (which may relate to employee knowledge of organizational norms and culture, degree of trust, etc.). Finally, there may be a competitive element to offering remote work if, for example, local competitors are offering such "non-traditional" work arrangements. Note that these and other factors are difficult to observe and measure, especially at scale, which presents a host of empirical challenges we will discuss at more length in the data section.

2.2 Employees and the decision to engage in remote work

Here, we discuss the employee perspectives which have not yet been mentioned in shaping preferences for remote work. One large theme is that certain demographic groups, such as women with young children and families (e.g., Mas and Pallais [2017]; Atkin et al. [2020]; Barbulescu and Bidwell [2013]) disproportionately value job flexibility, including the ability to conduct remote work, in part due to societal expectations of family roles. During the COVID-19 pandemic, in a survey of scientists, Myers et al. [2020] found that female scientists with child dependents experienced a substantial decline in time devoted to research. Given this disparity, it may be the case that the flexibility accompanying remote work may be especially appealing to women with young children at home. This preference may be compounded by the higher cost of commuting to work women typically face (Le Barbanchon et al. [2021]).[4]

A further reason women may be drawn to remote work is the well-documented empirical pattern (using a variety of methods and samples, including a natural experimental setting (Flory et al. [2015]) that women seek ways to avoid overt competitive dynamics. Applied to the focal context, one conjecture is that onsite work may entail more visible competitive dynamics as compared to remote work. A series of studies reveal a few reasons for this female preference. First, Gneezy et al. [2003] find that women may be less effective than men in competitive environments as proxied by a laboratory experiment tournament incentive scheme in which participants compete with an anonymous opponent in solving maze puzzles (but not in non-competitive contexts such as paying a task piece rate). This may be related to Barbulescu and Bidwell [2013]'s archival data finding that such expectations of success may be an important reason why female applicants may choose (or not) to apply in the first place for a job. Interestingly, in a lab experiment, Niederle and Vesterlund [2007] determined that while performance in a first stage task (solving relatively easy math problems) does not differ across genders, men are twice as likely to opt for a competitive tournament

[4] There is research which also analyzes individual- and time-of-day-specific compensating wage differentials more generally for work flexibility (not specific to gender or URM status) within the context of gig workers using data from Uber [Chen et al., 2019].

compensation scheme in a second stage of the same task, as compared to women (73 versus 35 percent). The authors interpret the results as evidence that men are overconfident about their relative performance as compared to women. As a result, it appears that the female preference to avoid competition is not simply an artifact of expectations of poor performance, but that there may be other issues at play.

Candidate explanations include differences in how seriously applicants take the qualification requirements of job postings, with women interpreting the requirements much more literally as compared to men (Coffman et al. [2020]). Consequently, this study finds that reducing ambiguity on required qualifications may help address the gender gap when it comes to "top of the funnel" considerations for getting more women involved in the labor force. There are of course many other barriers that could exist further down the funnel. For example, the lack of referral networks is especially harmful for racial minorities in their job search (Fernandez and Fernandez-Mateo [2006]). Even the language of job posts can be important: Abraham and Stein [2020], in a randomized control trial, vary the language of how demanding are the required qualifications for a given position. When positions are described with softer qualifications, the gender skills gap among applicants shrinks.

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.[5] This is due to 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. On this last mechanism, Campero [2021] reports that women disproportionately occupy the least competitive (and least prestigious and remunerated) segment of the software development industry: quality control in software. This pattern is consistent with the general literature about women avoiding competitive contexts. Furthermore, Murciano-Goroff [2022] find that women software programmers systematically under-report their skills relative to men (and employers do not adjust their expectations of the difference), pointing to a potential compounding factor in applicant self-selection behavior. Of 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 tend to be driven by choices of the worker, whereas underrepresentation of under-represented minorities (URM) tend to arise from employer choices. These results remind us that it is unlikely that one action will be sufficient to bring labor market underrepresentation among different groups in balance. However, discrete interventions can

[5] 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

help. For example, Del Carpio and Guadalupe [2022] find that women are more likely to participate in a technology skills training program when presented with an example of a female success story.

URM may similarly value flexibility afforded by remote-eligible job positions, also due to household care-giving responsibilities which, as a result of their relative economic status, may not be outsourced to the same degree as other demographic groups. Beyond women and URM, employees in aggregate seem to value flexible work arrangements such as remote work, but how much? 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.

These forces would suggest that a shift in jobs becoming remote-eligible would lead to higher rates of female and URM applicants. However, based on the Bloom et al. [2015] field experiment (among other studies) on remote work and lower promotion rates, women and URM job candidates may be less likely to apply to remote-eligible positions given potential concerns about future promotion prospects. Furthermore, Emanuel et al. [2022] find that onsite work for software engineers in a Fortune 500 firm, particularly women, is associated with greater feedback on their work and lower incidence of separation from the firm. We therefore consider the question of whether the exogenous introduction of remote-eligible jobs will increase the number of gender and ethnically diverse applicants to be largely an empirical one. Remote work flexibility and lowering the competitive context would suggest a positive predicted relationship, while diminished career promotion prospects would predict a negative relation between remote work and female/URM applications. Similarly, how experienced applicants may react to remote-eligible listings is also an empirical matter given the lack of prior literature.

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 rely 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 to improve our understanding of the applicant pool consequences of offering remote-eligible work.

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.[6] 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. While we do have further information about the applicants, including their applications to job postings from within the platform, we have limited information through the platform about activity further down the hiring funnel (such as mutual employer interest). However, we do not 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 the date the US declared a federal emergency, March 13, 2020). 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 resume data, including educational history, skills, prior employers and LinkedIn page) and this is the source of our demographic information. They are then 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 from within the platform. 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

[6]The COVID-19 induced recession of 2020 has affected workers unequally, with demand for technology workers, such as those working for venture capital-backed startups, remaining very 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]).

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 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.[7]

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 three locations being San Francisco, New York City, and London 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 posted jobs designated as being remote, where each x-axis tick corresponds to a week. The illustration spans the period beginning about two 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

[7] Organizations 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%).

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, where the fraction of jobs listed as remote in each category is shaded a different color. The figure indicates that the distribution of listings posted on this platform shifts slightly 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. Unlike in the pre-closure period, management positions are listed with a higher incidence of remote eligibility than engineering or operations jobs, at least 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 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 is the reference category). Consistent with the descriptive patterns, management roles are less likely to be remote, as are product and operations roles. 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 not to have large physical footprints and to offer remote jobs are also those more likely to grant more equity for early stage employees.

The second column in the table broadens the sample to the middle of 2022, and interacts each organizational and job characteristic with an indicator of whether the job was listed before or after the pandemic closures. From this column, we can see that many of the trends from column (1) reverse direction. That is, as a larger part of the economy was forced towards remote work, the correlations between firm size, job type, and compensation are pushed towards zero.

3.2.2 Empirical design examining the consequences of remote listing on applicant behavior

In the focal analysis of the paper, 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 and appendix 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) helps improve inference attributed to the *remote* job effect on applicant characteristics.

