"WORKING" REMOTELY?

SELECTION, TREATMENT, AND THE MARKET PROVISION OF REMOTE WORK

Natalia Emanuel · Emma Harrington[1]

This version: April 9, 2021
Latest Version: Click here

Abstract: Why was remote work rare prior to the pandemic? One possibility is that remote work reduced worker productivity. Another is that it attracted less productive workers. We test these possibilities in the call-centers of a Fortune 500 online retailer. We find that working remotely increased call-center workers' productivity. When previously on-site workers took up opportunities to go remote in 2018, their hourly calls rose by 7.5%. Similarly, when COVID-19 closed on-site call centers, a difference-in-difference suggests that the productivity of workers who switched to remote work rose by 7.6% relative to their already remote peers. Despite these positive productivity effects, remote workers were 12pp less likely to be promoted. If better workers are more concerned about being overlooked in remote jobs, remote workers will be adversely selected. Consistent with this theory, we find evidence that remote work attracted latently less productive workers. When all workers were remote due to COVID-19, those who were hired into remote jobs were 18% less productive than those who were hired into on-site jobs. Extending remote opportunities to on-site workers similarly attracted less productive workers to on-site jobs. The sorting of workers by ability meant some workers opted out of remote work because they did not want to pool with less productive workers. This led to deadweight losses in the market for remote work. Looking forward to life after the pandemic, our model suggests that COVID-19 could attenuate these losses if firms have learned how to better evaluate remote workers or workers' tastes for remote work have become more heterogeneous.

[1]Contact: Harvard University, 1805 Cambridge Street, Cambridge, MA 02138, eharrington@g.harvard.edu. We thank Nathan Hendren, Claudia Goldin, Lawrence Katz, Edward Glaeser, Louis Kaplow, Amanda Pallais, Elie Tamer, Jeff Liebman, and participants of the Public Finance and Labor Economics Workshop at Harvard for helpful comments. We are grateful to our colleagues, Lisa Abraham and Jenna Anders, as well as Alex Albright, Dev Patel, Ashesh Rambachan, Ljubica Ristovska, and Hannah Shaffer. This project would not have been possible without the curiosity and commitment to research of our colleagues at the firms who shared data: Lauren and Trevor. We are grateful for financial support from the National Science Foundation [Natalia] and the Lab for Economic Applications and Policy. The findings and conclusions expressed are solely those of the authors and do not reflect the opinions or policy of the organizations that supported this work.

I INTRODUCTION

Prior to Covid-19, only 5% of Americans worked remotely all of the time.[2] A few months into the pandemic, most workers who could work remotely did so (Brynjolfsson et al., 2020).[3] So, what we can expect once the pandemic subsides? The answer to this question hinges on why remote work was so rare prior to COVID-19.

One possibility is that working remotely reduces productivity. Another possibility is that remote jobs attract latently less productive workers. If better workers are more concerned about being overlooked for promotion when out of earshot of their managers, remote workers will be adversely selected.

In this paper, we test these possibilities using data from the call-centers of a Fortune 500 online retailer. Call-center work is an easily "remotable" job and one that has been the focus of existing scholarship on remote work.[4] The largest experiment on preferences for remote work was run in a US call-center (Mas and Pallais, 2017) and the largest experiment on its productivity effects was run in a Chinese call-center (Bloom et al., 2015).

Mas and Pallais (2017) find that call-center workers are willing to accept 8% lower wages to have the option to work remotely. Given the low rates of remote work among call-center workers, this high willingness to pay suggests remote work is costly for firms. However, Bloom et al. (2015) finds no such costs, with remote work increasing productivity by 14% (Bloom et al., 2015).

We see the same disconnect during the pandemic. The majority of workers report being happier and more productive working remotely, but few firms have started to advertise jobs that will be permanently remote (PwC, 2020; Morning Consult, 2020; Barrero et al., 2020; Ovide, 2021).[5]

We argue the missing piece to this puzzle is career concerns. Indeed, in Bloom et al. (2015)'s experiment, remote work halved workers' chances of promotion, opening up the possibility of unraveling: remote jobs may become dead-end jobs that only attract workers who are unlikely to

[2] In the 2018 American Community Survey (ACS), 5.3% of workers reported working from home (based on the authors' calculations). In the American Time-Use Survey between 2013 and 2017, 20.5% of workers reported spending some time working from home and 11.4% reported spending the entire day working remotely that day (Brynjolfsson et al., 2020). Since 2000, the share of workers working remotely has been fairly constant despite continued innovation in communication technologies, as illustrated in Figure A.1 (Mas and Pallais, Forthcoming). Oettinger (2011) found that home-based work expanded substantially between 1980 and 2000, albeit starting from a very low base.

[3] These estimates are also complementary with those from the Bureau of Labor Statistics, where 35% of workers report working remotely because of the pandemic.

[4] As has become clear during the pandemic, a large share of jobs are remotable. While many jobs cannot be done from afar — it's hard to clean houses, move furniture, or flip burgers remotely — other jobs don't involve any of these needs to be in-person. Dingel and Neiman (2020) use O*Net information about occupations to classify each occupation as either having an in-person need or being remotable. This analysis suggests that 37% of jobs were remotable in the US at the outset of the pandemic.

[5] In two surveys of 2,207 remote workers, 32% report wanting to remain fully remote after the pandemic and 56% report wanting to work remotely for more than half their week (PwC, 2020; Morning Consult, 2020). Similarly, in a ZipRecruiter survey, 45% say they want a job that would let them be permanently remote. However, only 8 or 9% of jobs are actually permanently remote, up just 6pp from before the pandemic.

advance.

Our paper builds on the nascent literature on remote work by providing new estimates of the productivity effects and promotion penalties of remote work in the US context. We develop a two-period model that incorporates workers' career concerns. This model predicts that differences in promotion opportunities will lead to differences in worker selection. We then test this prediction empirically, offering the first analysis of productivity differences between workers who are hired into comparable remote and on-site jobs. Natural experiments at the retailer allow us to separately identify the treatment and selection effects of remote work. We pull these estimates together to quantify the inefficiencies that arise from adverse selection into remote work. We organize these analyses into three parts.

The first part provides descriptive evidence on the relationship between remote work and promotion, which then motivate a model of the choice to be remote or on-site that turns on career concerns. At our retailer, workers who chose remote jobs were 12pp less likely to be promoted than those who chose on-site jobs, complementing the results of Bloom et al. (2015).[6] We find suggestive evidence that these promotion differences reflect differences in the information that managers acquire about remote and on-site workers. Managers' evaluations help predict the future performance of on-site workers but not remote workers.[7] In our model, this motivates the hypothesis that working remotely increases the probability that capable workers will be overlooked for promotions. Since better workers have more to lose from being overlooked for promotion, they will only choose remote work if they have more extreme tastes for remote work.[8] As a result, fewer high-ability workers choose remote work, leading to adverse selection into remote jobs. This adverse selection raises the average cost of remote work above its marginal cost, which causes the market price of working remotely to exceed the efficient incentive.[9]

In the second part of the paper, we identify remote work's marginal cost — or the treatment effect — and its selection effect using natural experiments at the retailer.

At the retailer, some workers transitioned from on-site to remote work, first due to a voluntary program and later due to COVID-19. These transitions allow us to identify the treatment effect of remote work on productivity. In 2018, the retailer posted opportunities for on-site workers to go remote. Workers had discretion over whether they went remote but not when they did so, which depended on a spot opening up on a remote team. We exploit the quasi-random timing of workers' transitions to remote work in an event study design. We find productivity sharply

[6] If the estimates from Bloom et al. (2015) apply, two thirds of this gap is due to the causal effect of remote work on promotion opportunities.

[7] These differences in information would compound any direct effect of face-time on managers' perceptions on workers' level of dedication (Bailyn, 1993; Elsbach et al., 2010).

[8] This builds on the literature suggesting that promotions and bonuses can affect worker selection. Bender et al. (2018) find that better management practices that identify and reward good performances tend to attract better workers to the firm. Similarly, Brown and Andrabi (2020) find that performance pay induces better worker selection. Manchester (2012) finds a similar phenomenon for tuition reimbursement programs.

[9] Our model is most similar to Einav et al. (2010)'s formulation of Akerlof (1978). It also shares many features of classical labor market models of adverse selection (Salop and Salop, 1976; Miyazaki, 1977; Weiss, 1995).

rose by 7.5% when workers went remote with no sacrifice in customer satisfaction reviews. This complements Bloom et al. (2015)'s experimental results and similar quasi-experimental findings of Choudhury et al. (2020) in the US Patent Office.

We find similar patterns around COVID-19's lockdown: when the offices closed down, on-site workers were forced into remote work while already remote workers continued at home. In a difference-in-difference design, on-site workers' productivity rose by 7.6% relative to that of their already remote peers. This result extends the findings of Bloom et al. (2015) and Choudhury et al. (2020) since we estimate the treatment effect for workers who do not necessarily want to work remotely. This also contributes to the growing literature on COVID-19's productivity effects, which has, for example, found that remote work decreased time spent in meetings using similar designs (DeFilippis et al., 2020; Yang et al., 2020).

These natural experiments also identify the selection effect of remote work. Identifying remote work's selection effect is usually challenging because workers who choose remote and on-site jobs typically end up in different roles, often at different firms. Thus, productivity comparisons are often infeasible and uninterpretable. Our setting is unusual in that the retailer hired workers into remote and on-site jobs and randomly routed calls between them.

After COVID-19's lockdown, all workers worked remotely. But some had originally chosen to be remote and some had not. We find that those who originally chose remote jobs took 18% fewer calls per hour than those who originally chose on-site jobs during the lockdown. Using a caregiver survey conducted by the retailer, we find that productivity differences between workers who initially chose to be on-site and remote are not driven by child-care responsibilities, suggesting career concerns rather than constraints at home drive adverse selection.[10]

The introduction of a remote work program in January of 2018 yields another lens on selection. Among workers who ultimately took up opportunities to go remote, some were initially unaware that they could go remote and others were recruited with this possibility in mind. Workers who were hired before January of 2018 chose a job thinking that it would always be on-site; workers who were hired after January of 2018 chose a job that they knew could go remote. Thus, only later hires could be selected because of their readiness to go remote. Consistent with this story, the selection effects of remote work only show up for those who were hired after the policy change. Among these later cohorts, workers who went remote were 11.8% less productive than their peers who opted to remain on-site. By contrast, among earlier cohorts, workers who went remote were 7% more productive than their peers who remained on-site, consistent with a positive treatment effect of remote work. The difference in these differences suggest that the offer of remote work attracted workers who were 18.8% less productive, despite receiving same pay and always being trained on-site. This design complements the analysis in Linos (2018), which finds similar patterns in the US Parent Office around its introduction of a work from home program.

[10]This is in some contrast to Adams-Prassl (2020), which finds that women working on MTurk who have an infant at home are more likely to have work interruptions. Concerns about such distractions are reinforced by popular articles that tout remote work as especially beneficial for parents (see, for example, Schulte, 2020).

In the third part of the paper, we pull these estimates together to understand how they shape the market for remote work and the surplus that remote work generates. Using the estimated demand for remote work in our context, we find that adverse selection reduces the provision of remote work by 10pp. This distortion reduces the surplus from remote work by 29% which amounts to about $1 billion lost annually across the 3.5 million call-center workers in the US.[11] This is an especially notable effect since the career costs of remote work may be lower in call-center jobs than other lines of work because productivity is partially tracked electronically and many call-center workers only plan to stay in this occupation temporarily. Thus, in other settings, the distortion from the signaling component of remote work may be larger.

