Startup labor markets and remote work: Evidence from job applications

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

July 26, 2021

Abstract

Does offering remote work allow startup firms to attract more experienced and more diverse (gender and race) talent? We examine job listings and job applicant behavior on a leading platform in this space, AngelList Talent, amid the COVID-19 pandemic-induced shutdowns. We first characterize the jobs and organizations offering remote work before the shutdowns. We then leverage the context to help address the empirical confound of job design (including offering remote jobs) as co-determined with unobserved job and firm characteristics. By doing so, we estimate the change in applicant characteristics to job postings which are (exogenously) shifted to being remote. This design is a window into evaluating a managerial choice (offering remote work) which will likely become more salient in post-pandemic job design. We find that offering remote-eligible work attracts more experienced and diverse job applicants.

Keywords: remote work; startup labor market; talent acquisition; diversity and inclusion; COVID-19; technical workforce.

*Correspondence to: dhsu@wharton.upenn.edu & tambe@wharton.upenn.edu. We are grateful for feedback from Matthew Bidwell, Peter Cappelli, Raj Choudhury, Vilma Todri, and seminar audiences at the Asian Innovation and Entrepreneurship Association (AIEA), Boston College, CU Boulder, UC Berkeley, George Mason University, the Workshop for Information Systems and Economics, INSEAD, Tulane, and USC. We thank Kevin Laws at AngelList for data access. 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. A much wider share of the world experienced remote work in 2020-21, with the COVID-19 pandemic due to government-mandated shutdowns. 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 that they want employees back in the office as soon as possible after widespread vaccine access. Other organizations, such as Facebook and Google, have taken a hybrid stance, allowing some fraction of their workers the ability to conduct all of their work remotely and/or have the flexibility to 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 mostly (though not exclusively, as we detail in a below literature review) find a positive individual productivity effect caused by the remote work condition within the subject organizations. Left open, however, are two issues related to the labor *market*, which in turn inform our research questions in this paper: (1) what organizational and job characteristics shape the likelihood a given job posting is listed as "remote"?, and (2) do employers attract more experienced and more diverse (gender and race) talent when they designate a job as remote-eligible?

To study these questions, we analyze a new data set, from AngelList Talent, that covers activity from both sides of the startup labor market, 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 online 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, such as skill requirements associated with each job. We find and characterize the empirical patterns which are associated with the likelihood a given job is listed as remote-eligible. This sets up the second analysis we conduct to address the 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 improve causal inference. We do so by analyzing applicant behavior in the midst of the COVID-19 pandemic shutdowns, which forced a quick transition to remote work conditional on hiring aspirations (and therefore made this feature of job design less of a managerial choice). This allows 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. We also present (in the appendix) two separate empirical strategies based on sample matching and contrasting automatic and manually posted job listings to triangulate our results. Across the estimation approaches, we broadly find that offering remote work affords organizations more applicants (and with more experience), as well as applicants drawn from more diverse gender and race backgrounds.

2 Literature

Two streams of prior work relate to our research question. While both relate to the shifting organization of work, the first domain takes the employer perspective about considerations to offering remote work. The second domain takes the employee perspective and examines drivers of their preference for/against remote work. 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[1], 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.

With the caveat that the field experimental evidence is drawn from only a handful of organiza-

[1]The literature pre-dates the COVID-19 pandemic. For example, see the work on "telecommuting" (e.g., Dutcher [2012]; Gajendran and Harrison [2007]), geographically-dispersed teams (Hinds and Mortensen [2005]; Gibson and Gibbs [2006]), and more generally, alternative job arrangements (e.g., Mas and Pallais [2020])

tions, 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), Bojinov 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. 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 prior 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: 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. 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 on-site 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. Women also face higher costs of commuting [Le Barbanchon et al., 2021]. A second theme is that infrastructure preparedness, both with regard to space in the home as well as computing resources, including access to a stable broadband internet connection, is an important precondition to remote work. Indeed, Barrero et al. [2020] argue that one of the reasons WFH is likely to persist in the long term is because of the sunk infrastructure arrangements many made in these resources during the course of the COVID-19 pandemic. A third theme is employee psychological and social considerations (such as isolation, mental health, and well-being) associated with remote work (e.g., [Gajendran and Harrison, 2007]). The total effect of these considerations is not straightforward, though, because remote work can act asymmetrically on employees' relations with their family as compared to their co-workers, and so intersects with work-family dynamics (and the unbalanced impact across genders) [Rothbard, 2001].

On balance, employees 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.

An open set of issues relate to labor market dynamics (attracting talent, especially with racial and gender diversity) aligned with the managerial decision of making jobs remote-work eligible. There are two empirical difficulties of addressing this managerial problem. First, there are a number of unobserved and unmeasured factors on both sides of the marketplace identified in the literature 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*.

3 Data

In this section, we describe the data, variables, and empirical strategy we use to examine the correlates and consequences of offering remote 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.[2] Importantly for our study, prospective employers on the platform are asked to designate whether the job being posted is "remote-eligible", 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 do not have direct information through the platform of any 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, our access to daily site activity from job *applicants* runs from February 5 to June 18, 2020, and so we use this time window for our job listings analysis (though importantly, when we analyze the changes to applicant quality induced by making a job posting remote-eligible, we use pre-2020 job listings to "train" our *remote* prediction model). These dates book-end the adoption of employers' pandemic policies. 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, skills required, and compensation details (wages & equity) for a subset of the jobs. From the applicant data, we have information on users' platform activity (e.g., searches, clicks, applications) such as listings viewed and applied to. To engage on the platform, job candidates had previously filled in their profile information (such as resume data, including educational history, skills, prior employers and LinkedIn page). 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

[2] 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. [2020]), while the economic situation for all small businesses on average characterized as much more precarious and financially fragile (e.g., [Bartik et al., 2020a]). This mirrors the findings of Dingel and Neiman [2020] that while approximately 37% of jobs in the US can be conducted remotely (which favorably compares to other countries), the comparable figures for jobs in management and computing in the US (those closest to the jobs on the AngelList Talent platform) are greater than 80%.

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 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 would like to characterize the types of organizations listing on the AngelList Talent platform (not just in our sample window). As an overview, consider the following information about the top four employers in the overall data: SpaceX (849 job listings), Slack (452), Twilio (374), and DoorDash (372). In total, of the approximately 62,000 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.[3]

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: (1) factors shaping the likelihood a job is listed as remote-eligible (before the COVID-related shutdowns in the US, which we take to be March 13, 2020 in accordance with U.S. federal government designation of a national emergency and associated quarantine orders), and (2) how WFH job status influences job applicant characteristics to job listings (number of applications, average experience, and gender & racial diversity).