A supplemental (separate) empirical approach (which we present in Appendix C) leverages the natural experiment afforded by the COVID-19 discontinuity in the workplace and its implications for hiring (due to government-mandated shutdowns, the decision to offer remote work was largely taken away from managers). We examine applicant pool characteristics to jobs and organizations which were mainly on-site in the pre-shutdown regime but were mainly remote in the post-shutdown period (details of the empirical estimation strategy are contained in the appendix). In this manner, we estimate the change in applicant characteristics (diversity and experience) stemming from the exogenous switch to remote work in a way separate from our main matching approach to triangulate our main results.

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 appears 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).

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). The top panel is the fraction of applications to job listings that are from female applicants, separated by whether listings are remote or onsite. The number of female 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 remote-eligible 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 difference in the applicant industry experience variable also shows a clear pattern: experience is higher across both time periods for remote-eligible listings as compared to on-site listings (though it is hard to tell if that difference is diminishing in the post-closure regime). As is true for all patterns shown in this table, no statistical or 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 includes a window of time that includes the pandemic closures. The key specification we estimate is $y_i = remote_i + log(salary)_i + log(equity)_i + log(numapps)_i + \gamma_i + \epsilon_i$ where γ is a vector of 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. 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 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 the 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. 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. Further details about the matching and the covariate balance within the matched sample are shown in Appendix A.

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). 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.

In Appendix B, we report 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 7, 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 the figure illustrates the week-by-week patterns (with the point estimate designated with a dot, and the vertical black bar displaying the 95% confidence interval). Figure 7 gives a visual representation of the temporal evolution of applicant response to remote-eligible jobs. The pattern seems most pronounced for URM candidates.

In addition to these tests, Appendix C presents supplementary results from an alternative estimation strategy which exploits the sudden shutdowns as a way to focus our analysis on jobs that have onsite characteristics, but were forced into the remote condition due to changes in the business environment.

4.2 Remote effects by geographic location

We also explore remote-diversity effects by local labor market pool diversity in Figure 8. To prepare the data, we sort city labor markets by local applicant diversity, proxied by the composition of applicants to onsite jobs. 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 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 is as may be expected because the literature has principally emphasized geography and migration costs in the context of ethnicity.

Figure 9 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 8. 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 9 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 wage fields in these data allow us to compute the revealed compensating wage differential (if any) that workers associate with remote jobs as compared to onsite positions. We can compute this number in two different ways, by a) analyzing preferences as directly reported by workers, and by) inferring this value from workers' application behavior.

The results of these analyses are shown in Table 3. Columns (1)-(3) show the results of workers' self-reported preferences. When workers join the platform, they submit profile information about themselves, including the types of jobs they would prefer, expected salary, their preferences for remote work, and as discussed above, their demographic information. Column (1) analyzes the relationships between demographics, expected wages, and preferences for remote work. A logit specification that tests which workers prefer remote work (only) shows that women and URM candidates both have a preference for remote work, conditional on job role and experience levels. Overall preferences for remote work also rose after the pandemic closures, as might be expected.

Column (2) shows that controlling for experience, but not job role, a preference for remote work is associated with about 6.4% lower expected salary in the overall worker sample. Women and URM candidates also expect lower salaries, regardless of remote designation. Some of these differences may be due to differences in the types of jobs sought by different demographic groups. The differential grows significantly larger for women, to over 20% when adding the remote main effect with the remote-gender interaction, suggesting that women trade greater wages for WFH flexibility.

In column (3) which includes job role fixed effects, workers expect a 6.8% lower salary for remote work flexibility, even after controlling for occupational differences. The coefficient estimates on the URM and female indicators decline in economic magnitude after including job role fixed effects, suggesting that part of this differential is due to women and URM candidates applying to different job types, which is consistent with what prior work has found. Column (3) also shows that the salary expectations of female employees significantly negatively interacts with remote work (as it did in column (2)), but that the parallel interaction with URM is positive. The estimates suggest a compensating differential for remote work that ranges from about 3% (URM) to 17% (women).

Column (4) uses an alternative approach for estimating compensating wage differentials based on the application behaviors 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 (4), the unit of observation is a job application, and we estimate how much lower posted wage levels tend to be when the job being applied to is remote. Our compensating differential main effect estimates for remote work are similar, indicating posted wages that are about 7.5% lower when applicants are submitting to remote jobs. The muted effects on the interaction terms between remote and URM/gender would suggest that conditional on jobs being remote, shifting the wage level does not have a large impact on the composition of the applicant pool.

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 gave many 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 in the category of intangible assets and knowledge-based), our findings on women and URM job applications spurred by the job design decision of making a job remote-eligible are notable. Our compensating wage differential analysis quantifies this preference.

One large issue these results raise is whether female and URM preferences for remote work are beneficial to the individual's productivity and ultimately, to the individual's promotion prospects at an organization. While the empirical results from the literature on the first area are unclear in that there is support for positive and negative effects (as discussed in our literature review), the evidence for dampened promotion prospects from less onsite work seems to be less controversial. As a result, from the standpoint of designing the workplace for an equitable future, follow-on work may examine the extent to which workplace policies and interventions, such as engineering social interactions for remote and hybrid remote 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.

An area for result interpretation in our study is the extent to which potential hiring firms in our data decide to wait out the COVID shutdowns and forestall their hiring activities.[8] 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 2019-2022 pandemic may still be yet to be revealed. Another fruitful domain of 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 particularly in competitive labor markets, employers may be more likely to offer such job flexibility. 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. Other such 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

Abraham, L. and Stein, A. (2020). Words matter: Experimental evidence from job applications. *Unpublished manuscript, Harvard University.*

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.

[8]If 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.

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.

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.

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.

Barrero, J. M., Bloom, N., and Davis, S. J. (2020). Why working from home will stick. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2020-174).

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*.

Bernstein, S., Townsend, R., and Xu, T. (2020). Flight to safety: How economic downturns affect talent flows to startups. *NBER Working Paper*, (w27907).

Blinder, A. S. et al. (2009). How many US jobs might be offshorable? *World Economics*, 10(2):41.

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.

Campero, S. (2021). Hiring and intra-occupational gender segregation in software engineering. *American Sociological Review*, 86(1):60–92.

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.

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.

Coffman, K. B., Collis, M., and Kulkarni, L. (2020). *When to Apply?* Harvard Business School.

Del Carpio, L. and Guadalupe, M. (2022). More women in tech? evidence from a field experiment addressing social identity. *Management Science*, 68(5):3196–3218.

Dill, K. (2021). Wework CEO says least engaged employees enjoy working from home. *Wall Street Journal*.

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.

Flory, J. A., Leibbrandt, A., and List, J. A. (2015). Do competitive workplaces deter female workers? a large-scale natural field experiment on job entry decisions. *The Review of Economic Studies*, 82(1):122–155.

Gneezy, U., Niederle, M., and Rustichini, A. (2003). Performance in competitive environments: Gender differences. *The Quarterly Journal of Economics*, 118(3):1049–1074.

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.

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.