Looking forward to life after the pandemic, our analysis suggests that the lasting impacts of the mass experiment in remote work will hinge on how this experience has affected worker sorting into remote and on-site work. If workers learned about their idiosyncratic tastes for remote work during the lockdown, workers' decisions about whether to be remote or on-site may be less driven by private information about ability going forward. Further, if firms have learned how to better identify remote workers to promote, this will attenuate workers' incentives to sort into remote and on-site jobs by ability.[12]

The rest of the paper proceeds as follows. Section II briefly introduces our data and context. Section III offers descriptive evidence on the differences in promotion between remote and on-site workers. Section IV presents our model of how career concerns can shape the market for remote work. Section V builds intuition for our estimation strategies. The next three sections estimate the average and marginal costs of remote work using natural experiments at the retailer. Section IX estimates workers' demand for remote work in our context. Section X details how these forces shape the market for remote work and the inefficiencies in its provision. Section XI concludes and discusses unanswered questions about remote work.

II DATA

Our data come from the call-centers of a Fortune 500 online retailer between 2018 and 2020. During this time, we observe 3,440 call-center hires, of whom 84% were recruited into on-site jobs and 16% were recruited into remote ones. These workers answered incoming calls — such as "when will my order arrive?" or "can I change my shipping address?" The typical call lasted 9.5 minutes (standard deviation = 4.0 minutes) and in a typical hour, workers answered 3 calls, with the remainder of the hour spent waiting for incoming calls and taking breaks (standard deviation =

[11]This finding builds on the literature that investigates how selection can lead to an under-provision of workplace amenities. Tô (2018) finds evidence that taking parental leave is a negative signal about a worker's subsequent productivity (Goldin et al., 2020). The possibility for self-selection into jobs with certain amenities has been stressed as a motivation for government mandated benefits broadly (Summers, 1989) and in the context of workers' compensation insurance (Gruber and Krueger, 1991) and maternity leave specifically (Gruber, 1994; Ruhm, 1998).

[12]However, as noted by Juhász et al. (2020), management practices often take time to adapt to possibilities opened up by new technologies. Suggestively, job posts for new remote positions have remained low in the pandemic, indicating, perhaps, that much of the pandemic's effects on remote work may not last (Ovide, 2021).

4.4 calls).[13] After new recruits finished their three weeks of training, they were all randomly routed calls from the same pool, regardless of whether they were on-site or remote.[14] With sufficient experience, some workers were promoted into more specialized roles; to insure fair comparisons of workers' productivity, our analyses focus on the first six months of workers' tenure before these promotions occur.

Workers could handle most calls by themselves. For trickier calls, they could ask their managers for help in an online chat or an in-person conversation. They could also forward calls to more experienced or specialized teams. As illustrated in Figure 1, workers answered calls faster as they gained experience at the retailer because they had more answers on hand. This growth occurred for both remote and on-site workers, suggesting remote work did not preclude on-the-job learning.

Every call was tallied in the retailer's database. In addition to tallying calls, the retailer tracked a proxy of their quality: after each purchase, the customer was asked about her satisfaction with her experience (from one to five stars). This was then ascribed to all the workers who spoke with her. Electronic monitoring of call quantity and quality gives managers (and econometricians) considerable insight into the productivity of workers, even when they were remote.

However, these metrics are imperfect. A quick call might be efficient, curt, or just lucky. A satisfied customer might not leave a review (the participation rate is 11.5%); a dissatisfied one might leave a 5-star review to be polite (the mean review is 4.9 out of 5); and an irate customer might leave a 1-star review regardless of what the worker says.[15]

Thus, despite this detailed data, the firm is often uncertain which workers did a good job. This creates pitfalls in performance pay. In customer service, there is a tradeoff between the quantity and quality of calls — faster calls tend to end with less-satisfied customers.[16] Call quantity is easy to record, while call quality is more difficult to measure. If the firm paid piece-rates per call, then workers would optimize for speedier calls rather than more satisfied customers. Given this inevitability, firms often choose weaker performance incentives (Holmstrom and Milgrom, 1991; Baker, 2002).[17] At the online retailer, bonuses accounted for at most 17% of annual compensation,

[13]In the raw data, entry-level remote workers answered 0.089 more calls/hour than on-site workers answered: this comparison is misleading because remote workers were hired later on average when the online retailer was fielding more calls.

[14]Calls were randomly routed between workers online at the same time. One natural concern is that remote and on-site workers work at different times. Working hours were determined by time-zone and the retailer employed workers on-site and remote workers in all four continental time-zones. Thus, our productivity analyses include date by time-zone fixed effects.

[15]Each call is recorded, allowing managers to do additional quality-assurance checks. However, the incentive system at the retailer makes it difficult to trust these reviews. Managers are judged based on their workers' performance. Thus, some managers might turn a blind eye to a bad call since negatively rating their workers reflects poorly on them too.

[16]Across workers, a one standard deviation increase in calls handled is associated with a 0.1 standard deviation decrease in average customer satisfaction reviews, controlling for the month and time-zone in which the worker was hired (se = 0.0378, p-value= 0.0081).

[17]In conversation, upper management at the retailer also expressed other reservations about increasing incentive pay. Line managers always complained workers were gaming the system. Workers always complained that the retailer was rigging the system (and that their coworkers were cheating). There may have been both bad behavior and bad

with the remaining 83% a fixed hourly wage.[18]

The information frictions in the job also have implications for the econometrician, who must worry that the metrics at hand are misleading. However, the firm's choice to make compensation primarily hourly rather than piece-rates makes the quantity of calls a less problematic barometer of productivity. When workers do not have strong incentives to sacrifice quality for quantity, the quantity of calls is a more useful metric of productivity.[19]

The same challenges that made piece-rates problematic led to uncertainty about whom to promote. When a manager observed a worker handling many calls, there were always two possibilities. One, the worker was handling calls both quickly and kindly and operating on a higher production possibility frontier (PPF). Alternatively, the worker was handling calls quickly but curtly, trading quality for quantity along a lower PPF. Since the firm wanted to promote the worker under the first scenario but not the second, this created noise in the promotion process. Beyond the promotion choice, this uncertainty created noise in managers' regard for workers. This mattered to workers for two reasons. One, promotions increased pay by $2/hour, or 13% of wages given the average base of $15/hour. Two, many workers needed reference letters for new jobs — indeed, on average, the entire call-center workforce turned over about four times within a single year (see row 9 of Table 1).[20]

When metrics are complete, face-time can matter. Indeed, those who were within earshot of their managers were promoted more frequently: in row 8 of Table 1, 17% of on-site workers were promoted compared to less than 5% of remote workers (difference = -12.3pp, se = 1.15). If better workers — who operate on higher PPFs — gain more from being observed more closely, then they may disproportionately opt into on-site jobs.

The difficulty of ascertaining worker quality also meant there was no foolproof test for new recruits. Since many workers were coming to this job with little prior experience — indeed, the average age is 32 in row 2 of Table 1 — resumes were often of limited help. This gave scope for a worker's choice of whether to be on-site or remote to provide additional information about her likely ability.

If workers were paid according to their average product, then differences in worker selection would lead to difference in initial wages between on-site and remote jobs. Indeed, those hired

outcomes from the suspicion of bad behavior. The potential for incentives to lead to cheating is noted by Cadsby et al. (2010) and the potential for these schemes to reduce intrinsic motivation is noted by, for example, Benabou and Tirole (2003); Bénabou and Tirole (2006).

[18]By contrast, within the same retailer, sales' workers have about half of their pay in bonuses because the firm can directly observe a salesperson's revenue and profits.

[19]As Goodhart's law warns, once a useful number becomes a measure of success, it can cease to be a useful number. Thus, call volumes can be a good measure of productivity that is nonetheless problematic to use as the basis of pay because it it incomplete.

[20]The differences in turnover between remote and on-site workers is economically meaningful (albeit statistically imprecise): the firm faced a bigger risk of losing remote hires quickly and bearing the costs of a three-week training with no subsequent return. While economically meaningful, these turnover differences do not drive the productivity analyses, which give less weight to workers who spent less time answering calls at the retailer.

into on-site jobs were paid $1/hour more at the time of hire (row 11). The initial wages of on-site workers ranged from $14/hour to $16/hour and were set to reflect the pay in the local outside options in customer-service (CSR) in the metropolitan statistical area (MSA) (row 12). By contrast, all remote hires were paid $14/hour regardless of where they lived.[21] We will argue that this difference in pay is the result of equilibrium differences in worker quality between on-site and remote jobs, even after accounting for pay's direct effect.

Workers' career concerns were not the only drivers of their choice to take on-site or remote work. Women were more likely to be remote than men — 88% of remote workers are female compared to 69% of on-site workers (row 1 of Table 1). This may reflect women's more limited geographic mobility (e.g., Caldwell and Danieli (2018); Le Barbanchon et al. (2021)) or more extensive responsibilities in the home. Remote workers were also more likely to have caregiving responsibilities, as reported in a retailer-wide survey in June 2020. While nearly 60% of the on-site population had caregiving responsibilities, 74% of the remote population has these responsibilities, most of which reflect caring for children.

These patterns are emblematic of those in the American workforce more broadly. As detailed in Table A.1, among employed prime-age workers without college degrees in 2018, remote workers were 5pp more likely to be female, 4pp more likely to have children at home, and 0.4pp more likely to have children under 5, with starker differences in childcare responsibilities among female workers.

The need to juggle caregiving with working may be one of the reasons that remote work was often part-time work in the workforce as a whole. Within the retailer, the vast majority of both remote and on-site workers work full-time (94% to 95% in row 13 of Table 1).

In the American workforce, remote workers were also more likely to have physical disabilities — but not cognitive ones — suggesting limited mobility may be another driver of workers' preference to be remote.

Commute time is another salient factor. At the retailer, when a remote work program was introduced, on-site workers who went remote were 13.5pp more likely to live at least 15 miles from the office than their peers who remained on-site (se = 0.14).[22]

To summarize, despite the detailed information that the firm collected, there were important informational frictions in this job. The firm was consequently uncertain which workers deserved higher pay and promotions. This uncertainty had two potential implications for remote work: (1) it could incentivize better workers to choose an on-site job to be in earshot of their managers

[21]In theory, remote work could allow the retailer to recruit workers from less expensive labor markets with lower average pay. In practice, the retailer located its physical call-centers in low-wage labor markets, while recruiting remote workers from across the country. Even though the remote workers at the retailer were drawn from lower paying labor markets than the average customer service worker (with an average local wage of $16.40/hour), the remote call-center workers still had better outside options than the retailer's on-site workers on average.

[22]This comes from a regression that includes hiring month by call-center fixed effects with standard errors clustered at the individual level.

and (2) it could incentivize all workers to go on-site to signal higher ability. In addition to career concerns, there seem to be other drivers of workers' desire to be remote, such as commute time and caregiving. Thus, a worker's choice to be an on-site or remote job could only be an imperfect signal of likely ability.

III Descriptive Evidence on Promotion

Before presenting the model, we provide two stylized facts about promotion.

One, on-site workers are promoted in greater numbers and more quickly within their tenure at the retailer. Figure 2 depicts the percent of on-site workers (in blue) and remote workers (in orange) who have been promoted on the y-axis as a function of their tenure at the retailer on the x-axis. Among workers who persist at the retailer, 33% of on-site workers are promoted compared to 15% of remote workers. The differences are even starker within the first year. At the year-mark, nearly 30% of on-site workers have been promoted compared to just over 10% of remote workers.

These differences in promotion rates could either be due to differences in latent skill between remote and on-site workers or to differences that do not reflect the workers' ability to do the high-skill job. These ulterior differences could reflect (a) managers' knowledge about whom to promote on-site, (c) managers' bias toward face-time, or (d) managers' preferences for those in the office. If skill is still a prerequisite for promotion, any of these mechanisms could lead to an incentive for better workers to choose on-site work.