[3]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%).

3.2.1 Description of remote listings

In our job listings dataset, the top four job titles are: software engineer (321), senior software engineer (294), product manager (233), and sales development representative (177), while the top four skills most requested in job postings are: *javascript* (1366), *python* (1273), *sales* (1204), and *react.js* (1141). More generally, Figure 1 provides information about the distribution of job postings by employee headcount (left panel) and by highest attained funding round (right panel), both disaggregating on-site (denoted in light blue-colored bars) and remote-eligible jobs (orange-colored bars). Taken as a whole, smaller and earlier funding round stage companies are the ones most actively listing jobs, and these type of firms are also the ones disproportionately listing their jobs as WFH.

We begin by describing factors associated with remote-eligible jobs.[4] In Table 1, we present descriptive organizational characteristics, separated by remote-eligible as compared to on-site jobs. A final column reports the difference in the average values of the conditional means. We see that the difference between on-site and remote-eligible jobs increases with organizational headcount and attained financing milestones (note that while the percentage of remote-eligible job listings for pre-seed financing stage firms is high compared to other attained financing rounds, the absolute number of listings in this data bucket is relatively small) and is slightly more prevalent outside of Silicon Valley.[5]

In addition to organizational attributes correlated with the likelihood that a job listing is remote-eligible, the appendix also contains job title and skill-level factors as well (together with a more extensive discussion).

3.2.2 Variables and empirical design examining the consequences of remote listings

On the applicant side, the descriptive pattern of fraction of applications sent to remote-eligible jobs pre- and post-shutdown is dramatic. Figure 2 demonstrates that before the shutdowns, the fraction of applications sent to remote job listings as a share of all applications is essentially flat and

[4]To give a broader context using the AngelList Talent job-posting data since its inception in 2011, the time series of remote-eligible jobs peaks in 2013 and 2017, when over 30% of the jobs are designated WFH. Since 2017, the fraction of remote-eligible jobs has been declining. Of course, the absolute number of jobs listed on the platform has been increasing over time as well.

[5]The Silicon Valley designation includes the following cities: San Francisco, San Jose, Fremont, Santa Clara, Sunnyvale, Oakland, Berkeley, Santa Monica, and San Mateo.

stable, at approximately 50%. After the shutdowns, the corresponding percentage is rising with the elapsed weeks since the shutdowns, and ending at approximately 90% by the end of our observation window.

For the second analysis, we examine how a host of job applicant characteristics relate to (and are induced by) the managerial decision to make jobs remote-eligible. For applicant characteristics analyzed and reported in this paper, we use a sub-sample drawn from all of the platform applicant activity data, so figures such as number of applications reflect statistics for the sample of applicant activity, not total platform statistics. In this analysis, our outcome variables center on applicant reactions to job postings, and span the following: (1) *application count* is the count of number of applications sent to a job listing (mean in sample=3.1 for jobs receiving at least one job application); (2) *applicant experience* is the number of years of work experience in the current job role among applicants to a job (an average of 4.4 years); (3) *number of female applicants* is the average number of applications sent by female applicants to job postings (average is 0.7), which we infer using the popular *gender-guesser* package that infers gender from name; and (4) *number of URM applicants* is the average number of applicants to a given job which are likely to belong to an underrepresented minority group (mean=0.3), using the *wru* package in R (see Imai and Khanna [2016] for applications of this package in social science research). Table 2 contains the summary statistics for these variables. The explanatory variables are *post-closure*, an indicator for after the COVID-induced shutdowns (mean=0.81); and *remote* is a dummy variable for whether a job is listed as remote-eligible (mean=0.21).[6]

To characterize the evolving job listing and application landscape before and after the shutdowns, consider the following figures. Figure 3 shows the dramatic shift in the fraction of new job listings designated remote-eligible by week before and after the shutdowns (designated by the vertical red line) in the AngelList Talent job listings. The figure indicates volatility around the middle weeks of March 2020. As the transition was occurring to the pandemic economy, employers were perhaps attempting to fill local positions and may have refrained from posting. Shortly thereafter, given what we know ex-post about the relentless demand for technical and technical-adjacent talent even during the height of the pandemic, the fraction of remote listings settles higher, this time with

[6] We are unable to strictly distinguish between "remote-only" jobs as compared to a hybrid arrangement of "partially remote" with the balance of a job being on-site. We therefore adopt the term "remote-eligible" throughout. We nevertheless believe that the job design contrast with "on-site only" is meaningful.

less volatility. By May, listings experienced an unmistakable shift towards remote listings.

To address the endogenous nature of remote-eligible job listings, our principle empirical strategy is to leverage a natural experiment afforded by the COVID-19 discontinuity in the workplace and its implications for hiring. We also offer two alternative empirical strategies (contained in the appendix) to improve our inference. Since we will explain details of the operationalization below, we give only a high-level overview of our approaches here.

For our natural experiment, we leverage the exogenous COVID-19 government-mandated shutdowns in which the decision to offer remote work was largely taken away from managers. We employ models to predict jobs that are on-site in the pre-shutdown regime and use them to identify job listings we would have predicted to also be on-site in the post-COVID regime, but are instead remote as a result of the COVID-19 shutdowns. In this manner, we estimate the change in applicant characteristics (quantity, experience, and diversity) stemming from the exogenous switch to remote work.

As independent and alternative approaches to triangulate our results, we first implement a (coarsened exact) matching (CEM) procedure in which we balance the "treated" (remote) observations with the "control" observations on (exact) job title. This matching approach helps improve inference attributed to the *remote* job effect on applicant characteristics. A second approach exploits a feature of the digital data from AngelList talent: we are able to detect whether job postings are "automatically" re-posted as compared to "manually" listed. Under the assumption that auto-posted jobs involve a different managerial decision-making process as compared to those manually posted, we estimate how manually listing jobs as remote-eligible after the shut-down compares to a range of alternative configurations. As a result preview, we find directionally similar statistical and economic effects using all three approaches.

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. Consider Figure 4, in which we plot the mean number of applications, their quality (measured by applicants' industry experience), and diversity (number of female applicants and number of URM applicants) over time (again, the dashed vertical line denotes the closure boundary). The mean for on-site job listings is marked in a solid line, while the remote job listing mean is marked by a dotted line throughout the figure.