Jensen, J. B., Kletzer, L. G., Bernstein, J., and Feenstra, R. C. (2005). Tradable services: Understanding the scope and impact of services offshoring [with comments and discussion]. In *Brookings trade forum*, pages 75–133. JSTOR.

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.

Murciano-Goroff, R. (2022). Missing women in tech: The labor market for highly skilled software engineers. *Management Science*, 68(5):3262–3281.

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.

Niederle, M. and Vesterlund, L. (2007). Do women shy away from competition? do men compete too much? *The Quarterly Journal of Economics*, 122(3):1067–1101.

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*.

Srikanth, K. and Puranam, P. (2011). Integrating distributed work: comparing task design, communication, and tacit coordination mechanisms. *Strategic Management Journal*, 32(8):849–875.

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*.

Yang, Y., Tian, T. Y., Woodruff, T. K., Jones, B. F., and Uzzi, B. (2022). Gender-diverse teams produce more novel and higher-impact scientific ideas. *Proceedings of the National Academy of Sciences*, 119(36):e2200841119.

Figure 1: Time series of jobs listed as remote

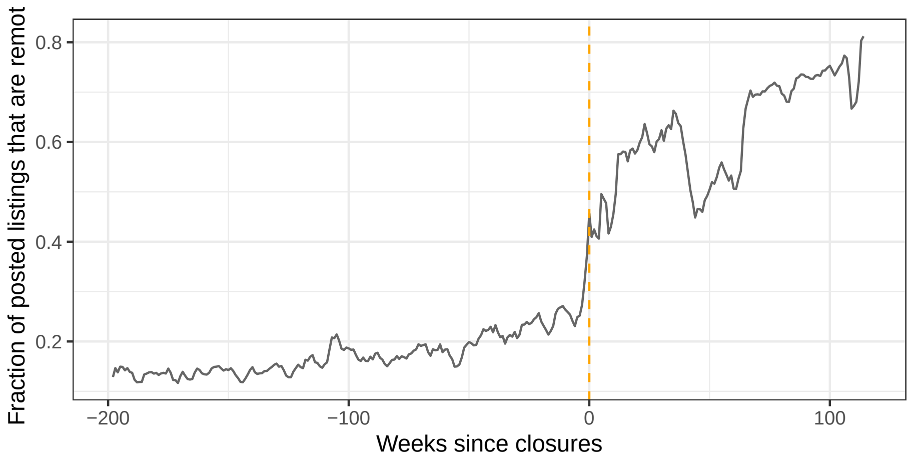

Figure Notes: This chart illustrates the fraction of full-time, US-based job listings in our sample that are tagged as remote-eligible. The x-axis denotes the number of weeks since the March 15th mandated COVID closures with the orange line indicating the date of closure. The y-axis is the 15 day rolling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

Figure 2: Employer size and remote-eligible jobs

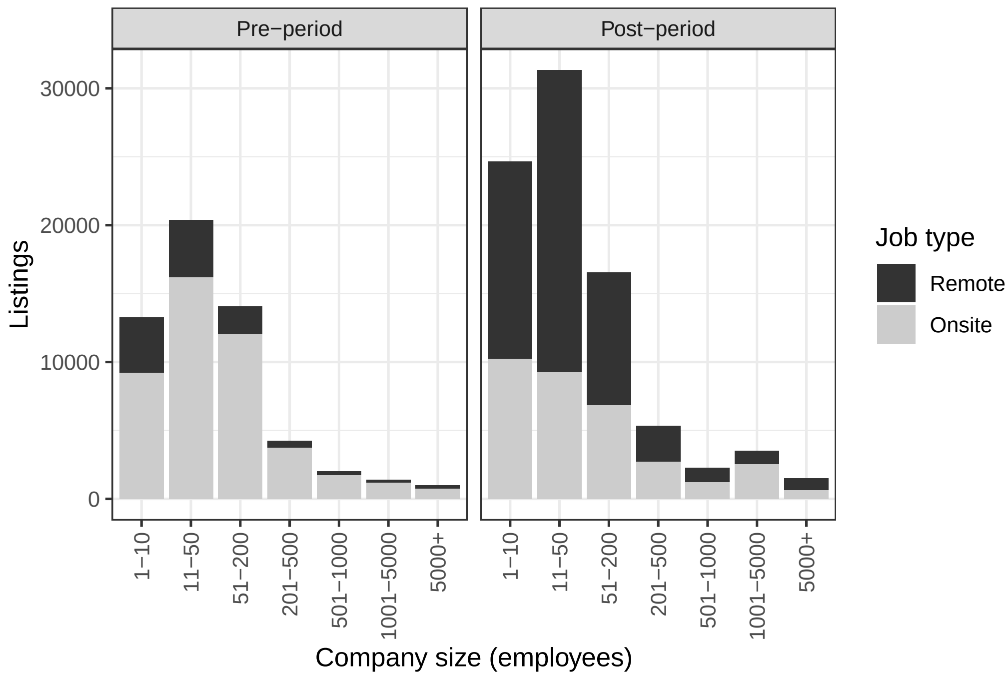

Figure notes: These charts report the distribution of full-time, US job listings by company size before and after the pandemic-induced shutdowns in March 2020. The pre-period includes dates from Jan 1, 2018 through March 15, 2020. The post-period are dates after March 15, 2020. The gray bars designate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

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

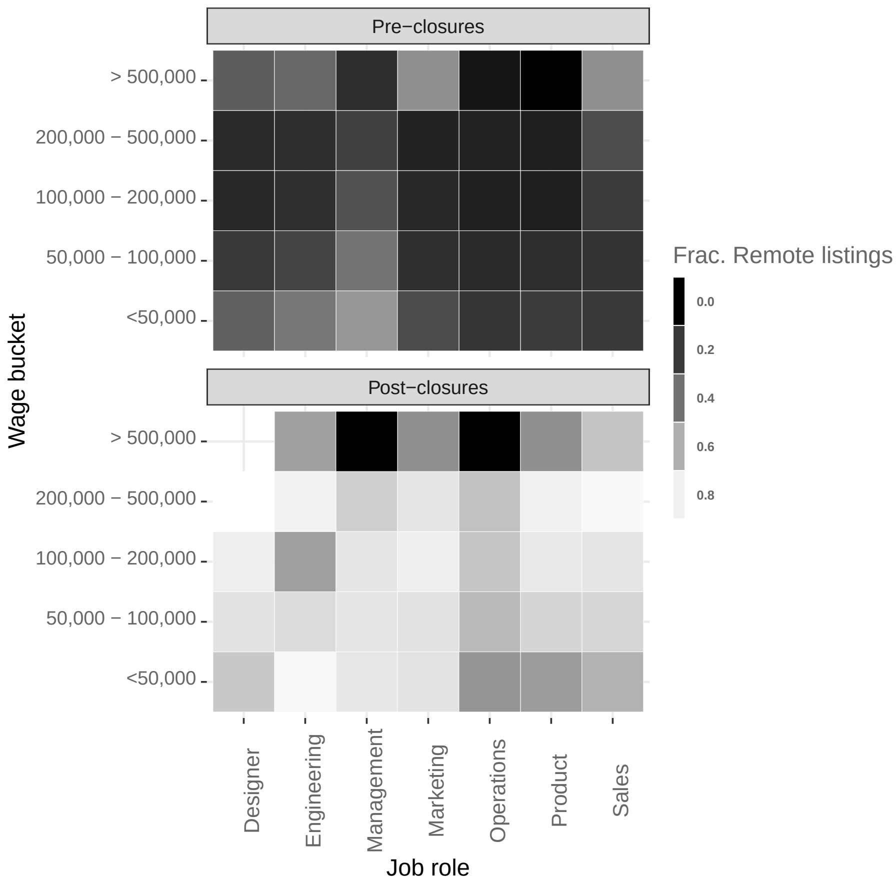

Figure Notes: These heat maps illustrate the fraction of full-time, US-based job listings that are designated as remote. Each cell illustrates the fraction of job listings that are remote by wage bracket and job role. The top panel includes job listings posted before or on March 15, 2020, and the bottom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