Two, we provide suggestive evidence on managers' knowledge about whom to promote. To investigate this question, we consider the predictive power of managers' reviews of workers' performances. In December of 2018, the retailer's managers were asked to rate each worker on a scale from one to five — with one signifying unsatisfactory performance; two, inconsistent performance; three, meeting current expectations; four, highly effective performance; and five, exceptional performance.

When making these evaluations, managers could consider the workers' track record in the data. Mangers could also consider their additional information about the workers' skills beyond what was recorded in the data. This information could come from overhearing calls — either purposefully or passively — or talking to workers about their strategies and experiences.

If managers have additional information, then they should give better-than-expected evaluations to unlucky workers whose metrics understated their skills. If measurement error is idiosyncratic rather than systematic, then unlucky workers should have better metrics in the subsequent months. Thus, if managers have additional information, their evaluations should predict workers' subsequent metrics conditional on their past metrics. If on-site managers have more additional information, their reviews should be more predictive.

We estimate how the manager's evaluation in December of 2018 predicts subsequent customer satisfaction, which is the metric that the reviews are largely aimed at encapsulating. Letting i index the worker; t, the date; z, the time-zone; and j, the worker's job title, we then have:

$$\text{Customer Satisfaction in } 2019_{i,t,z} = \psi_1 \text{Manager Evaluation in Dec } 2018_i + \psi_2 \text{Avg. Cust. Sat. July-Dec } 2018_i + \mu_{t,z,j(i)} + \epsilon_{i,t,z}. \quad (1)$$

The first column of Table 2 estimates this for all workers. Managers appear to have additional information about worker performance beyond the recorded metrics. Each additional point — which is equivalent to about one standard deviation in the distribution of evaluations — predicts a 0.053 standard deviation increase in subsequent customer satisfaction, holding constant past customer satisfaction (95% CI = [0.0060, 0.10]).

The second column limits the analysis to on-site workers. Consistent with on-site managers being able to passively overhead the calls of their on-site workers, managers' evaluations are especially predictive of subsequent performance on-site. By contrast, in the third column, which limits to remote workers, managers' evaluations are not significantly predictive of subsequent performance.

The fourth column evaluates the difference in predictive power of managers' evaluations between remote and on-site workers. The interaction term suggests that an additional point on the evaluation tends to translate into 0.08 standard deviations higher subsequent customer satisfaction for on-site workers relative to remote ones (95% CI = [-0.0043, 0.16]).

These differences in predictive power come from both sides of the distribution of perceived skill. Workers who received negative evaluations were more likely to do well if they were remote, consistent with their managers underestimating them.[23] Workers who received positive evaluations were more likely to do poorly if they were remote, consistent with their managers overestimating them.[24] Further, the distributions of evaluations in Appendix Figure A.2 are similar, suggesting that the differences in predictive power do not reflect differences in how managers are reporting performance.[25]

[23]On-site, workers who received below average evaluations tended to do 0.1 standard deviations worse in future customer satisfaction (p-value = 0.098). By contrast, among remote workers, those who received sub-par evaluations tended to improve by 0.066 standard deviations (p-value = 0.95).

[24]On-site, workers who received above average evaluations tended to do 0.16 standard deviations better in future customer satisfaction (p-value = 0.001). By contrast, among remote workers, those who received positive evaluations tended to regress by 0.040 standard deviations (p-value = 0.65).

[25]The granularity of the reviews seems to lead to similar distributions of evaluations — on-site managers tend to evaluate their workers more positively but the differences are marginal and do not impact the variation in evaluations given the gradations in possible scores. In the only other review cycle when call-center workers were evaluation, their performance was judged on a three point scale. The majority of workers were told they were meeting expectations. This limited the potential signal in these evaluations. Further, the distribution of evaluations was significantly different for on-site and remote workers. On-site workers were more likely to receive high marks and less likely to receive low marks as illustrated in Appendix Figure A.3. These differences in the distribution of reviews limits our ability to compare the information of on-site and remote managers. As a result, we decided ex-ante to focus exclusively on the winter cycle.

IV MODEL

This section models the market for remote work and shows how information frictions can lead to an under-provision of remote work. The friction in the model is informational; workers know more about their ability than firms. Over time, the hiring firm learns about worker ability, but only imperfectly. Particularly, some workers who merit promotion are overlooked. This happens less frequently in the office where it's easier to observe worker ability. Since better workers gain more being noticed, the office is complementary with ability. We ask how this complementarity affects the market for remote work.[26]

We illustrate this dynamic in a simple two-period model: each job lasts for two periods and features two possible tasks — one low-skill and one high-skill.[27] If workers are promoted to the high skill task in the second period, they receive a bonus. Jobs can be one of two types — on-site or remote. Each firm posts a menu of these jobs — one on-site and one remote that may differ in their wages and their bonuses. We assume firms must post bonuses that they will credibly fulfill in the second period given the returns to promoting a high-ability worker to a high-skill task. All firms have the same, additive production function and operate in competitive markets. Workers vary along two dimensions — ability and taste for remote work.

The timing of the model is as follows. In period zero, each worker receives a private signal about her ability, revealing the probability that she will merit promotion. Each firm posts a menu of jobs, designed in light of the asymmetric information between the worker and the firm. Each worker chooses a job on the basis of her idiosyncratic tastes and her likely ability. In period one, workers perform the low-skill tasks. At the end of period one, the firm learns that some workers are of high-ability and merit promotion. These signals are never misleading but are incomplete: some good workers will be overlooked, especially remote. In period two, workers execute their assigned tasks and receive bonuses if promoted to high-skill tasks. We assume each employment relationship lasts for exactly two periods — workers cannot quit and firms cannot fire workers.[28]

[26]One might imagine other sources of complementarity between the office and ability. Better workers may benefit more from what the bonds they form with their bosses or the acquisition of new skills. Alternatively, those with more drive may be more eager to go to the office and of higher ability. Appendix I.B discusses these generalizations of the model and ways to diagnose them empirically.

[27]Our stylized model features two-periods and two rungs of the career ladder. This is a good approximation of our empirical context where there are two tiers and a fairly fixed, short time to promotion. The richer N-period problem also collapses to this two-period case if the firm can only learn about a worker's suitability for the next role from her performance in the current one. Allowing information about a worker's performance to still be pertinent many roles later makes the problem more complex. We leave this richer case to future research that analyzes jobs with lengthier career ladders, in which workers may face changing incentives to work on-site to prove themselves.

[28]On the worker side, assuming away quits comes without a great loss of generality: given the importance of references, the opinion of the firm still matters for workers who quit before they can be promoted. On the firm side, assuming away firing is often a good approximation of reality given the costs of retraining, the losses of morale, and the potential legal headaches. In our context, workers become about 15% more productive in calls handled per hour over the course of their first eight months at the retailer, suggesting it is not always a good bet to replace experienced underperformers with new hires (see Figure 1).

IV.A The Worker's Problem

Workers vary in their tastes and abilities.

Each worker has an idiosyncratic taste for remote work, $\epsilon_i \sim \mathscr{L}(\mu, \mathfrak{s})$, where $\mathscr{L}(\cdot)$ is logistic. This taste could reflect the needs of the worker's children or the length of her commute to an on-site job. We assume this taste is orthogonal to ability.

Worker i's ability is either high or low, $\Theta_i \in \{H, L\}$. In period zero, each worker i receives a noisy signal, $\theta_i \sim U[0,1]$, that reveals the probability that she will be high-ability and merit promotion.

In each period, the worker's utility is quasilinear in her income and her utility from her work arrangement. We normalize worker utility from on-site work to zero.

Worker i's income in the first period is determined solely by the wage in her chosen job, j. Her income in the second period will also include a bonus, b_j, if she is promoted.[29] The worker will be promoted if her performance is both good and noticed. Thus, the promotion probability for worker i in job j is given by $\theta_i p_j$, where θ_i is the probability worker i performs well and p_j is the probability that her good performance is noticed in job j. Since this bonus accrues in the second period, its value is discounted according to the discount factor δ.

Worker i chooses job, $j \in \{o \equiv \text{on-site}, r \equiv \text{remote}\}$, to maximize expected utility over the two periods. Let $s_{ij1} \in \{L, H\}$ denote the firm's signal about whether to promote worker i at the end of the first period. We then have:

$$\max_{j \in \{o,r\}} u(j, \epsilon_i) + \delta \mathbb{E}[u(j, \epsilon_i, s_{ij1}) \mid \Theta_i, j] = \begin{cases} w_r + \epsilon_i + \delta\left(w_r + \epsilon_i + \theta_i p_r b_r\right) & \text{if remote} \\ w_o + \delta\left(w_o + \theta_i p_o b_o\right) & \text{if on-site} \end{cases}. \tag{2}$$

This implies a threshold rule for the prices, or wage penalties, under which worker i chooses remote work:

$$\textbf{Worker } \boldsymbol{i} \textbf{ Chooses Remote Work if:} \qquad P \equiv w_o - w_r \leq \epsilon_i - \frac{\delta}{1+\delta}(b_o p_o - b_r p_r)\theta_i$$

Let D_i denote worker's demand or threshold price for taking the remote job. We show in the firm's problem that the bonuses are the same in the on-site and remote jobs, $b_o = b_r$. Thus, we have:

$$D_i = \epsilon_i - \underbrace{\frac{\delta}{1+\delta} b(p_o - p_r)}_{\equiv \kappa} \theta_i. \tag{3}$$

[29] When taking the model to the data, we assume that the worker receives the bonus even if she leaves the firm: this matches the reality that the firm's opinion determines the quality of her references for the next job. The model also assumes that the gains from promotion are purely pecuniary. In reality, there may be additional gains from personal or social validation: these additional gains will magnify our model's mechanism.

There are two determinants of a worker's demand for remote work — her idiosyncratic tastes and her expected ability. A worker who has a stronger taste for working remote (higher ϵ_i) will have a greater demand for remote work. A worker who has is more likely to merit promotion (higher θ_i) will have a lower demand for remote work because its more likely that being overlooked will cost her a promotion. The weight on a worker's promotion prospects in her choice to be remote depends on the potential career penalty of remote work, κ, which reflects (i) the net present value of the bonus, $\frac{\delta}{1+\delta}b$ and (ii) the gap in the probability that merit will be recognized in a on-site and remote job, $p_o - p_r$

To make this more concrete, consider a worker ℓ who is certain that he will perform poorly ($\theta_\ell = 0$). Being overlooked never costs him a bonus because he never merits promotion. Thus, the promotion penalty is zero. At the other extreme, consider a worker h who is certain that she will merit promotion ($\theta_h = 1$). Being overlooked always costs her a bonus since she always merits promotion. Thus, her career costs of remote work are $\kappa \equiv \frac{\delta}{1+\delta}b(p_o - p_r)$. Since she must sacrifice more to go remote at any given price, the career cost of remote work shifts down her demand curve for remote work (as illustrated in Appendix Figure A.4). At any given price, a lower share of workers with the best possible prospects will choose remote work than of those with no prospects. Between these two extremes, the greater the chance that a worker will merit promotion, the greater the chance that she will lose from being overlooked. Thus, the career costs of remote work scale with a worker's expected ability, θ_i.