With regard to applications, at the aggregate level, we see that while applications sent to on-site job listings are higher than those sent to remote jobs in the weeks leading up to the shutdowns, this pattern reverses in the post-shutdown period. The level of applications being sent to job listings in the post-shutdown period diminishes relative to the pre-period. The difference in the applicant industry experience variable 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). Finally, on the diversity front, while the number of female applicants to job listings appears to be somewhat higher for on-site as compared to remote jobs before the closures, that pattern appears to reverse thereafter, with higher female applicants to remote-eligible jobs (though the average number of female applicants to jobs appears lower after closures). For the number of URM-status applicants, the time trend between on-site and remote job listings appears comparable until a few weeks ahead of closures, after which it appears that remote listings are associated with higher numbers of URM applicants per job. 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 in an OLS regression framework in which we do not yet tackle the issue of endogenous remote job listings. In Table 3, we examine the correlates of applicant characteristics received to job listings, where the sample is confined to pre-closure observations only (recognizing that the post-closure time period likely represents a regime shift). In each specification, we include fixed effects for startup, job title, location, and week. Across the five applicant composition outcomes, we see that *remote* is positively related with all of the outcomes, and with large estimated economic effects (for example, a discrete shift to remote-eligible jobs is associated with a 10.7% increase in URM applications).

4.2 COVID-19 shutdowns as a natural experiment

The analysis in Table 3 carries the implicit assumption that remote listings are exogenously determined. However, a significant confound to this interpretation is that remote listings are more likely shaped by factors associated with the firm and the job position itself (in ways which are unmeasured and/or co-determined).[7] In order to move beyond correlating applicant characteristics with offering remote work, we now discuss our natural experimental approach.

The unanticipated COVID-19 shutdowns afford us a quasi-experimental empirical strategy to address the endogeneity of remote job listings. Our goal is to identify "treated" listings that are made remote by employers, but absent work closures, would have otherwise been listed as on-site. Because many listings are remote in the post-closure data, it is useful to classify those that, in a normal business environment, would have been on-site listings. To do this, we build a prediction algorithm that, using pre-closure data, classifies jobs as being remote or on-site given the characteristics of the job. We then use it to identify those remote jobs in the post-closure period that – had business conditions not shifted towards remote work – we would have otherwise expected to be on-site.

For the prediction task, we use a supervised machine learning classifier that performs particularly well on high-dimensional data, a Random Forest algorithm [Breiman, 2001].[8] Table 4 reports classification accuracy when identifying whether a listing is likely to be listed as on-site or remote-eligible, given features of the listing data, which include: job titles, skills listed in jobs, time-of-year, location, employer, and various employer attributes including headcount and funding status. We use 70% of the jobs listed in 2019 (N=7,150) as our training sample (panel A). The model's accuracy on the training sample is 96.3%, reflecting 65.4% of the sample classified as onsite when the job is listed as onsite and 30.9% of the sample classified as remote when a listing is remote. To validate the performance of this prediction model on out-of-sample data, we use the remaining 30% of jobs posted in the year 2019 (N=3,016). The accuracy rate of prediction drops to 82.2% in classification (shown in Panel B, corresponding to 60.9% correct onsite predictions and 21.3%

[7] 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.

[8] While we use a Random Forest algorithm for predicting remote-eligible jobs, it is possible to use other supervised learning methods, include logistic regression, or even simpler qualitative dependent variable models which use heuristic methods to identify jobs likely to be listed as onsite or remote. Doing so yields qualitatively similar results, but our main specifications use Random Forest model to classify listings because these models are best suited to the higher dimensionality in our data.

correct remote predictions). About 9% of jobs are incorrectly classified as being onsite when they are listed as remote.

In panel C, using the prediction model trained on the 2019 data, we identify jobs in the 2020 COVID-19 closure period that, based on their characteristics, we would have predicted to be onsite under normal business conditions. Here, our accuracy drops significantly, as expected given the shift in the business environment, to 53.5%. The listings in the bottom left quadrant, comprising 38.3% of listings, form the subject of our natural experiment, subject to a degree of measurement error on the order of 9% as suggested by Panel B.

In Table 5, based on the use of the Random Forest algorithm to classify the remote-eligible jobs as those that always would have been remote and those that were made remote due to the office closures, we categorize post-closure jobs into one of four groups: (1) predicted remote & actually remote; (2) predicted remote & actually onsite; (3) predicted onsite & actually onsite; and (4) predicted onsite & actually remote. We define the "treated" group as category (4), which are job-listings which are predicted to be onsite, but because of the COVID shutdowns, are instead observed as remote. The other groups serve as "controls". In Table 5, the baseline category is (3), and the focal coefficients of interest ("treated") are in the first row. All columns include fixed-effects for startup, location, job title, and week, with the exception of column (2) which omits startup fixed effects because we have fewer observations for job listings with applications where we can compute applicant experience measures. The *remote* coefficient is positive and significant in all of the columns, and with economic magnitudes which are somewhat lower than those estimated when treating remote-eligible as exogenous (in the case of the effect on female and URM applications), though not uniformly so (for example, predicted applicant experience is substantially higher in the natural experiment). In the natural experiment, being in the "treatment" group is predicted to lead to a 5.3% increase in female applications and a 9.6% boost in URM applications as compared to the baseline group (predicted onsite, actually onsite).

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 quantity, experience and diversity, should employers elect to make jobs remote-eligible.

Of course, our empirical work cannot capture the full dynamic equilibrium of labor market adjustments on both sides, which may occur in the aftermath of the 2019-2021 pandemic. An additional inference issue is the extent to which potentially hiring firms decide to wait out the COVID shutdowns and forestall their hiring activities.[9] We also do not measure productivity or other worker-level output measures. Ideally, future work in this domain would compare the quality of talent recruited as well as the productivity and creativity of employees. Doing so would provide greater insight into the potential broader impact of the shifting applicant behavioral patterns we document. 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. We suspect that the concept of "jobs which can be effectively performed remotely" is not exogenously-given and fixed. 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. As a result, we hope that the work presented here is a window to understanding remote work and the startup labor market, but recognize there is much work to be done to understand this aspect of the future of organizing work.

[9] 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.

References

Atkin, D., Schoar, A., and Shinde, S. (2020). Worker sorting, work discipline and development. Technical report, MIT Sloan working paper.

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. (2020a). The impact of COVID-19 on small business outcomes and expectations. *Proceedings of the National Academy of Sciences*, 117(30):17656–17666.

Bartik, A. W., Cullen, Z. B., Glaeser, E. L., Luca, M., and Stanton, C. T. (2020b). What jobs are being done at home during the COVID-19 crisis? evidence from firm-level surveys. Technical report, National Bureau of Economic Research.

Benoit, D. (2021). Jamie Dimon on booming economy and finally getting off zoom. *Wall Street Journal*.