Table 1: Logit predictions of remote designation for job listings

Sample	2018-2019 Only		2018-2022	
	Coef	Std. Err.	Coef	Std. Err.
Model:	(1)		(2)	
Variables				
Employees (11-50)	-0.425^{***}	(0.054)	-0.396^{***}	(0.051)
Employees (51-200)	-0.824^{***}	(0.081)	-0.802^{***}	(0.078)
Employees (201-500)	-0.823^{***}	(0.110)	-0.810^{***}	(0.107)
Employees (501-1000)	-0.339^{**}	(0.170)	-0.326^{**}	(0.161)
Employees (1001-5000)	-0.843^{***}	(0.166)	-0.703^{***}	(0.161)
Employees (5000+)	0.412^{**}	(0.170)	0.497^{***}	(0.163)
Engineering role	0.034	(0.070)	0.045	(0.066)
Management role	-0.071	(0.107)	-0.071	(0.102)
Marketing role	-0.124^{*}	(0.074)	-0.107	(0.071)
Operations role	-0.381^{***}	(0.072)	-0.313^{***}	(0.069)
Product role	-0.532^{***}	(0.103)	-0.500^{***}	(0.096)
Sales role	0.037	(0.077)	0.035	(0.074)
Log(Salary)	-0.300^{***}	(0.016)	-0.285^{***}	(0.016)
Log(Equity)	0.487^{***}	(0.032)	0.464^{***}	(0.032)
Post × Employees (11-50)			0.418^{***}	(0.081)
Post × Employees (51-200)			0.409^{***}	(0.155)
Post × Employees (201-500)			0.166	(0.237)
Post × Employees (501-1000)			0.618^{*}	(0.369)
Post × Employees (1001-5000)			0.085	(0.447)
Post × Employees (5000+)			-0.450	(0.355)
Post			1.755^{***}	(0.160)
Post × Engineering role			-0.084	(0.106)
Post × Management role			-0.275^{*}	(0.160)
Post × Marketing role			0.092	(0.116)
Post × Operations role			-0.451^{***}	(0.116)
Post × Product role			0.433^{***}	(0.139)
Post × Sales role			-0.067	(0.122)
Post × Log(Salary)			0.151^{***}	(0.029)
Post × Log(Equity)			-0.367^{***}	(0.057)
Fit statistics				
Observations	36,900		59,421	
Squared Correlation	0.06792		0.31865	
Pseudo R^2	0.06166		0.25015	

Clustered (startup_id) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table notes: The models in both 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 uses listings posted from 2019 onwards and includes an indicator for listings posted after the closure date (*Post*).

Table 2: Characteristics of job listings and applicants

	Average	Std.Dev.	Max	Min
Number of apps	90.427	252.332	33,827	1
Number of female apps	16.980	50.726	9,572	0
Number of URM apps (when race is reported)	4.583	11.411	1,177	0
Mean experience in years	4.893	2.271	10	0
Salary	165.717	204.735	1,000	20
Equity	3.399	9.009	99.999	0
Post-closures	0.771	0.420	1	0
Remote	0.268	0.443	1	0

Table notes: This table reports statistics for job listings in the analysis sample. The first row reports statistics on the total number of applications sent to job listings in the sample. The second row reports statistics on the number of female applicants to each listing, from those applications where gender is reported. The third row reports statistics on the number of applicants to each listing that are submitted by URM candidates, from those applications where users report race (33% of the application sample). The fourth row reports statistics on the experience of applicants in years for job listings in the sample that received at least one application. The fifth and sixth rows report statistics on the salary (reported in thousands) and equity figures (reported in percentage points) posted with the job ads. The seventh row reports statistics on whether the job ad was posted after the pandemic-induced closures. The final row reports statistics about whether jobs in the sample were designated as remote.

Figure 4: Total applications sent to remote and onsite job listings

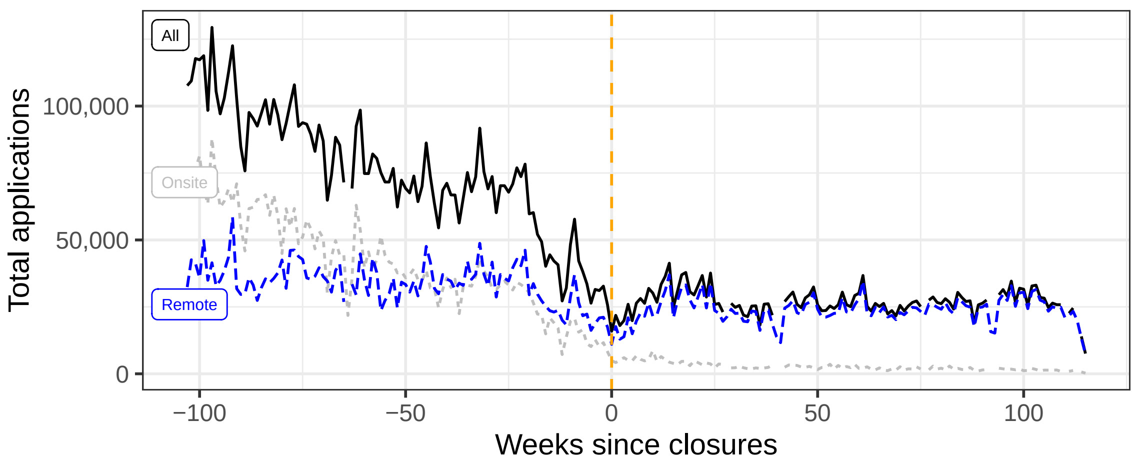

Figure Notes: This figure illustrates changes in numbers of applications to US, full-time job listings in our sample over the time window spanning two years before the pandemic-induced closures and about two years after the closures. The solid line is total number of applications to all job openings 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.