Since workers with poorer prospects lose less by being remote, choosing a remote job is a bad signal about a worker's ability. However, this signal is not fully revealing — some workers who are likely to merit promotion nonetheless choose remote work because of strong tastes. How much information is revealed by workers' choices depends on how much they hinge on workers' private information about ability versus their idiosyncratic tastes. The more important career concerns are, the more revealing these decisions will be; the more variable tastes are, the less revealing these decisions will be. The link between a worker's ability and her demand for remote work is the source of the selection problem: selection is more acute when career concerns loom large relative to preference heterogeneity and thus the link between a worker's ability and her demand for remote work is tighter. The next section turns to how this dynamic informs the menus set by the firm.

IV.B The Firm's Problem

Firms all share the same additive production function. Firms maximize profits and operate in competitive markets.

Firms assign workers to high- and low-skill tasks, $T \in \{L, H\}$. Let y denote the output of a low-ability worker in a low-skill task. High-ability workers ($\Theta_i = H$) have an absolute advantage in all tasks and a comparative advantage in high-skill tasks. High-ability workers produce an additional a units in low-skill tasks and an additional $A + a$ units in high-skill tasks, where $A > 0$.

The marginal product of promoting a high-ability worker to a high-skill task is A. Promoting a low-ability worker to a high-skill task, by contrast, costs the firm C.

Let τ denote the direct effect of remote work on productivity. This gives us the following production function for output Y generated by worker i in job j and task T in each period:

$$Y_{ijT} = y + \tau \cdot \mathbb{1}[j = \text{remote}] + a \cdot \mathbb{1}[\Theta_i = \text{H}] + \begin{cases} -C & \Theta_i = L,\ T = H \\ A & \Theta_i = H,\ T = H \end{cases}. \tag{4}$$

Firms have a couple of instruments for determining output. One, they choose whom to promote. Two, they set their initial menu of jobs.[30]

Each firm bases promotion decisions on its information about worker ability. The firm initially has no information about worker ability. Once the firm hires workers and employs them for a period, it observes that some workers are high-ability and merit promotion, $s_{ij1} = H$. The firm never mistakenly thinks that a low-ability worker is high-ability ($\forall\, i$ where $\Theta_i = L, s_{ij1} \neq H$), so promoting a worker who seems to be of high ability always increases output by A. However, the firm overlooks some high-ability workers ($\exists\, \Theta_i = H$ and $s_{ij1} \neq H$), so there are some other workers that the firm would like to promote but misses. However, to promote these workers, the firm would have to risk promoting low-ability workers and losing C. We assume that C is sufficiently high that firm only promotes worker whom it has learned are high-ability.

When setting its menu, each firm chooses the wages and bonuses for its on-site and remote job. We assume that each firm must post bonuses that it will want to fulfill in each period. In the second period, the firm would lose profits on each promotion if the bonus were any greater than A — the returns to promoting a high-ability worker to a high-skill task. Thus, the firm will only credibly fulfill bonuses of at most A.

This creates a commitment problem for a forward-looking firm. When setting its menu, the firm would like to post bonuses greater than A. Loading compensation into bonuses improves selection since workers who are more likely to succeed value bonuses more. Thus, the gains in worker selection from raising the bonus more than outweigh the lost profits from each promotion on the margin. However, if the firm sets the bonus above A, it would want to renege in the second period and not promote any workers. The firm would try to have its cake and eat it too. The firm would promise high bonuses to improve worker selection, but then refuse to pay out once good workers had been recruited to the firm. Foreseeing this, workers only find bonuses credible when the firm will want to fulfill them. The firm will consequently set the bonus at A, which leads to the best possible selection without posing the risk that the firm will renege.[31]

[30] We assume firms never fire workers after the first period because they either face legal constraints or prohibitively high retraining costs.

[31] Informational asymmetries may make it hard for the firm to commit to a high bonus. If the firm reneges, workers cannot prove their merit to either the court system or the pool of potential workers. Since it is difficult for workers to prove that the firm reneged on its promotion policy ex-post, it may be difficult for firms to commit to policies ex-ante. It's also worth noting that diminishing marginal utility (DMU) of income could lead to a similar result as imperfect

Now let's consider the firm's choice of base wages in the on-site and remote job. The firm will set its wages, such that it is indifferent between hiring workers into on-site and remote jobs. Thus, the expected profits from the two arrangements will be equal (and given competitive markets, will both be zero). Since the returns to promotions fully accrue to the worker with $b = A$, the problem collapses to the question of worker productivity in the low-skill task:

$$\mathbb{E}_j[Y_{i,j,T=L}] = y + a\,\Pr(\Theta_i = H \mid j) + \tau\mathbb{1}[j = r].$$

The difference in wages between the on-site and remote job — or the price of remote work set by the firm — will balance the difference in the average products:

$$P = w_o - w_r = \mathbb{E}_o[Y_{i,o,L}] - \mathbb{E}_r[Y_{i,r,L}] = -\tau + a\left[\Pr(\Theta_i = H \mid o) - \Pr(\Theta_i = H \mid r)\right]. \tag{5}$$

The first term reflects the treatment effect of remote work. The second term reflects the sorting of high-ability workers into on-site jobs, which creates the selection effect of remote work.

Figure 3 illustrates this market for remote work. The x-axis plots the share of workers who are working remotely and the y-axis plots the price of remote work. For simplicity, we consider the case where remote work has no direct productivity effect ($\tau = 0$). Thus, the marginal cost of switching a given worker from on-site to remote work is zero, pictured in dark green. However, the average costs of hiring a remote worker instead of an on-site one need not be zero since workers can sort between remote and on-site jobs on the basis of their private information about their ability.

The sorting of workers into remote and on-site work depends on the demand for remote work. The demand curve, drawn in yellow, aggregates individuals' demand from equation 3. Each individual's demand reflects her tastes and her likely ability. Consider the pool of workers who choose a remote job even at a high price (e.g. \$4/hr in Figure 3). A worker who knows she is likely to merit promotion (high θ) will have a lot to lose by taking a remote job and will only choose a remote job is if she has an extreme taste for remote work. By contrast, a worker who is less likely to merit promotion will need a less extreme taste to opt into remote work. Since tastes in the tails are less likely, a worker who is likely to do well will be less likely to opt into the remote job than a worker who expects to do poorly. As the price of remote work falls, workers who are likely to merit promotion need less extreme tastes to choose remote work, and, hence, the share of high-ability workers rises, so long as the taste distribution is unimodal. This causes the marginal product curve — illustrated by the light blue line of Figure 3 — to have a positive slope.

The slope of the marginal product curve depends on the sorting of high-ability workers across jobs and the returns to high-ability in the low-skill task (a). If tastes are more heterogeneous

commitment. DMU would make it costly for firms to load compensation into bonuses because workers would want consumption smoothing across the two periods and across the two states of the world, where they are seen as either high or low ability. These costs would limit the extent to which firms used bonuses to sort workers by ability. Thus, diminishing marginal utility of income would also lead firms to use the choice to be on-site or remote as a signal of ability.

(higher $\mathfrak{s}$), then workers will sort more on tastes and less on ability, causing the marginal product to be flatter. If career concerns loom larger in a worker's choice, then worker will sort more on ability and less on tastes, causing the marginal product to be steeper. The force of workers' career concerns depends on the potential career penalty from remote work, $\kappa = \frac{\delta}{1+\delta} b[p_o - p_r]$, and the variation in workers' signals about their likelihood of promotion, $\text{Var}(\theta) = \frac{1}{12}$. Thus, the slope of the marginal product curve is approximately:

$$\text{Slope of MP} \approx 2 \cdot a \cdot \frac{\kappa \text{Var}(\theta)}{\mathfrak{s}} = \frac{1}{6} a \frac{\kappa}{\mathfrak{s}}$$

When setting wages, the firm cares about the productivity of the average remote and on-site worker rather than those on the margin of this choice.

At each point, the pool of remote workers include both marginal workers and inframarginal remote workers, who choose remote work even when the remote wage is lower and fewer others choose remote work. Thus, the average output of remote workers (in orange) integrates the light blue line from left to right in Figure 3. Since the marginal product is rising, workers on the margin of remote work (in light blue) are always more productive than the average remote worker (in orange). Thus, opting into remote work means marginal workers pool with *less* productive workers.

Symmetrically, at each point, the pool of on-site workers includes both marginal workers and inframarginal on-site workers, who only choose remote work when the remote wage is higher and more workers choose remote work. Thus, the average output of on-site workers (in grey) integrates the light blue line from right to left in Figure 3. Since the marginal product is rising, those on the margin of on-site work (in light blue) are always less productive than the average on-site worker (in grey). Thus, choosing on-site work means marginal workers pool with *more* productive workers.[32]

To summarize, as the wage penalty — or price — of remote work falls, remote jobs become less adversely selected in keeping with classic selection models. At the same time, those who remain on-site become more advantageously selected. Thus, the average product in both remote and on-site work rise. As a result, the difference in average products — or the average cost of hiring a remote worker in navy — remains constant. This is approximately given by:

$$\text{AC} \approx a \cdot \frac{\kappa \text{Var}(\theta)}{\mathfrak{s}} - \tau = \frac{1}{12} a \frac{\kappa}{\mathfrak{s}} - \tau,$$

where the first term reflects the effect of worker sorting and the second term reflects the treatment effect of remote work.

[32] We can also think about this in a slightly different way. Let's read Figure 3 from right to left. The x-axis then becomes the share of the market choosing on-site work. The expected output of the marginal worker is decreasing in the share choosing on-site work. The average output of on-site workers integrates this marginal output so will be everywhere decreasing but also everywhere weakly above the marginal output.

In equilibrium, the price of remote work balances the difference in average products — or the average cost of hiring a remote worker instead of an on-site one. The market equilibrium will consequently be determined by the intersection of the average cost and the demand curve in Figure 3, leading to a market quantity of q_{mkt}. Since the equilibrium quantity is not determined by marginal cost, too few workers are remote, $q_{\text{mkt}} < q_{\text{eff}}$. The workers between q_{mkt} and q_{eff} only choose on-site work because of externalities from others' choices. These workers would prefer to be remote despite its potential career consequences. However, they choose to be on-site because the other on-site workers tend to be more productive. Thus, going on-site allows these workers to signal that they too are of high ability and deterred by remote work's career consequences. This signaling motive causes workers to enter a work arrangement that they prefer less and makes them no more productive. Thus, social surplus falls by the amount of the red shaded region.

Comparative Statics. The five forces driving the market for remote work are: (a) the treatment effect of remote work on productivity, (b) the average taste for remote work, (c) the variation in these tastes, (d) the potential career costs of remote work, and (e) the returns to ability in the low-skill task. Figure 4 illustrates the effect of perturbing each of these forces, focusing on the direction that will increase the market provision of remote work.

The first row of Figure 4 depicts changes in the market for remote work that make remote work more attractive through (a) productivity or (b) preferences.

Panel (a) pictures an increase in the treatment effect of remote work on productivity τ. This makes it less costly for the firm to hire remote workers and shifts the marginal and average cost curves down. This, in turn, increases both the efficient and equilibrium quantity of remote work. Since these quantities increase by approximately the same amount, the deadweight losses from adverse selection remain approximately constant. Thus, technologies that make remote work more productive increase its quantity but not the efficiency of its provision.

Panel (b) depicts an increase in the average taste for remote work, μ, which shifts out the demand curve for remote work. This outward shift in demand increases the market provision of remote work as well as the efficient allocation of remote work. Since these increases in quantity are of equal size, this shift does not impact the deadweight losses from adverse selection.

The second row of Figure 4 reveals what happens when the choice to accept a remote job relies more on tastes and less on ability.

Panel (c) considers what happens when tastes become more variable. First, the demand curve becomes more inelastic because when tastes are more diffuse, a given change in price leads to a smaller change in quantity. Holding costs constant (in the solid lines), making demand more inelastic would close the difference between the efficient and market quantity of remote work. When workers have more extreme tastes, they are less likely to choose an on-site job simply to pool with more productive workers.