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

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.

Bojinov, I., Choudhury, P., and Lane, J. N. (2020). A field experiment on virtual social interactions and performance of remote workers. Technical report, Harvard mimeo.

Breiman, L. (2001). Random forests. *Machine learning*, 45(1):5–32.

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.

Choudhury, P., Foroughi, C., and Larson, B. (2021). Work-from-anywhere: The productivity effects of geographic flexibility. *Strategic Management Journal*, 42(4):655–683.

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.

Dutcher, E. G. (2012). The effects of telecommuting on productivity: An experimental examination. the role of dull and creative tasks. *Journal of Economic Behavior & Organization*, 84(1):355–363.

Emanuel, N. and Harrington, E. (2020). 'Working' remotely?: Selection, treatment and the market provision remote work". Technical report, Harvard mimeo.

Gajendran, R. S. and Harrison, D. A. (2007). The good, the bad, and the unknown about telecommuting: meta-analysis of psychological mediators and individual consequences. *Journal of applied psychology*, 92(6):1524.

Gibson, C. B. and Gibbs, J. L. (2006). Unpacking the concept of virtuality: The effects of geographic dispersion, electronic dependence, dynamic structure, and national diversity on team innovation. *Administrative Science Quarterly*, 51(3):451–495.

Gompers, P. A., Kaplan, S. N., and Mukharlyamov, V. (2020). Private equity and COVID-19. Technical report, National Bureau of Economic Research.

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.

Hinds, P. J. and Mortensen, M. (2005). Understanding conflict in geographically distributed teams: The moderating effects of shared identity, shared context, and spontaneous communication. *Organization Science*, 16(3):290–307.

Imai, K. and Khanna, K. (2016). Improving ecological inference by predicting individual ethnicity from voter registration records. *Political Analysis*, pages 263–272.

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.

Mas, A. and Pallais, A. (2017). Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–59.

Mas, A. and Pallais, A. (2020). Alternative work arrangements. *Annual Review of Economics*, 12:631–658.

Rothbard, N. P. (2001). Enriching or depleting? the dynamics of engagement in work and family roles. *Administrative science quarterly*, 46(4):655–684.

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.*

Table 1: Fraction of remote and onsite job listings by startup characteristics

	Remote	Onsite	Difference
Employees			
1-10	0.555	0.445	-0.110
11-50	0.137	0.863	0.726
51-200	0.076	0.924	0.848
201-500	0.062	0.938	0.876
501-1000	0.093	0.907	0.814
1001-5000	0.060	0.940	0.880
5000+	0.022	0.978	0.956
Funding round			
Pre-Seed	0.513	0.487	-0.026
Seed	0.232	0.768	0.536
Series A	0.089	0.911	0.822
Series B	0.063	0.937	0.874
Series C	0.040	0.960	0.920
Series D	0.059	0.941	0.882
Series E	0.174	0.826	0.652
Series F	0.041	0.959	0.918
Series G	0.012	0.988	0.976
IPO	0.112	0.888	0.776
Acquired	0.051	0.949	0.898
Location			
Outside SV	0.145	0.855	0.710
In SV	0.164	0.836	0.672

Table notes: This table compares the frequency at which remote and onsite listings appear in firms with different characteristics. For each section, the first two columns are the fraction of remote and onsite listings for firms in the category, and the third column is the difference between the two.

Figure 1: What types of employers issue listings for remote-eligible jobs?

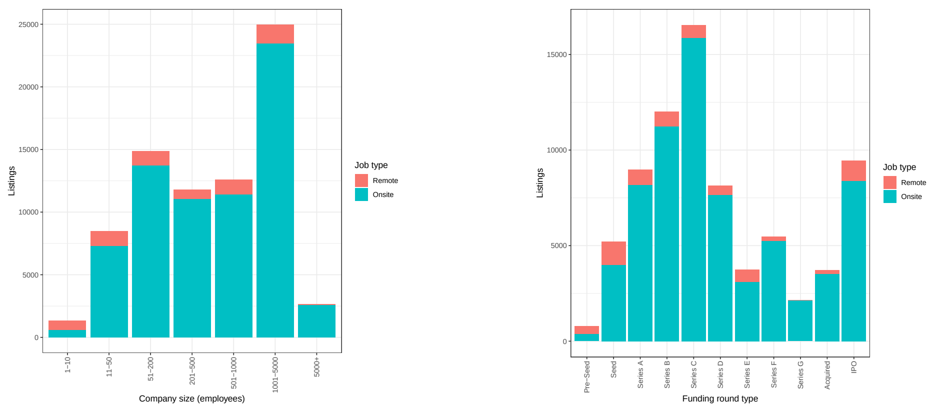

(a) Employer size

(b) Funding round type

Figure notes: These charts report the distribution of job listings by company size and funding round. Both pre and post COVID shutdown listings are included in the sample used to make these figures. The light blue bars designate onsite listings, and the orange bars designate remote-eligible listings.

Table 2: Applicant characteristics for job listings in sample

	Mean	Std. Dev.	Maximum	Minimum
Number of apps	3.092	4.229	103	1
Mean experience	4.351	3.065	10	0
Number of female apps	0.690	1.151	23	0
Number of URM apps	0.276	0.714	19	0

Table notes: This table reports application statistics for job listings in the sample. The first row is average number of applications for job listings in the sample that received at least one application. The second is average experience of applicants. The third row is the average number of female applicants to each listing. The fourth row is the average number of applicants to each listing from underrepresented minorities (URM).

Figure 2: Fraction of applications sent to remote openings

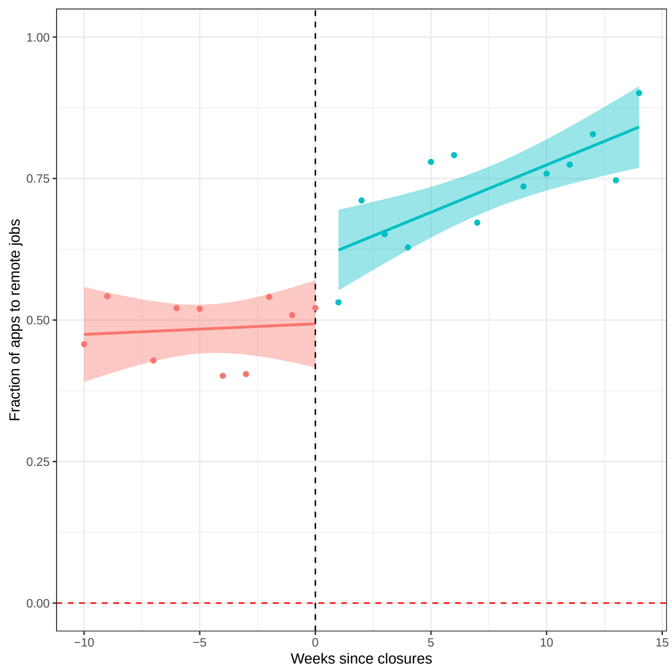

Figure Notes: This chart illustrates changes in the fraction of applications to listings that are remote-eligible. The x-axis denotes the number of weeks since the March 13th mandated COVID closures. The y-axis is the fraction of total applications sent to listings that are remote-eligible.