Figure 5: Attributes of applicants to remote and onsite jobs

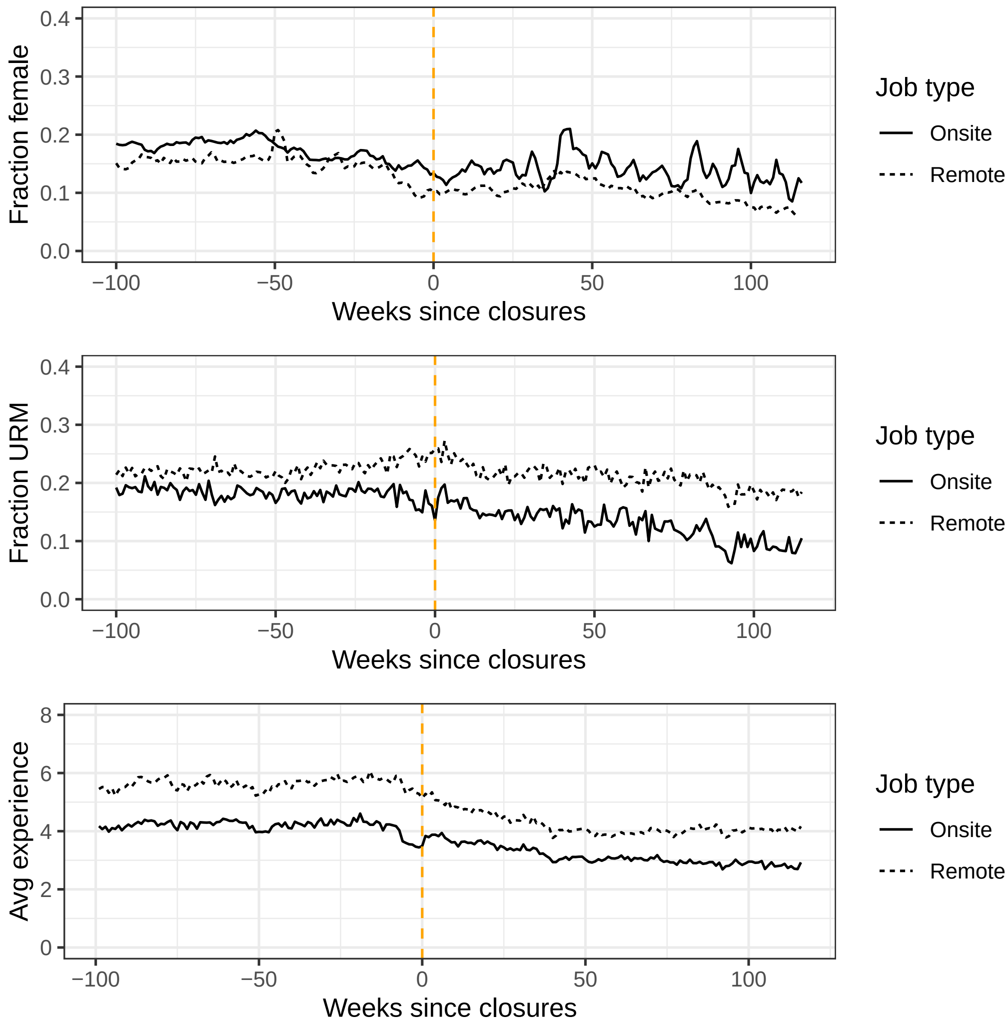

Figure Notes: These trend lines illustrate the time-series of the demographic composition of the labor pool to onsite and remote jobs. In each panel, the sample of jobs is limited to full-time, US listings. The top panel plots the fraction of female applicants to listings. The second panel is the fraction of URM applicants to listings, where the application sample is limited to cases where applicants report race. The third panel is the average industry experience in years of applicants to listings. In all panels, the x-axis represents days since closures and the orange line is March 15, 2020, the date of the pandemic-induced closures. In all panels, the solid line depicts the trend line for listings designated as onsite jobs, and the dashed line depicts the trend line for listings designated as remote jobs.

Figure 6: OLS and Matched sample estimates of remote coefficient

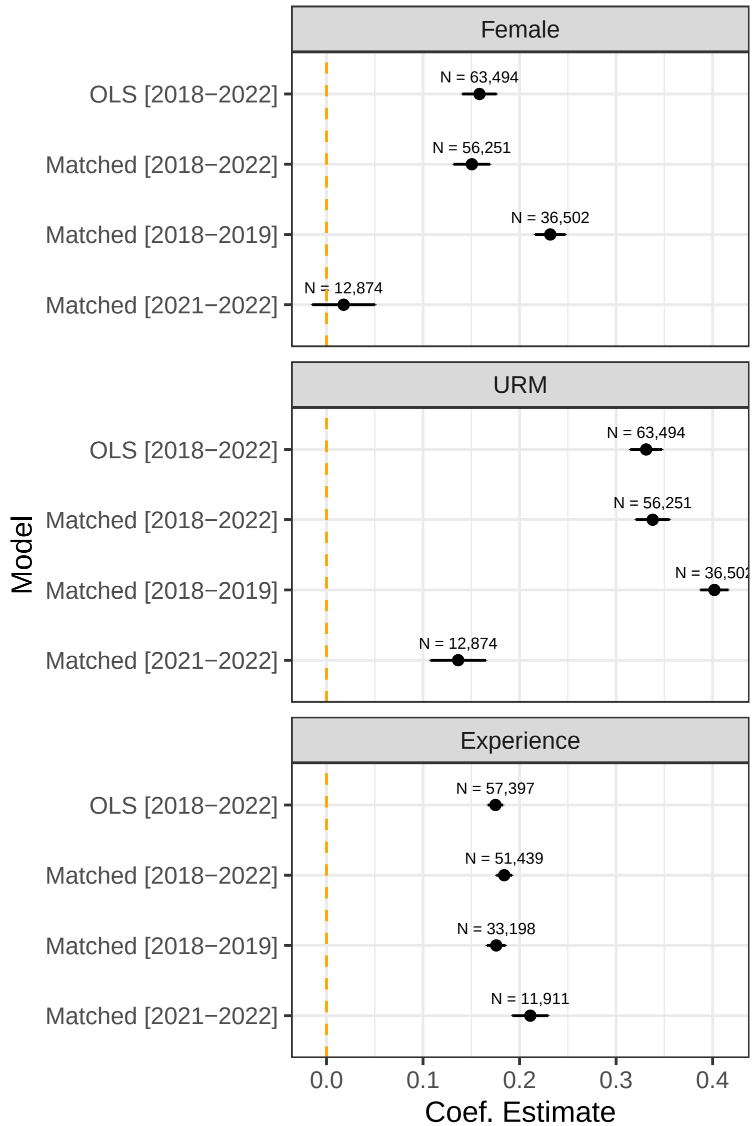

Figure Notes: These figures illustrate the results of analyses from OLS and matching tests. In all specifications, the sample is limited to US, full-time job listings from 2018-2022. The point estimate that is shown is the coefficient estimate on the *remote* indicator in the following specification: $y_i = remote_i + log(salary)_i + log(equity)_i + log(apps)_i + \gamma_i + \epsilon_i$. Here, y is the measure shown in the panel header, γ is a vector of fixed effects which includes 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 standard error bars on each point estimate indicate a 95% confidence interval. Standard errors are clustered on week of posting.