Now, consider the cost curves. When tastes are more heterogeneous, it's more likely that a worker's choice to be remote reflects strong tastes rather than poor prospects. As a result, sorting into remote work starts to track ability less strongly, causing the marginal product to flatten and the average cost to fall towards the marginal cost. The narrowing gap between the average and marginal costs causes the market provision of remote work to increase towards the efficient provision, reducing the deadweight losses. Since the demand and cost curves act in concert to reduce the market inefficiency, greater taste variation is a powerful force to reduce deadweight losses from adverse selection.

Panel (d) considers what happens when the potential career costs of remote work fall because being on-site becomes less essential for being noticed for promotion (lower $p_o - p_r$) or there are lower returns to promotion (lower A and thus b). Let's first consider the effects on demand. Reducing the career costs of remote work is more valuable to workers who are more likely to merit promotion. Since these workers tend to be reluctant to enter remote work, they are concentrated to the right of the graph. Thus, the demand curve tilts up rather than shifts. Holding costs constant (in the solid lines), this tilt would increase the efficient provision of remote work by more than the market provision since the average costs exceed the marginal costs. Thus, on its own, this tilt in demand would increase the deadweight loss. However, this tilting reflects a weakening of the relationship between worker ability and the demand for remote work. When career concerns are weaker, workers' choices to be remote rely more on tastes and less on ability. This causes the marginal product to flatten and the average cost of hiring a remote worker to fall towards the marginal cost. As a result, the market provision of remote work rises and the deadweight loss shrinks.

The third row of Figure 4 imagines reducing the returns to ability in the low-skill task. This change makes it less costly for firms to attract workers of lower ability, reducing the average cost of remote work. This increases the market provision of remote work and reduces the deadweight losses. This implies that the market failure in remote work will be lower where latent ability matters less — either because workers' resumes are more useful in screening or because the tasks themselves are less sensitive to worker ability.

Covid-19 and these Comparative Statics. Given these comparative statics, it's worth considering the potential effects of Covid-19's massive experiment in remote work.

Remote work may have become more attractive if workers (and their firms) bore fixed costs of setting up home offices, learning new digital technologies, or adopting new working strategies. After bearing these fixed costs, remote work may be both more productive as in panel (a) of Figure 4 and more pleasant as in panel (b).[33] These changes would increase both the efficient and market quantity of remote work. Under these scenarios, the pandemic would have a persistent effect on the prevalence remote work but would not eliminate the market failure in its provisions.

The pandemic might have also made choosing remote work a weaker signal of worker ability.

[33] Workers may have also come around to remote work if their beliefs were initially too pessimistic — e.g. because they had only worked remotely when they were sick, had caregiving duties, or were running behind on a deadline.

The pandemic may have increased the variation in workers' tastes for remote work as workers learned more about their own preferences. During the pandemic, many workers experienced full-time remote work for the first time. As workers became more certain what their tastes were, some learned they liked it less than expected while others learned they liked it more. Thus, belief updating would suggest COVID-19 would increase preference heterogeneity.[34] An increase in preference heterogeneity makes choosing a remote job a weaker signal of ability as in panel (c) of Figure 4.

Prior to the pandemic, few workers experimented with remote work because they did not internalize the effects of their learning on others' choice sets. All workers may have been better off if everyone experimented with remote work because this would weaken the signaling incentive to be on-site and reduce the market inefficiency (as illustrated in panel c). However, few workers may have had an individual incentive to try out remote work given the costs of taking a remote job. By forcing all workers to learn about their tastes, the pandemic may have pushed the market into a new equilibrium where workers are more certain of their tastes, tastes are more heterogeneous, and choices to be remote are less indicative of low-ability.

The experience of COVID-19 may have also reduced the career costs of remote work if companies learned how to better assess the productivity of remote workers. This would increase the market provision of remote work and reduce the DWL as in panel (d).[35] In the initial equilibrium, firms may have under-invested in learning how to evaluate remote workers because they only considered the private returns and not the downstream effects on the feasibility of remote work.

In these cases, the initial market may have been in a "bad equilibrium" because of the externalities associated with experimenting with remote work.

Conclusions of the model. We consider what happens when the firm learns less about remote workers' abilities. This informational asymmetry is more costly for better workers who are more likely to merit promotion. Workers consequently sort into remote and on-site jobs based on private information about their ability, with workers who are likely to do poorly sorting into remote work and those likely to do well sorting into on-site work. Thus, a central prediction of the model is that remote workers will be adversely selected (and on-site workers will be advantageously selected).

Adverse selection into remote work causes the average cost of remote work to exceed the marginal cost and too few workers to work remotely. Workers on the margin of choosing remote work opt to be on-site because they do not want to pool with less productive workers. The choices of these workers are distorted by the externalities of others' decisions, leading to a deadweight loss in the market. The model's central implication is that adverse selection will lead to an under-provision of remote work.

[34]In ethnographic work, Ford et al. (2020) find evidence of this polarization in preferences toward remote work. In a pilot survey we have run, which we discuss in more detail in Section X, we find complementary results.

[35]However, Juhász et al. (2020) argue that management practices often take decades to adjust, judging from the historical transition to factory work, suggesting that the pandemic may (thankfully) be too short to transform the career costs of remote work.

The model, however, stresses that adverse selection into remote work is not immutable. This selection problem is less acute when workers' choices to be remote rely more on tastes and less on ability. When workers' tastes are more variable — for example, because of learning — this weakens adverse selection into remote jobs and ameliorates the market failure.

V IDENTIFICATION STRATEGIES: INTUITION

In the model, the costs of remote work depend on two factors. One, the treatment effect — or how remote work affects a given worker's productivity. Two, the selection effect — or the extent to which workers sort into remote and on-site jobs based on private information about their ability. To build intuition for how we identify these two central effects empirically, it's useful to consider the ideal experiment. This experiment would involve two random assignments: one of workers' choice sets — or offers at the time of hire — and the other, their actual job.[36]

In such an experiment, workers would be randomly offered either a remote job or an on-site job and could choose whether or not to accept the offered job. Of those who accepted the on-site offer, some would be randomly surprised with remote work while others would continue on-site. Similarly, of those who accepted the remote offer, some would be randomly surprised with on-site work. The resulting treatment cells are summarized in Figure 5.

The average cost of hiring a remote worker instead of an on-site one is determined by the productivity difference between workers who are offered a remote job and work remotely (cell 1) versus those who are offered an on-site job and work on-site (cell 4). The average cost is composed of the selection effect (in blue) and the treatment effect (in red) (consistent with the expression in equation 5).

To identify selection into remote jobs, we would focus on the sample in the first column of Figure 5: within workers who worked remotely, we would compare the productivity of those who accepted an offer for a remote job (cell 1) to the productivity of those who accepted an offer for an on-site job (cell 3). This would isolate the causal effect of offering remote work on the selection of workers who accept the offer, by limiting comparisons to those with the same actual job.

A central prediction of the model is that remote workers are adversely selected, so workers who select remote jobs (cell 1) are less productive than those who select on-site jobs (cell 3) when all workers are in the same actual job.

To identify the treatment effect of remote work, we would focus on the sample in the bottom row of Figure 5 who initially accepted an on-site offer: within this sample, we would compare the productivity of those who were remote to the productivity of those who were on-site (cell 3 versus 4). This would isolate the causal effect of working remotely on worker productivity holding

[36]Such an ideal design would be similar to Karlan and Zinman (2009)'s in the credit market: first, they randomly advertised different interest rates and then they randomly gave discounts to borrowers who initially had high rates. This allowed them to separately identify adverse selection into loans with high rates and the moral hazard these rates created for repayment.

selection constant, by limiting comparisons to those initially offered the same job.

In the model, this treatment effect determines the marginal cost of remote work and, thus, would be the socially efficient price, or wage penalty, for remote work.

We leverage three quasi-experiments at the retailer to approximate this ideal experiment.

First, the retailer introduced a program that let on-site workers apply to go remote in 2018. This program changed the choice set of new hires, while surprising existing on-site workers. For workers who ultimately work remotely, this allows us to estimate the selection effect of remote work, by comparing those who were offered a job with the potential to be remote to those who were offered a job that they expected to stay on-site (approximating the comparison of cells 1 and 3).

This program also allows us to identify the treatment effect of remote work. When on-site workers took up the opportunity to go remote, we observe their productivity both on-site and remote. This allows us to hold the offered jobs constant and compare productivity when the actual jobs differs. For workers who were surprised by the opportunity to go remote, this approximates being offered an on-site job and then switching from on-site to remote work (cell 3 versus cell 4); for workers who knew about the opportunity to go remote at the time of hire, this approximates being offered a remote job and then switching from on-site to remote work (cell 2 versus cell 1).

Second, beginning in July 2018, the retailer hired entry-level workers directly into remote jobs if they lived far from one of the retailer's call-centers. Thus, depending on where a worker lived, she would be either offered remote work and actually work remotely (cell 1) or be offered on-site work and actually work on-site (cell 4), yielding an estimate of the average cost of remote work.

Third, COVID-19 caused the retailer's call-centers to close. During the lock-down, all workers' actual jobs were remote regardless of the type of job they were offered at the time of hire, allowing us to identify the selection effect, by comparing the productivity of those who were initially offered remote and on-site jobs (cell 1 versus 3).[37] When the call-centers closed, we also observe workers who were initially offered on-site jobs transition from on-site to remote work. By comparing the change in their productivity to that of workers who were already remote, we isolate the treatment effect of remote work (cell 3 versus 4) net of the common shocks of COVID-19.

In addition to quasi-random variation in the offer of remote work, our setting features quasi-random variation in the offered price, or wage penalty for remote work. The retailer's uniform $14/hour remote wage compares more favorably to local on-site options in some places than others. This variation lets us estimate the elasticity of demand for remote work in our setting.

[37] The retailer only offers jobs that start remotely in places that are not proximate to its physical call-centers. This means that those who are offered on-site jobs are not offered fully remote ones and vice versa. The offer set consequently differs across workers.

VI OPPORTUNITIES TO GO REMOTE IN 2018

In the beginning of 2018, the online retailer began to run out of desks in some of its call-centers. To open a pressure valve, the retailer started posting remote job openings that were open to existing workers. Workers who took up these opportunities to go remote saw no change in their wage or in the content of their work.

VI.A ESTIMATING THE SELECTION EFFECT

For workers hired before January of 2018, these remote possibilities came as a surprise. When they accepted the job offer, these opportunities to go remote were unknown. These workers did not accept the job because of their demand for remote work.

By contrast, workers who were hired in the first few months of 2018 knew they could apply for remote work when they accepted the job offer. Thus, these workers may have accepted the job because of their demand for remote work.[38]

Other than their difference in initial information, these workers had similar experiences. They were all paid the same wages on-site and remote. They were all trained on-site.[39] They were all drawn from the same labor markets. They all used the same software, fielded calls coming from the same pool, and were given the same performance incentives both on-site and remote.

To identify the selection effect, we compare the productivity of remote workers who were effectively offered different jobs. Specifically, we compare the productivity of later cohorts who accepted job offers with remote possibilities (akin to cell 1 in Figure 5) to that of earlier cohorts who accepted jobs offers that they expected to stay on-site (akin to cell 3 in Figure 5). To hold the actual job constant, we focus on the time when these workers were remote.

While this seems to approximate the ideal experiment, a simple comparison of later and earlier cohorts could be misleading for a couple of reasons. One, at a given point in time, later cohorts would tend to be less experienced and, thus, might look spuriously less productive. Two, at a given level of experience, later cohorts would take calls on different dates: given the fluctuations in consumer demand, this could make later cohorts look spuriously more (or less) productive.