Figure 3: Remote-eligible listings posted before and after COVID shutdowns

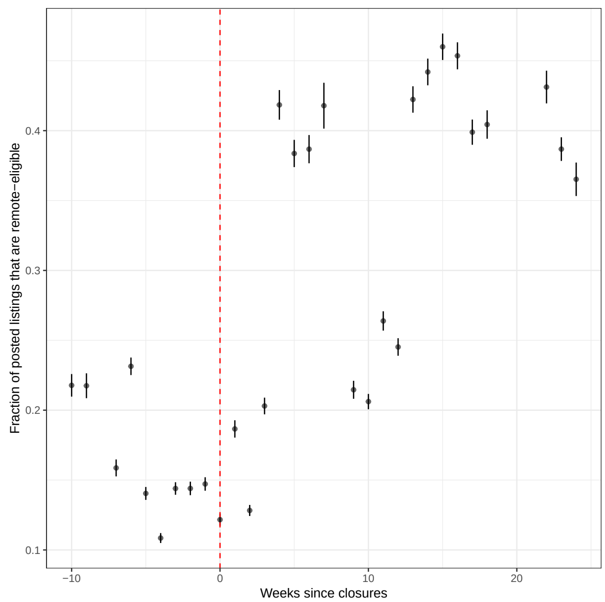

Figure Notes: This chart illustrates mean values and standard errors of the mean for the fraction of job listings in each week that are remote-eligible opportunities. The x-axis denotes the number of weeks since the March 13th mandated COVID closures. The y-axis is the fraction of total job listings that are denoted as remote-eligible.