Figure 7: Distributed lag model estimates of remote effects

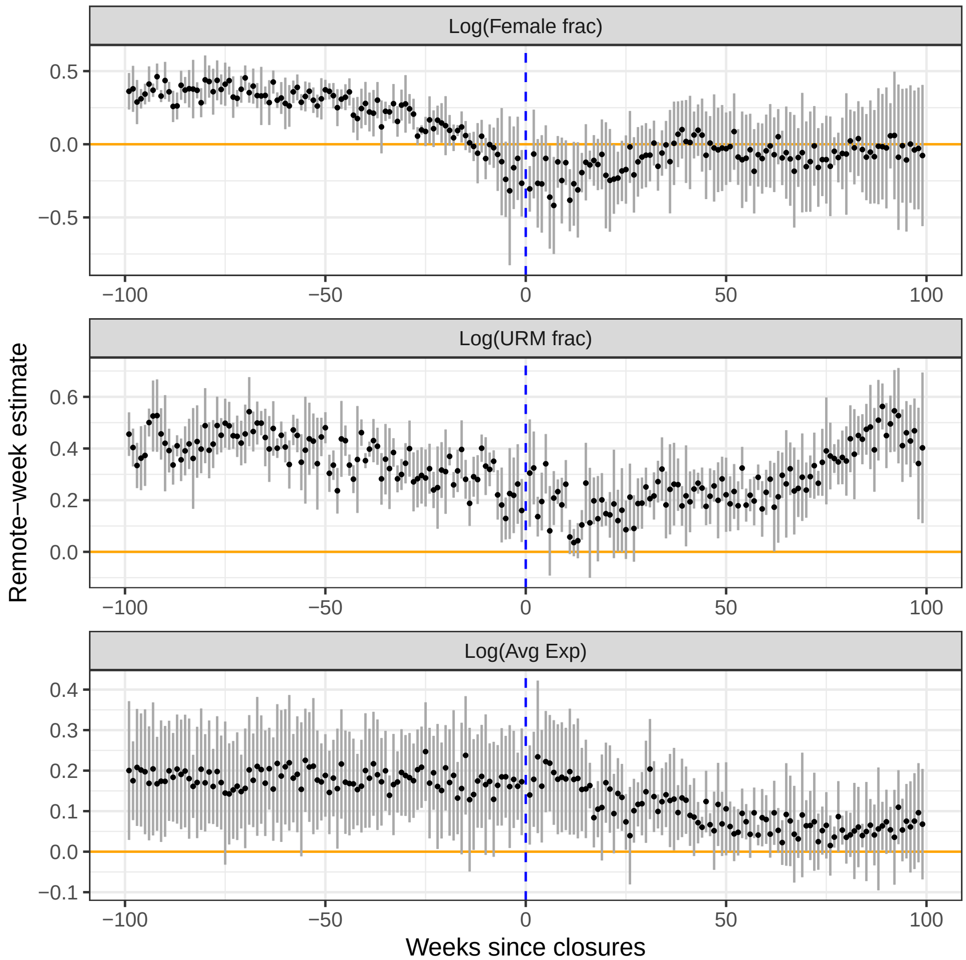

Figure Notes: This figure illustrates the coefficient estimate on $remote \times week$ from the following specification: $y_i = remote_i \times week_i + log(salary)_i + log(equity)_i + log(apps)_i + \gamma_i + \epsilon_i$ where $week_i$ is the number of weeks since the pandemic closures and γ_i includes 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 8: Remote effect by quantile of local market diversity

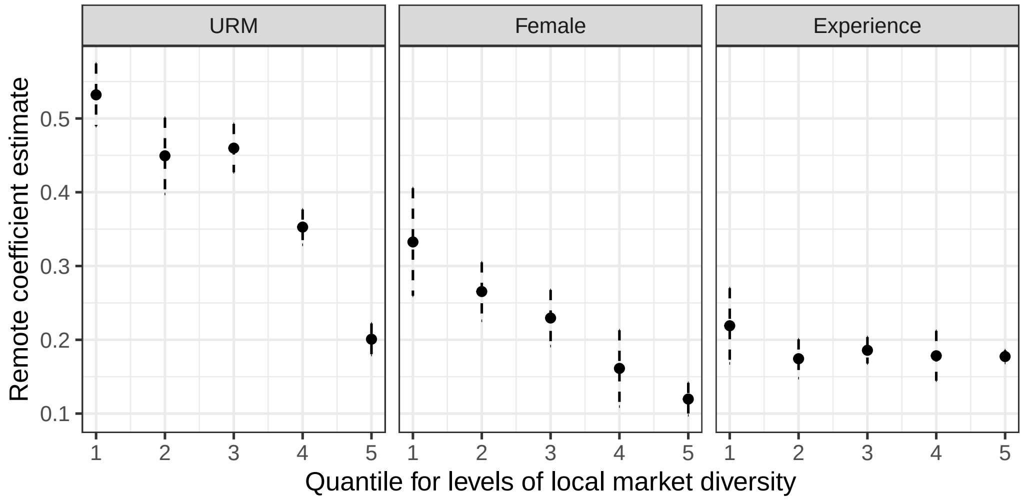

Figure Notes: This chart indicates the *remote* coefficient estimate from the specification used in Figure 6 but with the sample divided into five quantiles according to levels of diversity in the local market, proxied by the demographic content of applications to onsite jobs. For this analysis, we only 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 9: Top URM donor cities for least diverse local markets

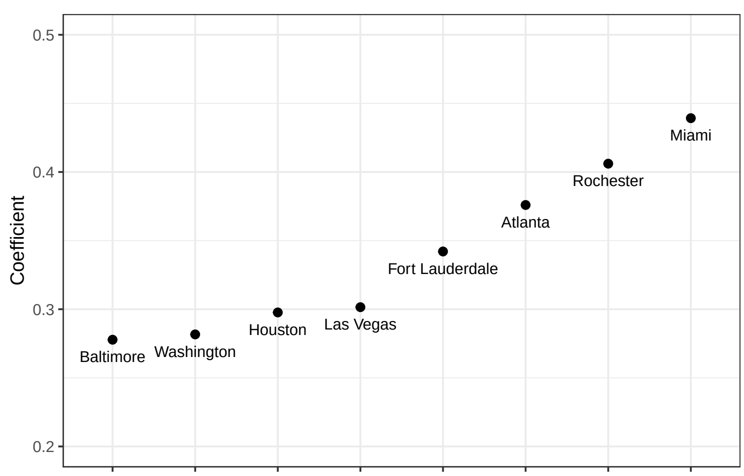

Figure Notes: This figure plots the eight largest city fixed-effects from the regression $urm_i = jobtype_i + city_i + \epsilon_i$ on a sample of all applications to remote jobs posted in the least diverse cities (Quantile 1) from Figure 8. The sample is also restricted to applications that originate 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. Variable	Remote?	Log(Salary)		
Family	*Logit*		*OLS*	
Model:	(1)	(2)	(3)	(4)
Variables				
URM	0.162***	−0.145***	−0.106***	−0.036***
	(0.025)	(0.020)	(0.012)	(0.007)
Remote × URM		0.025	0.040**	−0.016***
		(0.017)	(0.017)	(0.004)
Female	0.571***	−0.210***	−0.082**	−0.043***
	(0.054)	(0.035)	(0.027)	(0.006)
Remote × Female		−0.157***	−0.098***	−0.010
		(0.031)	(0.022)	(0.008)
Post-closure	0.252***			
	(0.035)			
Remote		−0.064***	−0.068***	−0.075***
		(0.015)	(0.018)	(0.019)
Fixed-effects				
Year posted				Yes
Job role	Yes		Yes	Yes
Years experience	Yes	Yes	Yes	
Fit statistics				
Observations	50,594	50,594	50,594	3,328,354
R^2		0.22860	0.36811	0.16689
Pseudo R^2	0.03950	0.17707	0.31318	0.17523