Jointly controlling for time and experience effects requires a control group. We leverage the pool of workers who were offered on-site jobs and chose to remain on-site. For these workers, the change in the menu was largely irrelevant. The introduction of an option they did not choose should not impact their productivity.[40]

[38] It's reasonable to suspect that these opportunities would have been widely known for a couple of reasons. One, the jobs were posted on external job boards for workers with experience at other companies. Two, fully 15% of the workers at these call-centers transitioned to remote work in the first few months of 2018 so word of mouth could have made this possibility known.

[39] The type of training may be an important consideration given the high returns to training in the call-center context: De Grip and Sauermann (2012) find that on-site training improves productivity by 10% in a field experiment.

[40] One caveat to this is that the retailer did not accept all applications for remote positions. Thus, some of the workers

We isolate the selection effect, using a difference-in-difference design. This design compares the change in productivity of remote workers in earlier cohorts and later cohorts to the same change among on-site workers. Since the introduction of opportunities to go remote was not staggered over time, we can operationalize the difference-in-difference design with a two-way fixed effects model. Letting t denote the day, $h(i)$ denote the month in which the worker was hired, $\ell(i)$ denote the location where she was hired, and $r(i)$ denote whether or not she is remote, we estimate:

$$\text{Calls/Hour}_{i,t} = \beta_{\text{Selection},1} \cdot r(i) \cdot \mathbb{1}[h(i) \geq \text{Intro of Remote Program}] + \beta_r r(i) + \mu_{t,h(i),\ell(i)} + \epsilon_{i,t}. \quad (6)$$

The fixed effects $\mu_{t,h(i),\ell(i)}$ isolate comparisons to workers who were answering calls on the same day and who had been hired in the same month and location. These fixed effects capture (a) fluctuations in consumer demand, t, (b) fixed differences across workers who were hired in a particular time, $h(i)$, and place, $\ell(i)$, and (c) the returns to experience that may be place-specific.[41]

The control for remote work $\beta_r r(i)$ reflects the productivity differences between workers who took up opportunities to go remote when surprised by this possibility (akin to cell 3 in in Figure 5) to their peers who continued on-site (akin to cell 4). This captures the treatment effect of remote work. For this to apply equally well to those who were offered remote and on-site jobs, the treatment effect of remote work must be homogeneous across these groups: Sections VI.B and VIII.B provide suggestive evidence to this effect.

The main coefficient of interest is $\beta_{\text{Selection},1}$. Interpreting $\beta_{\text{Selection},1}$ as the selection effect of remote work requires a parallel trends assumption — if there had been no change in the information in job offers, then the productivity of on-site and remote workers in each cohort would have proceeded in parallel.

This, in turn, requires (1) parallel time effects and (2) parallel experience effects. First, the random routing of calls suggests time-varying shocks to consumer demand should have parallel effects on remote and on-site workers. Second, all workers were trained on-site and recruited from the same labor-markets making it more plausible that any differences in on-the-job learning between remote and on-site workers would be small and constant over time.

Figure 6 presents the difference between on-site and remote workers for each cohort. This estimates the dynamic version of equation 6:

$$\text{Calls/Hour}_{i,t} = \sum_{h \in [\text{Nov 2017,March, 2018}]} \beta_{h,r} \cdot r(i) \cdot \mathbb{1}[h(i) = h] + \mu_{t,h(i),\ell(i)} + \epsilon_{i,t}. \quad (7)$$

The x-axis of Figure 6 represents the month in which the worker was hired and the y-axis represents the difference in calls per hour between remote and on-site workers from each cohort. The

who did not go remote may have wanted to do so. This would bias the selection estimate upwards since some of the workers in the control group may taken the job because of their demand for remote work within later cohorts. Further, the workers who were allowed to go remote may have been positively selected from a negatively selected group.

[41]Specifically, the interaction of calendar time, t, and hiring time, $h(i)$, controls for experience effects.

vertical dashed line indicates the introduction of opportunities to go remote. The three earlier cohorts to the left of the dashed line were offered jobs that they expected to stay on-site. The three later cohorts to the right of the dashed line were offered jobs that they knew could go remote. The sample excludes workers' first three months because productivity data for the first three months of the earliest cohort hired in October 2017 is no longer available.

In all three cohorts hired before the opportunities to go remote were introduced, those who were remote answered more calls per hour than their peers who remained on-site. Thus, the three blue points to the left of the dashed line all lie above zero. Pooling across these cohorts, the remote workers who had been surprised by this possibility answered 0.19 more calls per hour than their peers who continued on-site (approximating cell 3 versus cell 4 of Figure 5). Relative to the mean of 2.64 calls per hour, this represents a 7.0% increase (95% CI = [-4.5%, 18.6%]). While imprecisely estimated, this suggests that the treatment effect of remote work is positive. In the context of the model, this would suggest that workers should receive a wage premium to work remotely since it increases their productivity ($MC < 0$).

By contrast, in all three cohorts hired after the remote program was introduced, those who were remote answered fewer calls per hour than their peers who remained on-site. The three blue points to the right of the dashed line all lie below zero. Pooling across these cohorts, the workers who went remote after being offered a job that they knew had this possibility answered 0.31 fewer calls per hour than their peers who continued on-site. Relative to the mean, this represents 11.8% fewer calls per hour (95% CI = [-18.6%, -5.0%]).

This comparison contrasts workers who were offered remote work and actually worked remotely (cell 1) with those who were preferred on-site work and actually worked on-site (cell 4 of Figure 5). Thus, this 11.8% gap reflects the average costs of remote work to the firm.[42]

The lower productivity of these workers is particularly striking given the seeming advantage of working remotely, suggesting this reduced-form difference understates the true difference in latent productivity.

The difference in these differences isolates the selection effect of remote work net of the treatment effect of actually working remotely. As reported in column one of Table 3, this estimate of 0.50 fewer calls per hour suggests that being offered remote work reduced the productivity of new hires by 18.8% (95% CI = [-31.8%, -5.8%]). Even relative to the wide variability in calls handled, this represents a reduction of 0.35 standard deviations (95% CI = [-0.59, -0.11]).

The next two columns of Table 3 considers the robustness of this result to alternative controls. The second column includes more granular geographic controls for call-center rather than time-zone: these controls do not appreciably change the estimates on call per hour while increasing the R^2 by 7pp (from 41% to 48%).[43]

[42]Under the assumptions of our model, we can focus on the productivity of entry-level workers because the returns to promotion accrue to the worker rather than to the firm.

[43]It also worth noting that in theory, we could have isolated selection by analyzing cross-cohort productivity differ-

The third column includes a control for the time the worker spent on-site, which could affect the returns to experience among those who went remote. Consistent with the parallel trends of calls handled by experience in Figure 1, this does not appreciably affect the results.[44]

The final three columns of Table 3 consider the impact of offering remote work on the quality of calls as measured by customer satisfaction reviews. Our point estimates in the final three columns of Table 3 suggest that offering remote work decreased the quantity of calls without improving their quality. The point estimates suggest that offering remote work increased reviews by 0.025 to 0.040 points or one star per every 25 to 41 calls. While small in absolute terms, these effect are relatively large compared to the limited variation in reviews, ranging from 0.05 to 0.1 standard deviations of the outcome. Further, the confidence intervals do not rule out increases of one star for every 11 calls or a 0.23 standard deviation increase in customer satisfaction. This illustrates the challenges of drawing conclusions about customer satisfaction given the limited variability in reviews.

This design suggests that workers who accepted job offers for remote work were less productive than those who accepted job offers for on-site work. Even when workers were all working at home, earning the same wages, and drawn from the same labor markets, those who were offered remote jobs were less productive than those offered on-site ones. This is consistent with the prediction of our model that remote jobs attract worse workers who are less deterred by the career consequences of remote work. Further, this suggests that the average costs of remote work — estimated to be 11.8% in this design — diverge from the marginal costs of remote work — estimated to be -7.0%. This suggests too few workers will choose remote work relative to the socially efficient solution. The next section turns to estimating the treatment effect of remote work more directly.

VI.B Estimating the Treatment Effect

When the retailer introduced a remote work program in 2018, many on-site workers were pleasantly surprised and transitioned from on-site to remote work (approximating cell 3 to cell 4 in Figure 5). Later, workers accepted the retailer's on-site jobs because they could become remote ones: these workers had to start on-site before they could transition to remote work (approximat-

ences on-site (column two of Figure 5) rather than remote (column one of Figure 5). This alternative analysis suggests that offering remote work had no effect on worker selection (difference in call per hour of = 0.05 (se = 0.13)). There are a couple of possible explanations of this null result in light of our core finding. One possibility is that offering remote work attracts workers who only shirk at home. Sections VI.B and VIII.B offer suggestive evidence that the treatment effect of remote work is positive even for workers who are attracted to remote work, suggesting that this is not the right explanation in our context. Another possibility is that workers with a high demand for remote work have an added incentive on-site to try to secure this opportunity. We hypothesize this added incentive masked differences in selection when workers were on-site.

[44] One might worry that this linear control is insufficient, particularly since only workers in later cohorts can have short stints on-site (see Appendix Figure A.5). A natural approach would be to exclude workers with short stints on-site from the analysis. However, workers who go remote quickly are also those with the highest revealed demand for remote work. Thus, under our selection theory, we would expect the exclusion of these workers to mute but not eliminate the estimated effect. Indeed, replicating column two of Table 3 without workers who spend less than 10 weeks on-site suggests that offering remote work reduces calls per hour by 12.9%, 95% CI = [-27%, 1.5%].

ing cell 2 to cell 1 in Figure 5).

All of these workers had discretion over whether they went remote but not when they did so. The timing of workers' transitions were dictated by the time spent processing their applications and finding them openings on remote teams. Thus, the timing of the transitions to remote work were quasi-random. Under this assumption, an event study design identifies the effect of working remotely for those who switch holding their selection constant.

Letting e denote event time in number of weeks from the switch to remote work, we estimate:

$$\text{Calls/Hr}_{i,e} = \sum_{e \in [-6,6]} \tau_{ES,e} \mathbb{1}[e_i = e] + \epsilon_{i,e}, \tag{8}$$

where we index to workers' last week on-site before the switch to remote work. In addition to this dynamic specification, we also pool across weeks to estimate the productivity difference before and after workers' transitions to remote work within one and six week bandwidths, according to:

$$\text{Calls/Hr}_{i,e} = \tau_{ES} \mathbb{1}[\text{Post Transition}] + \epsilon_{ie} \quad \text{if } |e| \leq \text{Bandwidth}. \tag{9}$$

To insure that workers were fielding calls from the same pool both before and after their transition to remote work, we limit to workers who were in the same roles throughout the twelve-week time-span. These restrictions insure that workers were handling calls from the same pool and paid the same wages regardless of whether they were on-site or remote. To insure our comparisons hold the selection of workers constant, we limit to a balanced panel of workers who were at the retailer for the full twelve-week time-span.[45]

Figure 7 considers the hourly call volumes of the 120 entry-level workers who transitioned from on-site to remote work while continuing to take calls from the same pool. The x-axis represents the event time in weeks from the switch from on-site to remote work. The y-axis represents the workers' calls per hour. The vertical dashed line indicates workers' transition to remote work.

In the six weeks before workers' transition to remote work, workers' hourly calls were steady. Productivity (in the black circles) was consistently close to the reference week (in the horizontal dashed line) and the grey 95% error ribbon includes this reference productivity level.