Figure 4: Applicant attributes by job remote status

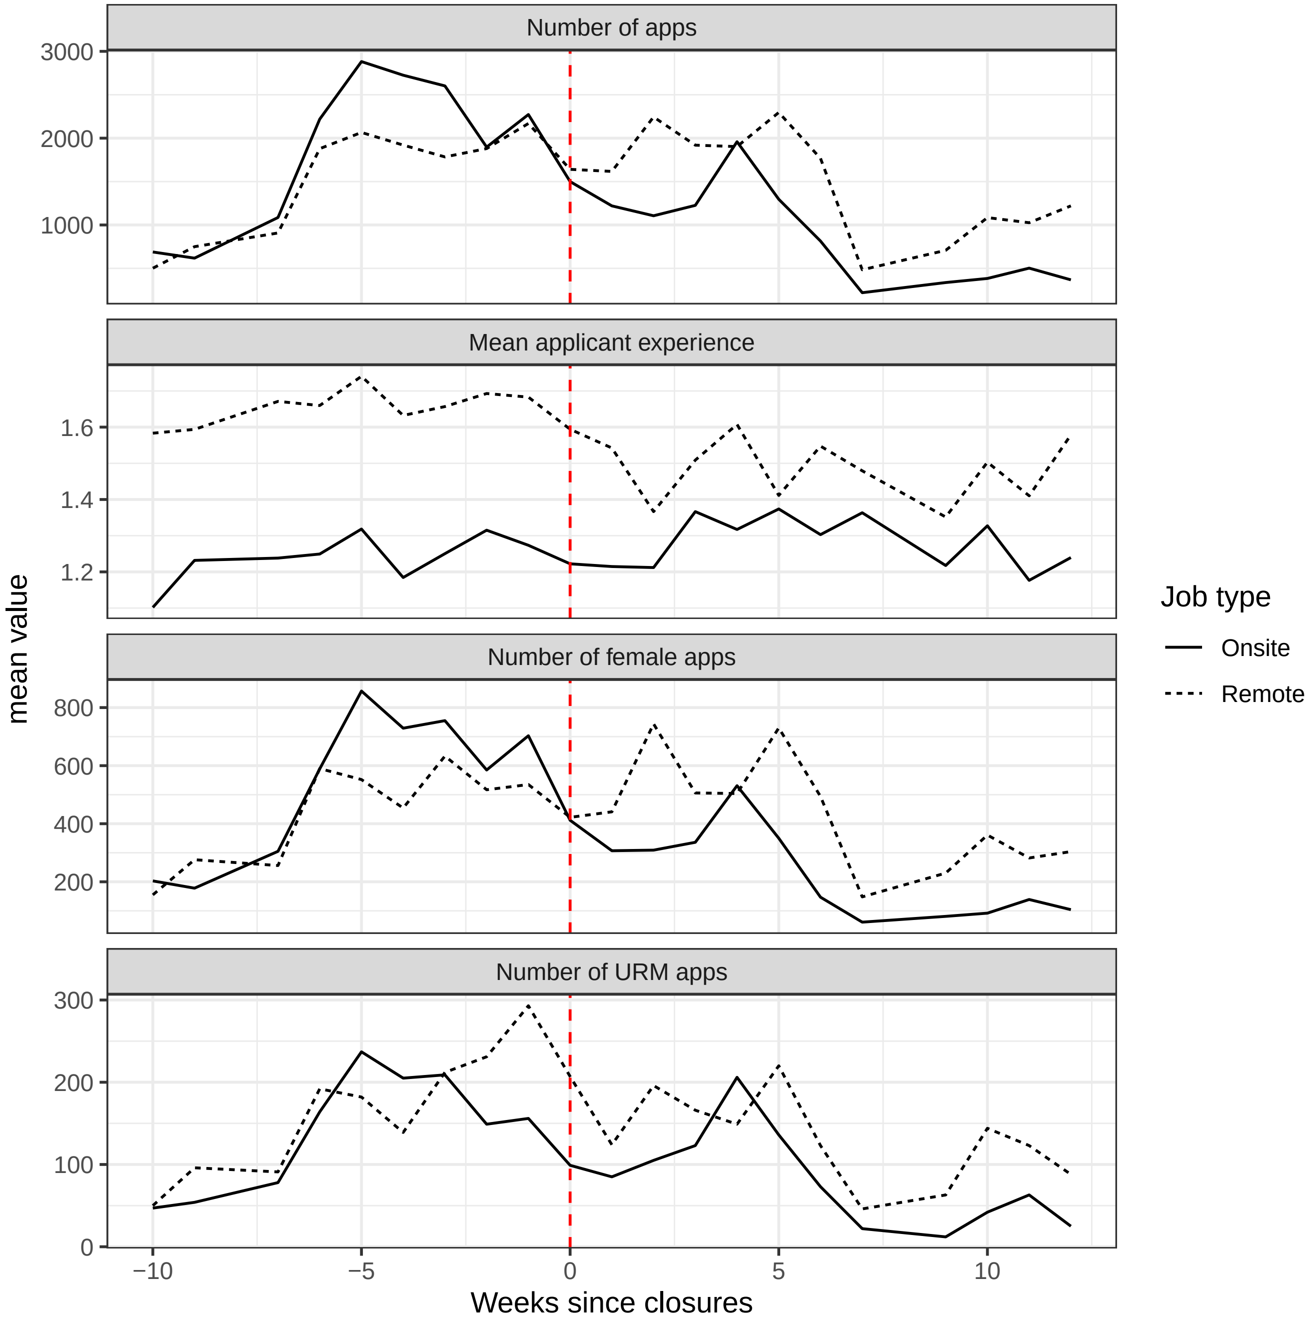

Figure Notes: The top panel plots the total number of applications against week since closures in panel. The second panel is the average experience of applicants against week in panel (experience measured as years of experience in their latest role). The third panel is the number of female applicants to listings. The fourth panel is the number of URM applicants to listings. In all panels, the solid line depicts the trend line for listings designated as onsite jobs, and the dotted line depicts the trend line for listings designated as remote jobs.

Table 3: Characteristics of applicants treating job listing remote status as exogenous (Pre-Covid sample)

	Dependent variable:			
	Log apps	Log exp	Log female apps	Log URM apps
	(1)	(2)	(3)	(4)
Remote	0.589***	0.486***	0.085***	0.107***
	(0.044)	(0.167)	(0.023)	(0.016)
Constant	−0.369	2.235**	−0.128	0.020
	(0.506)	(0.981)	(0.261)	(0.188)
Week FE	✓	✓	✓	✓
Job title FE	✓	✓	✓	✓
Location FE	✓	✓	✓	✓
Startup FE	✓	✓	✓	✓
Observations	5,754	875	5,754	5,754
R^2	0.624	0.895	0.467	0.445
Adjusted R^2	0.502	0.652	0.294	0.264

Table notes: This table reports OLS regressions on the numbers and characteristics of applications sent to job listings. The sample only includes applications posted prior to the COVID closures. The DV in cols 1-4 are logged number of apps, logged mean app experience, log number of female applications, and log number of URM applications, respectively. The number of observations in columns 1, 3, and 4 is all job applications in the sample. The number of observations in column 2 is job applications in sample that received at least one application. Standard errors are shown in parentheses, with *, **, and *** denoting significance at the 10%, 5%, and 1% level, respectively.

Table 4: Identifying onsite listings made remote-eligible by the Covid closures

	Predicted onsite	Predicted remote-eligible
Panel A: *2019 listings (Training sample)*		
Job is onsite	0.654	0.028
Job is remote-eligible	0.009	0.309
Panel B: *2019 listings (Test sample)*		
Job is onsite	0.609	0.088
Job is remote-eligible	0.090	0.213
Panel C: *2020 listings after closures*		
Job is onsite	0.289	0.082
Job is remote-eligible	0.383	0.246

Table notes: This table reports classification performance for a Random Forest algorithm that uses job and employer characteristics to predict whether a job is likely to be listed as remote-eligible or onsite. Panel A reports results from a training sample of a randomly selected 70% of the listings posted in 2019. Panel B reports results from a test sample of the remaining 30% of the listings from 2019. Panel C reports results for listings posted in the weeks after closure, which we use to identify "treated" jobs for the analysis shown in Table 5.

Table 5: Regressions based on natural experiment

	Dependent variable:			
	Log apps	Log exp	Log female apps	Log URM apps
	(1)	(2)	(3)	(4)
Remote & Predicted Onsite	0.258***	0.196**	0.077***	0.119***
	(0.040)	(0.094)	(0.020)	(0.018)
Onsite & Predicted Remote	−0.101***	−0.345***	−0.050***	−0.013
	(0.039)	(0.069)	(0.019)	(0.017)
Remote & Predicted Remote	0.055	0.481***	0.182***	−0.074***
	(0.056)	(0.093)	(0.028)	(0.024)
Constant	−0.829***	1.233***	−0.254**	−0.110
	(0.220)	(0.186)	(0.110)	(0.095)
Week FE	✓	✓	✓	✓
Job title FE	✓	✓	✓	✓
Location FE	✓	✓	✓	✓
Startup FE	✓		✓	✓
Observations	3,803	563	3,803	3,803
R^2	0.953	0.917	0.914	0.927
Adjusted R^2	0.940	0.880	0.891	0.907

Table notes: The sample is constructed using only the sample of jobs posted after the COVID closures. Applications are predicted to be remote or onsite using a random forest model trained on skill and title characteristics from job listing data posted in 2019. The DV in cols 1-5 are logged number of apps, logged mean app experience, log number of female applications, and log number of URM applications, respectively. The number of observations in columns 1, 3, and 4 is all job applications in the sample. The number of observations in column 2 is job applications in sample that received at least one application. Standard errors are shown in parentheses, with *, **, and *** denoting significance at the 10%, 5%, and 1% level, respectively.

A Correlates of remote job listings: addition detail

In this appendix, we provide more detail of our exploration of the organizational and job correlates of remote-eligible status. We start in Table A.1 by presenting a linear probability (OLS) regression framework to regress *remote* on organizational categorical variables which proxy for size (employee headcount) in column 1 (entering attained external funding status would be collinear with these size variables); location (*in Silicon Valley?*) in column 2; categorical skills in job variables (which we subjectively coded) in column 3; and all of the above in column 4.[10] Several patterns emerge, which are consistent with the patterns documented in the descriptive statistics: relative to the smallest organizational size category (employee headcount of one to ten), as the organizational size category increases, there appears to be a progressively lower likelihood of remote job listings. Silicon Valley jobs are less likely to be remote-eligible relative to jobs listed in other locations (column 2). Relative to the omitted skills outside of technical and business management, column 3 suggests that both technical and sales & marketing skill-oriented jobs are more likely to be remote-eligible. Management and other business-related skills are less remote-eligible, on average. Of particular note, however, is that with all organizational and job skills factors included in the regression specification (column 4), the overall variance the model explains (R-squared statistic) is only approximately 9%.