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

Table notes: This table analyzes preferences for remote work and remote-work based compensating differentials based on data directly 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 and third columns are regressions of applicant characteristics and remote preferences on the log of their self-reported salary expectations. The second column only includes fixed-effects for experience levels, and the third column includes fixed-effects for experience levels and job role. The fourth column estimates how the salaries posted in job postings that applicants choose to apply to differ when the jobs are designated as remote.

Appendix A Evaluation of match quality

This section presents results that assess balance in our matched sample. Figure A.2 shows the standardized mean differences of 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

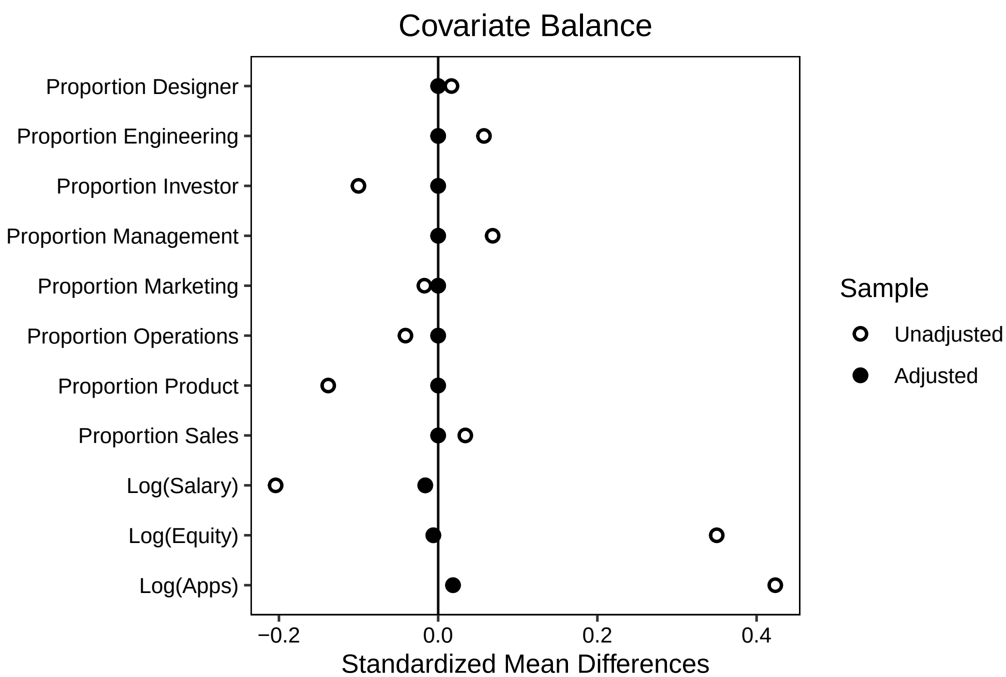

Figure notes: This figure illustrates standardized mean differences between the matched remote and onsite listings sample (Adjusted, filled circles), and for the raw sample (Unadjusted, hollow circles).

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 very similar for the matched sample.

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

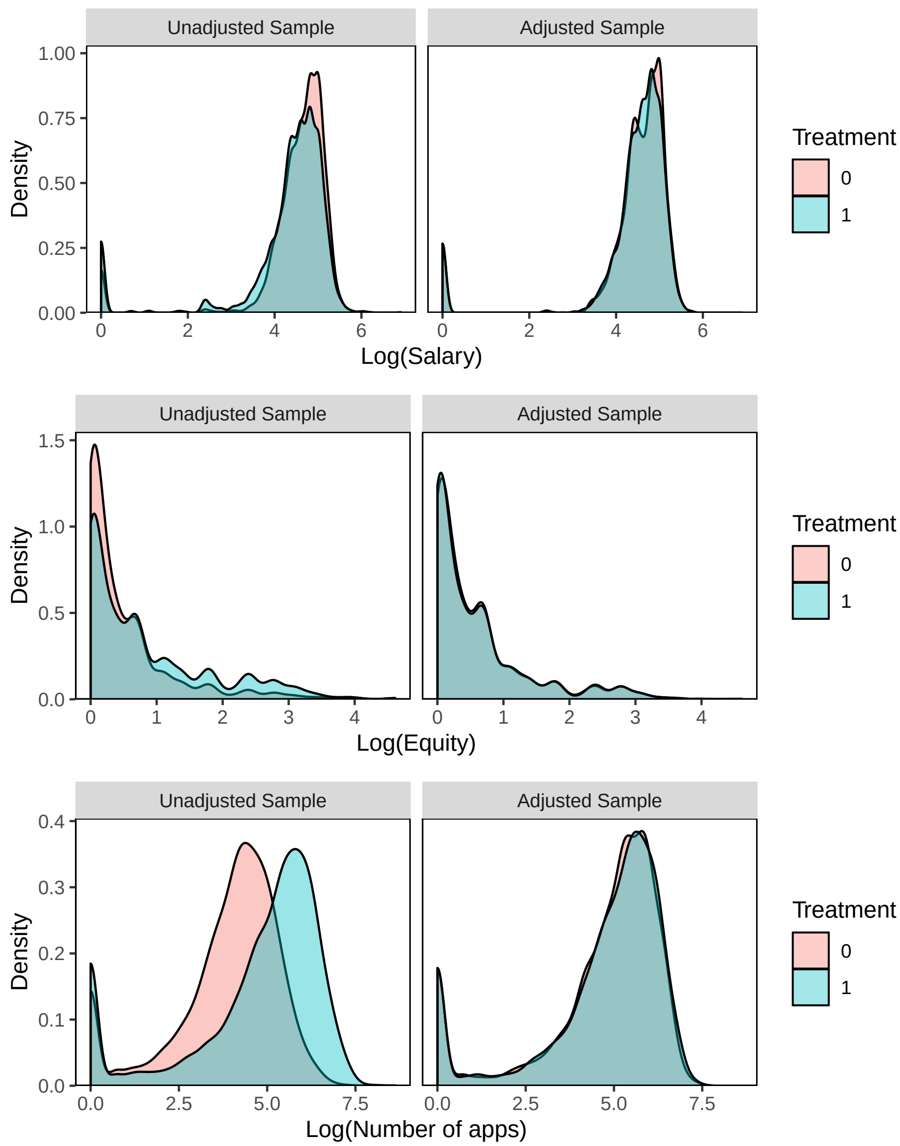

Figure notes: This figure illustrates the distribution of salary, equity, and number of applications for the remote sample, the matched onsite sample, and the complete onsite donor sample.