In the first week after workers' transition to remote work, productivity jumps up by 0.20 calls per hour (see column one of Table 4). Compared to the dependent mean of 2.93 calls per hour, this represents a 6.9% increase in calls handled (95% CI = [1.3%, 12.5%]). The composition of workers did not change, neither did the composition of calls. Instead the same workers became more productive when they were working at home rather than on-site. During the subsequent five weeks, productivity hovered around this higher level, with all the confidence intervals lying

[45] Particularly, of the 247 entry-level workers who transition to remote work in 2018, we exclude 40 who were not at the retailer in the full twelve weeks, 47 who were still in training during the first six weeks, and 40 who changed job titles.

above the reference level (in the horizontal dashed line).

Pooling across the six weeks before and after the transition to remote work, calls per hour increases by 0.22 or 7.5% of the dependent mean (95% CI = [3.1%, 11.9%]). These productivity changes within workers complement the productivity differences across workers within the early cohorts in Figure 6: those who went remote but were surprised by this opportunity answered 6.9% more calls per hour than their peers who continued on-site.

When workers transitioned to remote work, their additional calls did not come at the expense of lower customer satisfaction which remained stable around this transition (as detailed in columns two and three in Table 4). Transitioning to remote work also suggestively improved workers' reliability, with unapproved absent minutes falling by 2 to 4 minutes per day. This reflects a meaningful decrease relative to the base of 10 minutes per day. However, this effect is measured imprecisely.

The changes in calls per hour are strikingly similar across workers who were offered different jobs at the time of hire. Of the 120 workers, 71 knew about the opportunity to go remote when they were offered the job and the other 49 were surprised by this possibility. In their first six weeks working remotely, the 71 workers who were initially offered the opportunity to go remote answered 0.25 more calls per hour (se $= 0.0812$) than they had in their last six weeks on-site. Similarly, in their first six weeks working remotely, the 49 workers who were initially offered a purely on-site job answered 0.21 more calls per hour (se $= 0.102$) than they had in their last six weeks on-site. The difference in these differences of 0.04 calls per hour suggests the treatment effect of remote work is similar for workers who were offered remote work and those who were offered on-site work. However, this strategy cannot rule out large differences between these effects (the 95% CI spans from -0.21 to 0.30 calls per hour). We return to the question of the similarity of these effects when we analyze the natural experiment created by COVID-19 in Section VIII.B.

These findings suggest that call-center work is conducive to remote work. Holding the set of workers fixed, the firm would like workers to work remotely. However, our model suggests that offering remote work may attract workers who are less skilled and less deterred by the career consequences of being out of earshot of one's manager. The previous section bears out this prediction empirically: when the firm started to offer remote work, less productive workers took the job. In our setting, this adverse selection from offering remote work more than outweighed the positive treatment effect of having workers work remotely. Thus, on average, having workers work remotely was costly to the firm, even though, on the margin, it would be more productive for a given worker to work remotely. This suggests that the market price of remote work may penalize remote work even though the efficient price would promote it. The next section turns to estimating the average cost of remote work on a larger sample of remote and on-site hires.

VII HIRING REMOTE WORKERS DIRECTLY: JULY 2018 - COVID-19

In July of 2018, the retailer expanded its remote work program by posting ads for entry-level hires in labor markets where they did not have on-site call-centers. Depending on where workers lived, they would be offered either remote jobs or on-site jobs at the retailer.[46] As external recruiting replaced internal transfers into remote jobs, the retailer's workers soon fell into one of two types: either they were offered remote work and actually worked remotely (as in cell 1 of Figure 5) or they were offered on-site work and actually worked on-site (as in cell 4 of Figure 5). Between July of 2018 and April 6, 2020 when the retailer closed its call-centers due to the onset of COVID-19, the retailer hired 442 remote workers and 834 on-site workers at the entry wage of $14/hour.

The productivity gap between remote workers in cell 1 and on-site workers in cell 4 includes the average cost of remote work — which in turn captures both the treatment and selection effect of remote work as pictured in Figure 5. For this comparison to solely reflect the impact of remote work, the workers offered remote jobs and those offered on-site ones must be otherwise equivalent.

We compare remote and on-site hires who handled calls on the same day (t), were in the same time-zone (ℓ), and were hired in the same month (h), according to:

$$\text{Calls/Hour}_{i,t} = \beta_{ac,r} r(i) + X_i' \alpha + \mu_{t,h(i),\ell(i)} + \epsilon_{i,t}. \tag{10}$$

This identifies the average cost of remote work if two assumptions hold. One, the labor markets where the retailer offered remote work were similar to those where it offered on-site work conditional on the observable labor-market characteristics in X_i. Two, X_i captures the observable, relevant traits of workers that the retailer could use to screen workers.

Ex-ante, it's reasonable to suspect that the locations of the retailer's low-wage call-centers were conditionally random. The retailer based its location decisions on labor market indicators it received from Emsi, a statistical modeling company. We control for these same indicators when comparing the productivity of on-site and remote workers. This conditional comparison will reflect the average cost of remote work so long as the retailer's locational choices do not reflect additional, relevant information about labor markets that we cannot observe.[47]

If the labor markets where the retailer built call-centers were more productive than other observably similar labor markets, then we would expect the reduced-form difference in worker productivity between on-site workers and remote workers to be much larger than the average cost of remote work within each labor market. This, in turn, would suggest that $\hat{\beta}_{ac,r}$ would be much larger than the average cost of offering remote work within its on-site labor markets, which we estimated to be 11.8% in the previous section. If instead the labor markets where the retailer built its

[46] This policy change led to many more of the retailer's workers being hired directly into remote jobs, as illustrated in Figure A.6.

[47] When we've learned about idiosyncratic reasons that the retailer located call-centers in a labor market — e.g. being a founder's hometown — they seem plausibly orthogonal to local determinants of productivity.

call-centers were about as productive as observably similar labor markets, then we would expect $\hat{\beta}_{ac,r}$ to align with the average cost of offering remote work within a labor market.

A goal of this section is to test this in order to understand whether comparing workers drawn from different labor markets can reveal the effects of remote work. Drawing these comparisons is useful for two reasons: one, they can yield more precise estimates of the average cost of remote work based on larger samples of workers; two, together with the the shock of COVID-19, such comparisons can separately identify the treatment and selection effects of remote work as we investigate in the next section.

Table 5 details the productivity difference between remote and on-site workers. The first column considers the raw productivity difference. Remote workers answered 0.32 fewer calls per hour than their on-site peers who were hired in the same month, paid the same wage, and answering calls at the same time. Compared to the dependent mean of 3.1 calls per hour, this represents 10.2% fewer calls per hour (95% CI = [-15.5%, -4.9%]).

If the retailer only recruited remote workers from low-wage labor markets that were similar its on-site call-centers, it might be able to narrow the productivity gap. The second column includes a control for the median pay in customer service within each worker's metropolitan statistical area as estimated by Emsi. Each additional dollar in workers' local alternatives correlates with 0.034 fewer calls answered per hour (1.1% fewer). Since the retailer's remote workers have better alternatives on average than its on-site workers, the resulting difference in productivity is slightly smaller at 8.4% (95% CI = [-14.3%, -2.52%]).

The quality of a worker's outside option may depend on the prevalence of customer service work as well as the prevailing wage. In places where there are more customer service jobs, the retailer may find it more difficult to recruit and retain skilled workers. Indeed, in column three of Table 5, for every 1pp increase in the share of prime-aged people in customer service, workers answered 0.027 fewer calls per hour (0.9% fewer). Since the retailer's remote workers have more prevalent alternatives than on-site workers on average, the resulting estimate suggests that remote workers were 7.6% less productive than on-site workers.

These analyses suggest that the retailer could improve the selection of remote workers by targeting recruitment at particular types of labor markets. However, this strategy would only go so far in closing the productivity gap. Further, these estimates are similar to the average cost of remote work of 11.8% (95% CI = [-18.6%, -5.0%]) estimated in the previous section. Since this estimate compared workers within a labor market, the concordance of these results suggests that the retailer did not locate its on-site call-centers in latently more productive labor markets.

Another way to think about this exercise is to imagine hiring managers at the retailer leafing through workers' applications. These hiring managers could choose whether to screen out potential hires based on the characteristics of their local labor markets. The question then becomes, "Can these hiring managers fix the selection problem by only hiring remote workers from labor markets that look like those of the retailer's low-wage call-centers?" The answer to this appears to be "no". Even conditional on the observable characteristics of the local labor markets, remote workers are less productive than on-site workers. Consistent with our model, the selection effect appears to be inherent to the choice to be remote.

A similar line of reasoning would suggest that it might be possible to screen on other observable worker attributes, such as worker experience. Call-center work at the retailer is an entry-level job. As a result, many workers have limited relevant prior experience. Consistent with this, older workers are no more productive on average than younger workers. Gender is also a weak predictor of performance. Thus, even after conditioning on the available worker characteristics in column five of Table 5, remote workers are 7.7% less productive than on-site workers.[48]

It's worth noting that remote workers answer fewer calls because they take about a minute longer on each one (see column one of Table A.3). But longer calls do not appear to make for more satisfied customers: customer satisfaction is not statistically nor economically different between remote and on-site workers. The point estimates in Table A.4 suggest that remote workers average between 0.004 and 0.010 better reviews. This translates into earning one more star in every 100 to 250 calls. Further, the 95% confidence intervals of these estimates suggest it is unlikely that remote workers earn more than one extra star per 42 calls or lose more than one star every 79 calls.

The reduced-form gap in productivity is particularly striking given the advantage that working remotely seems to confer on worker productivity. Given a positive treatment effect of remote work, these reduced-form differences understate the latent productivity differences between workers offered remote and on-site jobs. Using the estimated 7.5% productivity advantage of working remotely estimated in Section VI.B, the 7.6% to 10.2% reduced-form difference suggests remote workers would be 15.1% to 17.7% less productive than on-site workers if they were all working in the same actual job.

COVID-19 allows us to see the selection into remote work more directly.

VIII Shock of COVID-19: All Workers Work Remotely Regardless of Initial Offer

On April 6, 2020, COVID-19 caused all of the retailer's call-centers to close their doors and send their workers home.

VIII.A Estimating the Selection Effect

During COVID-19's lockdown, workers who had initially been offered on-site jobs were working remotely(as in cell 1 of Figure 5) alongside those who had been initially offered remote jobs (as in cell 3 of Figure 5). A comparison between these workers captures selection into remote work

[48]These results are all qualitatively similar when we use the full sample of on-site workers rather than only those with $14/hour entry-level wages.

but not the treatment effect of working remotely. Thus, the lockdown offers a unique opportunity to isolate selection into remote work. We compare worker productivity among those with base wages of \$14/hour:

$$\text{Calls/Hour}_{i,t} = \beta_{\text{Selection},2}\mathbb{1}[\text{Offered Remote Work}_i] + X'_{i,t}\alpha + \mu_{t,h(i),\ell(i)} + \epsilon_{i,t} \tag{11}$$

during the lockdown. For $\beta_{\text{Selection},2}$ to identify the selection effect of remote work, three assumptions are necessary. One, remote and on-site workers must be drawn from conditionally similar labor markets. Two, remote and on-site workers must be similarly exposed to the productivity consequences of the pandemic. Three, remote training must be as effective as on-site training.

Figure 8 illustrates the distributions of call volumes of the retailer's workers during the lockdown. In this sample, all workers worked from home, earned the same wages, and took calls from the same pool. But some had initially been offered remote jobs (in blue) and others had initially been offered on-site jobs (in orange). During the lockdown, those who had initially been offered remote jobs answered 0.62 fewer calls per hour than those who had been initially offered on-site jobs (95% CI = [-0.88,-0.35]), as detailed in column one of Table 6. Compared to the mean calls per hour of 3.4 during the pandemic, this represents 18.1% fewer calls per hour (95% CI = [-25.9%, -10.2%]).[49]

The remaining columns of Table 6 consider the robustness of the results to the inclusion of controls. The selection estimates range from -0.52 to -0.62, which represent 15.3% to 18.2% of the dependent mean.