Especially disappointing from the standpoint of overall model performance is relating job characteristics, such as job title, to remote work (R-squared of about 0.02 based on coarse job skills), as our prior expectation based on the literature (such as Dingel and Neiman [2020]) is that occupation-level differences would be economically significant correlates of remote-work eligibility. We first examine in Table A.2 the frequency with which job titles are designated as remote-eligible. We show the job titles which are associated with the largest and smalles likelihood of a remote job listing by regressing *remote* on job titles and sorting the resulting estimates of each job title fixed effect. The left hand side of the table shows the job titles most associated with remote listings, such as "senior account manager of digital advertising," while the right hand side shows the 10 job titles least likely to be associated with remote listings, such as "computer vision engineer."[11] The second panel of Table A.2 shows an analogous list for requisite skills which are most (least) associated with remote job listings on the left (right) hand of the panel, using the same methodology as before (except replacing skills for job titles). Because of the large number of skills listed in the data, we confine our attention to the 250 most frequent skills in the data. The skills most associated with remote jobs are: "blockchain" and "spring" (a technical skill associated with backend and cloud engineering); the skills least associated with remote jobs are "Salesforce" and "C#" (a programming language).

We present an alternative approach to estimating how technical and business management skills relate to *remote* in a multivariate framework, without our manual classification. In Table A.3, we specify the 250 most frequent skills and employ a LASSO regression with a lambda value of 0.01. The top 30 skills with significant coefficient values (sorted in descending value) related to *remote* are listed, with the remaining 220 skills receiving no weight (we use the output of the LASSO model as inputs to skill-level fixed effects in other empirical tables).[12]

[10]These results are also robust to a logit model. We present OLS here to compare a model fit statistic with another table with a different estimation strategy and also because of the large number of fixed effects we will employ in subsequent analyses.

[11]This heterogeneity is consistent with broader surveys regarding cross-industry variation in remote work adoption in the face of COVID-19 [Bartik et al., 2020b] and how variation in the nature of work (particularly information work) correlates with remote work [Brynjolfsson et al., 2020].

[12]In those regressions, we predict *remote* using various categories of fixed effects. In descending order of importance as judged by R-squared statistics of overall model fit are startup (organizational) FEs; job title FEs; location FEs; skill FEs; and week FE (with a high exceeding 0.7 for startup FEs and a low of about 0.01 for just week FEs).

Table A.1: Predicting remote-eligible listings using organizational and job characteristics

	Dependent variable:			
	Remote-eligible job			
	Headcount	Location	Job skills	All
	(1)	(2)	(3)	(4)
11-50	−0.275*** (0.014)			−0.160*** (0.037)
51-200	−0.390*** (0.013)			−0.249*** (0.043)
201-500	−0.404*** (0.014)			−0.220*** (0.068)
501-1000	−0.267*** (0.014)			−0.295** (0.137)
1001-5000	−0.423*** (0.014)			
5000+	−0.441*** (0.018)			−0.455 (0.462)
in Silicon Valley		−0.037*** (0.006)		−0.056 (0.034)
Technical			0.019*** (0.004)	0.015*** (0.004)
Sales and Marketing			0.031*** (0.005)	0.028*** (0.005)
Business			0.0001 (0.023)	−0.010 (0.023)
Office & customer ops			0.027* (0.016)	0.021 (0.016)
Management			−0.039* (0.021)	−0.031 (0.021)
Constant	0.441*** (0.013)	0.096*** (0.003)	0.265*** (0.023)	0.440*** (0.036)
Observations	13,518	13,518	1,029	1,029
R^2	0.100	0.003	0.053	0.091
Adjusted R^2	0.100	0.003	0.048	0.081

Table notes: This table reports OLS regressions on whether a listing is remote. All listings are from the pre-COVID sample. The omitted variable in col (1) is 1-10 and in col (2) is outside SV. Columns (1) and (2) include all observations for which size and location are available. Columns (3) and (4) only include observations for which skills are also available. Standard errors are shown in parentheses, with *, **, and *** denoting significance at the 10%, 5%, and 1% level, respectively.

Table A.2: Most and least remote-eligible job titles and skills

Most remote-eligible jobs	*Least remote-eligible jobs*
Senior Account Manager - Digital Advertising	Computer Vision Engineer
Co-Founder	Senior Technical Recruiter
Social Media Intern	Salesforce Administrator
CTO	Sales Operations Analyst
Administrative Assistant	Finance Manager
Content Writer	HR Coordinator
Web Developer	Head of QA
iOS Developer	Partnerships Manager
Chief Technology Officer	Junior Customer Success Manager
Backend Developer	Director of Marketing

Most remote-eligible skills	*Least remote-eligible skills*
Blockchain	Salesforce
Spring	C#
Wordpress	SQL
Creative Writing	Operations Management
Back-End Development	Finance
Community Management	Business Operations
Copywriting	Scala
Blogging	Operations
Social Media Strategy	Microservices
Fluent in English	Software

Table notes: The top table lists the job titles that are most and least likely to appear in remote listings. Only titles that appear at least 200 times in the data set are used for this analysis. The bottom table lists the skills that are most and least likely to appear in remote listings. Only skills that appear at least 200 times in the data set are used for this analysis. Both tables only use data from the pre-COVID closure period.

Table A.3: Skills and remote status

Skill	Coefficient
Spring	0.513
Social Media Marketing	0.246
Blockchain	0.215
Social Media	0.197
Customer Service	0.178
Express.js	0.171
Sales and Marketing	0.160
Networking	0.156
Content Creation	0.146
Linux	0.108
PostgreSQL	0.070
MySQL	0.062
Graphic Design	0.048
Sales	0.043
React Native	0.042
Web Development	0.042
Copywriting	0.034
NodeJS	0.029
Java	0.023
Product Development	-0.002
Recruiting	-0.008
Microservices	-0.062
Embedded Systems	-0.121
SaaS	-0.121
SQL	-0.188
R	-0.237
Jenkins	-0.294
Business Operations	-0.385

Table notes: This table reports the results of coefficient estimates from a LASSO regression of major skill categories on whether a job is listed as remote-eligible. A positive (negative) coefficient indicates that a skill is more (less) likely to be associated with a remote-eligible job. Only the top 250 skills were used in the regression.