Appendix B Robustness tests on main specification

In this section, we present the results of additional robustness tests. Our baseline regression, presented in the main body of the paper, includes fixed-effects for job type (at a broad level of the job taxonomy, created by the data provider) 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 also test a specification that uses fixed-effects for detailed job title rather than broader job category.

B.1 Additional fixed-effects

First, we include fixed-effects for location (city) and employer. We first limit the sample to the top 1,000 cities and top 1,000 employers, in terms of the number of applications that appear 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 of these specifications, the coefficient estimate on the *remote* variable remains 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 Fixed-effects for specific job title

The specifications that have been used to this point use fixed-effects for broad job roles, which comes from a taxonomy specified by the job provider. In Figure B.2, we substitute fixed-effects for the fine-grained job titles included in the job listings. Due to a long tail of job listings in this sample, we include fixed-effects for specific job titles (e.g. Marketing Analyst, Software Developer IV) only after limiting the sample to listings that fall into one of the 1,000 most common job titles reported in these postings.

We replicate our baseline specification on a matched sample produced 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. Substituting fine-grained job title-based fixed effects in the limited sample does not impact the sign or significance of our *remote* coefficient, although the magnitude of the estimates are smaller than that of the estimated coefficients in Figure 6, which suggests that some, but not all, of the demographic differences in remote jobs can be attributed to heterogeneity in job title. 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 the experience of the application pool.

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

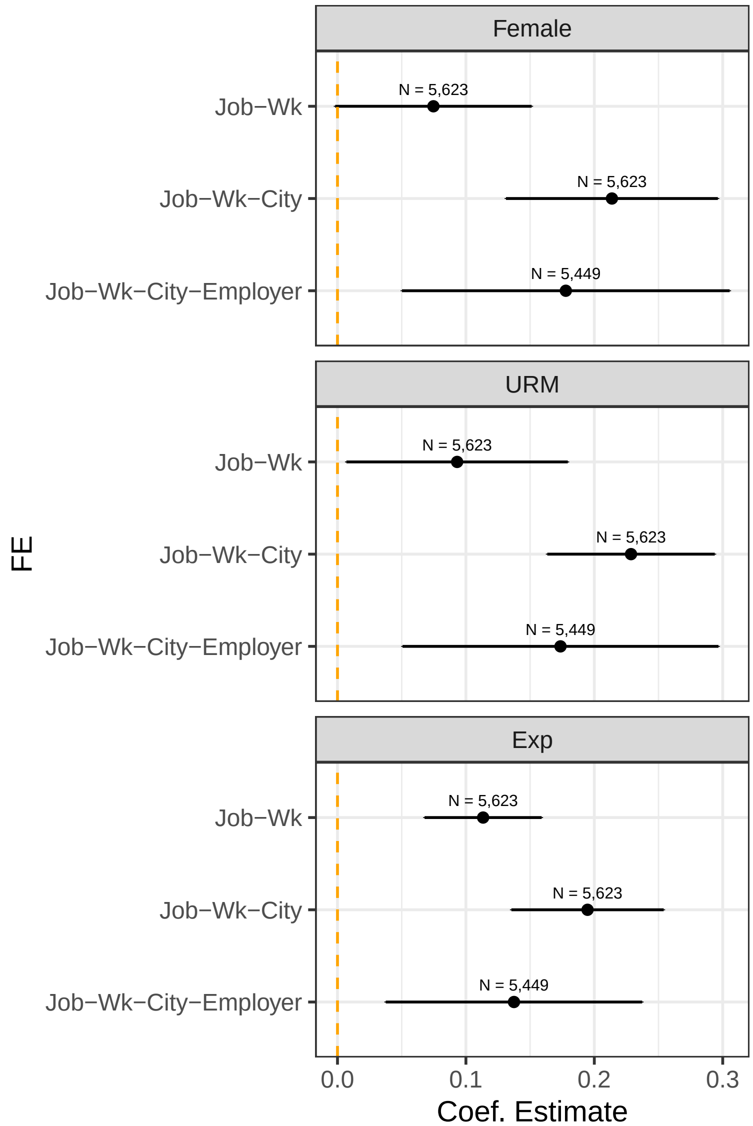

Figure notes: These charts show the remote coefficient when using the specification from Figure 6, except that we add additional fixed-effects. We limit the sample to postings located in the top 1,000 cities and top 1,000 employers in terms of number of listings posted in the city or by the employer, 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: Finer grained job title controls on matched sample

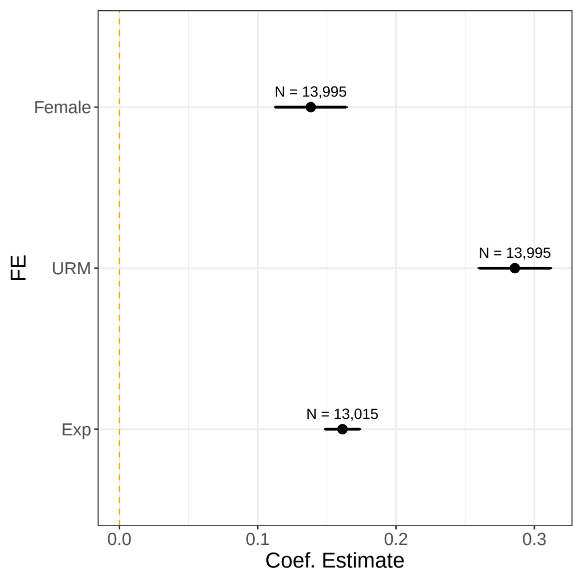

Figure notes: This chart illustrates the estimate on the remote coefficient using the same specification as in Figure 6, but instead of broad job categories, we substitute fixed-effects for specific job titles. The sample is limited to postings where the posted job title is one of the 1,000 most frequently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

Appendix C Alternative estimation method

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).[9] In the main body of the paper, we use a matching approach to address these confounds. Here, we discuss an alternative approach.

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 data 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 shifted to 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 the characteristics of an otherwise onsite job fixed, designating it as remote increases the diversity and experience of the application pool.

[9]To 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	Log(Female)	Log(URM)	Log(Exp)
Model:	(1)	(2)	(3)
Variables			
Remote	0.101***	0.209***	0.184***
	(0.031)	(0.027)	(0.016)
Log(Salary)	−0.216***	−0.166***	0.372***
	(0.034)	(0.038)	(0.034)
Log(Equity)	−0.094***	−0.059**	0.073***
	(0.026)	(0.027)	(0.011)
Log(Num apps)	0.775***	0.744***	−0.079***
	(0.013)	(0.011)	(0.006)
Fixed-effects			
Week of posting	Yes	Yes	Yes
Job role	Yes	Yes	Yes
Fit statistics			
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 $y_i = Remote_i + Log(Salary) + Log(Equity) + Log(Apps) + \gamma_i + \epsilon_i$ 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.