B Impact of remote-eligible job listings on applications: alternative estimation strategies

While our main empirical estimation strategy for how designating jobs as remote-eligible affects applicant quality is the post-COVID shutdowns contained in the main text, in this appendix section, we describe two alternative strategies to estimate the effect. A first approach matches treated and control observations by exact job title. This approach allows us to sharpen the sample to eliminate undesirable heterogeneity (if the treatment and control groups are unbalanced on key observable variables, the threat of unobserved third factors could generate spurious results).[13] Table B.1 presents the regression results of applicant characteristics in response to job listings using the matched sample. In Table B.1, the *remote* coefficient is positive and significant across the applicant characteristic outcome variables at the 1% level. In comparing the estimated economic effects in Table B.1 relative to the estimates in Table 3 (in which we treat remote jobs as exogenously-given), we see that the estimates in Table B.1 are smaller in magnitude (a discrete change in remote status corresponds to a 4-7% increase in female and URM applications). Relative to the natural experiment approach, any matching empirical strategy can only match on observable variables, leaving open the possibility that unobserved variables could correlate with *remote* listings, however.

A second alternative approach for estimating how applicant characteristics are shaped by designating jobs as remote-eligible centers on a feature of job postings we can detect: whether each job posting was posted "automatically" on a fixed schedule, or "manually." A natural question to ask is how do jobs posted in one manner differ from those posted in the other? Perhaps jobs which are newer or which require more fine-tuning are more likely to be manually posted. On the other hand, auto-posted jobs may arise from one or more of the following conditions: (1) these jobs are "important" enough and sufficiently difficult to fill that firms post these positions on a fixed schedule; and/or (2) managers believe that more recent job posting may be more noticeable by applicants who themselves may confine their search to more recent listings.

Under the assumption that manual job postings designated as remote-eligible and made after the COVID-induced workplace shutdowns reflect active managerial choice, a quasi-experimental approach would be to designate manually-posted remote jobs post-shutdowns as the "treated" condition as compared to alternative job design and time period configurations (we take as the baseline on-site jobs which are automatically posted in the pre-COVID era). The results suggest that the "treatment" condition as compared to the baseline yields largely similar applicant patterns as those we report in the natural experiment featured in the main text. The effect on the number of female applicants is projected to be increased by 10% as compared to the control. Under-represented minorities (URM) are predicted to increase 18.1% for the treatment condition as compared to the control.

There a number of interpretational issues associated with this auto versus manually job posted analysis. Among them is that the treatment condition could reflect jobs which were always remote-eligible, but were manually modified after the shutdowns to reflect updated work conditions or preferences (and the job listings were important enough for managers to make these updates). Another path of data generation for the treatment condition would be jobs which were formerly on-site, but which are now remote. With some maintained assumptions, we may be able to distinguish these in the data, but we follow that basic approach in the main natural experiment. Here, we bundle both, as we wish to simply compare against automatically-posted jobs.

[13]Relative to Table 3, which uses only the pre-closure sample, Table B.1 uses both pre- and post-closure job postings, though matched according to the above criteria.

Table B.1: Regressions using sample matched on job title

	Dependent variable:			
	Log apps	Log exp	Log female apps	Log URM apps
	(1)	(2)	(3)	(4)
Remote	0.233***	0.498***	0.046***	0.073***
	(0.030)	(0.111)	(0.015)	(0.011)
Constant	0.185	1.420***	0.049	0.004
	(0.388)	(0.407)	(0.192)	(0.139)
Week FE	✓	✓	✓	✓
Job title FE	✓	✓	✓	✓
Location FE	✓	✓	✓	✓
Startup FE	✓		✓	✓
Observations	4,209	457	4,209	4,209
R^2	0.581	0.682	0.360	0.442
Adjusted R^2	0.495	0.485	0.229	0.328

Table notes: This table reports OLS regressions on the numbers and characteristics of applications sent to job listings on a set of observations matched on job title and all matches are required to be exact. Listings from both before and after the closures are used in the analysis. The DV in cols 1-5 are logged number of apps, logged mean app quality, logged mean app experience, log number of female applications, and log number of URM applications, respectively. The number of observations in columns 1, 4, and 5 is all job applications in the matched sample. The number of observations in columns 2 and 3 is job applications in the matched sample that received at least one application. The *MatchIt* package in R was used for matching. Matching is done without replacement and matches are one-to-one. All regressions include fixed effects for title, week of job posting, and location. Standard errors are shown in parentheses, with *, **, and *** denoting significance at the 10%, 5%, and 1% level, respectively.

Table B.2: Regressions with auto-posted and manually-posted job listings on the POST-COVID sample

	Dependent variable:			
	Log apps	Log exp	Log female apps	Log URM apps
	(1)	(2)	(3)	(4)
REMOTE	0.884***	0.544	0.100***	0.181***
	(0.069)	(0.396)	(0.031)	(0.022)
MANUAL	−0.077***	0.135*	−0.020**	0.004
	(0.019)	(0.082)	(0.009)	(0.006)
POST-PERIOD	−0.140***	0.586	−0.024**	−0.018**
	(0.025)	(0.557)	(0.011)	(0.008)
REMOTE × MANUAL	−0.606***	−0.070	−0.077**	−0.146***
	(0.073)	(0.432)	(0.033)	(0.023)
REMOTE × POST-PERIOD	−0.816***	−0.053	−0.063**	−0.181***
	(0.069)	(0.433)	(0.031)	(0.022)
MANUAL × POST-PERIOD	0.113***	−0.024	0.012**	0.011***
	(0.013)	(0.116)	(0.006)	(0.004)
REMOTE × MANUAL × POST-PERIOD	0.648***	0.027	0.055*	0.178***
	(0.073)	(0.480)	(0.033)	(0.024)
Constant	0.046	1.708***	−0.002	−0.026
	(0.097)	(0.606)	(0.044)	(0.031)
Week FE	✓	✓	✓	✓
Job title FE	✓	✓	✓	✓
Location FE	✓	✓	✓	✓
Startup FE	✓		✓	✓
Observations	21,724	917	21,724	21,724
R^2	0.347	0.693	0.179	0.200
Adjusted R^2	0.299	0.474	0.118	0.141

Table notes: The sample is constructed using only the sample of jobs and applications from the post-COVID period. Applications are predicted to be remote or onsite using a random forest model trained on skill and title characteristics from job listing data posted prior to 2020. The DV in cols 1-5 are logged number of apps, logged mean app quality, logged mean app experience, log number of female applications, and log number of URM applications, respectively. The number of observations in columns 1, 4, and 5 is all job applications in the sample. The number of observations in columns 2 and 3 is job applications in sample that received at least one application. Standard errors are shown in parentheses, with *, **, and *** denoting significance at the 10%, 5%, and 1% level, respectively.