The controls in the second column consider the role of local pandemic conditions — particularly, the number of COVID-19 deaths per hundred thousand and number of cases per ten thousand in the worker's home county during the week that she was handling calls.[50] A one standard deviation increase in deaths is associated with a 2.0% decrease in productivity (95% CI = [-5.3%,1.3%]). While suggestive of a negative effect of the pandemic on worker productivity, the estimates on deaths and cases are not jointly significant (the p-value of the Wald test is 0.28 in this specification and 0.14 after including the full battery of controls). Further, these controls appreciably impact the estimated effect of offering remote work on worker productivity.[51]

The third and fourth columns control for the local labor market conditions before the pandemic. During the lockdown, workers might have been unable to turn their usual alternatives into actual job offers. The relevant alternative during the pandemic was non-participation for many workers.[52] Better expected alternatives after the pandemic might it more costly to leave the labor

[49]The dependent mean is higher than in previous analyses because the online retailer saw an uptick in consumer demand during the pandemic as consumers switched from brick-and-mortar to online shopping.

[50]This uses data from *The New York Times* on `https://github.com/nytimes/covid-19-data`.

[51]This relatively weak relationship between the intensity of the pandemic and worker productivity may reflect a couple of features of this setting. One, workers were working at home and thus not directly exposed to the risks of the pandemic during their working hours. Two, during these early months of the pandemic, many public-health precautions, such as school closures, were enacted across the country regardless of the local case count. The uniformity of policy may have limited the scope for the local pandemic to impact workers' day-to-day experiences.

[52]Given the disparate impacts of the pandemic across men and women, this might be especially true in this female-

market. Consistent with this, each additional dollar that the outside options in a worker's labor market usually pay correlates with 0.067 more calls answered per hour or a 1.9% increase (95% CI = [-0.30%, 4.2%]).

The fifth column introduces individual level controls to account for differential effects of the pandemic for women and men and younger and older workers.

These controls may be insufficient to handle the differences in care responsibilities that workers faced during the pandemic. In June of 2020, the retailer conducted a survey about workers' caregiving responsibilities. Table A.5 details the impact of including these controls. Within the sample who complete the caregiving survey, including the control for caregiving has a barely detectable impact on the estimated productivity gap between remote and on-site workers.[53]

As in the previous section, remote workers did not compensate for lower quantity with appreciably higher quality (see Table A.6). Once we condition on the observable characteristics of workers, remote workers averaged 0.009 higher reviews than formerly on-site workers or one additional star per every 111 calls during the pandemic (95% CI = [-0.026, 0.044]).

This analysis suggest that those who were offered remote jobs were less productive than those who were offered on-site jobs even when all workers were working at home, taking similar calls, and earning the same wages. This complements the natural experiment from Section VI.A that utilized the introduction of opportunities to go remote to identify the effect of offering remote work on worker selection. These estimates are not only the same directionally but also similar quantitatively with an estimated selection effect of around -18% in both contexts. This concordance is striking because these estimates are based on independent samples and different identifying assumptions. Thus, these designs provide independent evidence that workers who choose to work remotely are less able than those who choose to work on-site.

This bears out a prediction of our model that more productive workers will be more deterred by the potential career consequences of remote work. This further suggests that workers will likely face a wage penalty for remote work even if it makes them more productive since taking a remote job signals that a worker is less productive.

VIII.B Estimating the Treatment Effect

The pandemic closures of the retailer's call-centers on April 6, 2020 also shed light on the treatment effect of remote work. In the months leading up to April 6, 2020, workers who had been initially offered on-site jobs were working on-site (as in cell 4 of Figure 5). In the months after April 6, 2020, these workers worked remotely (as in cell 3 of Figure 5).

dominated job.

[53]Caregivers were slightly more productive than non-caregivers during the lockdown — consistent with the pattern that we see before the pandemic. Since remote workers were more likely to be caregivers than on-site workers in this sample (52.5% versus 45.1%), the inclusion of this control marginally increases the estimated selection effect of offering remote work.

The productivity change of formerly on-site workers before and after April 6, 2020 captures the treatment effect of remote work (τ_o). This change, however, also captures the personal shocks of the pandemic (γ_o), the increased demand for the online retailer (ρ_o), and the spillovers from others' productivity changes (σ_o). The change in productivity of on-site worker o is then given by:

$$\Delta\text{Calls/Hour}_o \equiv \text{Calls/Hour}_{o,\text{Post}} - \text{Calls/Hour}_{o,\text{Pre}} = \tau_o + \gamma_o + \rho_o + \sigma_o.$$

When the offices closed, some workers at the retailer were already working remotely. These workers were exposed to the same spurious shocks but saw no change in their work arrangement (they instead remained in cell 1 of Figure 5). For each remote worker r, the change in productivity around the office closures is then given by:

$$\Delta\text{Calls/Hour}_r \equiv \text{Calls/Hour}_{r,\text{Post}} - \text{Calls/Hour}_{r,\text{Pre}} = \gamma_r + \rho_r + \sigma_r.$$

Thus, to isolate the treatment effect of remote work, we compare the productivity changes of on-site workers forced into remote work to that of already remote workers who saw no change in their work arrangement. The difference in these differences gives us:

$$\Delta\text{Calls/Hour}_o - \Delta\text{Calls/Hour}_r = \tau_o + (\gamma_o - \gamma_r) + (\rho_o - \rho_r) + (\sigma_r - \sigma_o). \tag{12}$$

Our identifying assumption requires similar spurious shocks across remote and on-site workers. Our identifying assumption relies on three components: (a) similar shocks of the pandemic ($\mathbb{E}[\gamma_o - \gamma_r] = 0$), (b) similar responses to the uptick in consumer demand ($\mathbb{E}[\rho_o - \rho_r] = 0$), and (c) similar spillovers from changes in others' productivity ($\mathbb{E}[\sigma_o - \sigma_r] = 0$).

We estimate this difference-in-difference according to:[54]

$$\text{Calls/Hour}_{i,t} = \tau_{\text{DiD}} \cdot \mathbb{1}[\text{On-Site Offer}_i] \cdot \mathbb{1}[t \geq \text{April 6, 2020}] + \psi \cdot \mathbb{1}[\text{On-Site Offer}_i] + \mu_{t,h(i),\ell(i)} + \epsilon_{i,t} \tag{13}$$

where i indexes the worker and t indexes the date; $h(i)$ denotes the month worker i was hired and $\ell(i)$ denotes the time-zone in which she works.

Figure 9 illustrates this difference-in-difference design. The x-axis represents the week of work. The y-axis represents calls taken per hour. The vertical, red line indicates the office closures on April 6, 2020. To the left of this red line, workers who had been offered on-site jobs (in the solid line) were working on-site; to the right, they were working remotely. On both sides, workers who had been offered remote jobs (in the dashed line) were working remotely.

In the seven weeks before the office closures, on-site workers tended to answer more calls than remote workers. Pooling across these weeks, on-site workers averaged 0.20 more calls per hour than remote workers (95% CI = [-0.012, .42]). Compared to the pre-period mean of 2.6 calls per

[54] Note this two-way fixed effect design does not introduce negative weights because the office closures were not staggered over time.

hour, this represents 8.0% more calls taken per hour (95% CI = [-0.48%, 16.4%]). This is in line with our results in Section VII: in the cross section, on-site workers took more calls than remote workers.

In Section VI.B, the productivity of those who opted into remote work rose sharply when they started to work remotely. If these productivity gains also extend to workers who initially chose to be on-site, then the productivity gap would widen when on-site workers started to work remotely. Indeed, when the office closed and on-site workers went remote, the productivity gap widened. In the seven weeks after the office closures, formerly on-site workers answered 0.57 more calls per hour than already remote workers. Compared to the post-period mean of 3.75 calls per hour, this represents 15.2% more calls per hour (95% CI = [7.2%, 23.3%]).

The difference in these differences, also reported in the first column of Table 7, suggests that remote work increased the calls per hour of formerly on-site workers by 0.37 calls per hour or 9.8% (95% CI = [2.8%, 16.7%]).

Under our identifying assumption, this strategy indicates remote work caused those originally on-site to become 9.8% more productive.

We consider two ways to relax our assumption of similar shocks of the pandemic ($\mathbb{E}[\gamma_o - \gamma_r] = 0$). One, we consider the window in which the treatment effect is estimated. Two, we include controls for proxies of workers' exposures to the stressors of the pandemic.

The retailer closed its call-center on April 6, 2020. This was three weeks after most American schools shut their doors on March 16, 2020. Focusing on the six weeks around the office closures holds school policies constant. Differential effects of school closures on on-site and remote workers would then be netted out in the difference-in-difference. In theory, this would hold an important component of the pandemic's effects constant in the comparison. Zeroing in on this narrow bandwidth in the second column of Table 7 results in a larger estimated treatment effect of remote work of 0.68 calls per hour or 17.6% compared to the mean in this post-period (95% CI = [10.7%, 24.5%]).

However, in the weeks immediately leading up to the office closures, other impacts of the pandemic may have been differentially born by those in the office. During this time, working in the office may have been draining because of the fear of catching the virus and the safety precautions to avoid its spread. Indeed, in Figure 9, on-site workers became less productive relative to those already working remotely in the weeks leading up to the closures.[55] Thus, it might be reasonable to exclude the three weeks prior to the office closures from the analysis in order to compare productivity in remote work to a more typical baseline in the office. The resulting donut design in the third column of Table 7 suggests that remote work increased the call volumes of formerly on-site workers by 0.28 calls per hour or 7.6% (95% CI = [-0.002%, 15.4%]).

[55] Between March 16, 2020 and the office closures, on-site workers took only 4.8% more calls per hour than remote workers. By contrast, in the previous four weeks, on-site workers took 10.7% more calls per hour than remote workers, a significantly greater gap in calls (p-value on the difference = 0.072).

We then include controls for proxies of workers' exposures. The fourth column of Table 7 includes controls for COVID-19 deaths and cases in each worker's county during each week they were taking calls. As in the previous section, more deaths translated into fewer calls but this was not a strong predictor of productivity.[56]

The fifth column of Table 7 controls for workers' childcare responsibilities as reported in a June 2020 survey conducted by the retailer. The results suggest school and daycare closures reduced parents' productivity in April and May relative to late February and early March. Since these burdens were differentially born by remote workers, controlling for them reduces the estimated treatment effect. However, since the differences in child-care responsibilities are muted within this sample of entry-level workers —- 53% of remote workers and 44% of on-site workers care for children — the inclusion of this control has a marginal effect on the point estimate, reducing it from 0.29 to 0.28 or 7.3% (95% CI = [-3.6%, 18.2%]). Similarly controlling for differential impacts of the pandemic along gender and age do not appreciably affect the estimated treatment effect, as reported in Table A.7.[57]

Our second identifying assumption is that remote and on-site workers responded similarly to the uptick in consumer demand at the online retailer ($\mathbb{E}[\rho_r] = \mathbb{E}[\rho_o]$). During the onset of the pandemic, consumers switched from brick-and-mortar shopping to online retailers like this one. If remote and on-site workers responded differently to the uptick in consumer demand, then the widening gap in productivity could be due to these differential responses and not the causal effect of working remotely. To assess this possibility, we consider the cross-sectional patterns and the case study of the 2020 holiday rush.

In the cross-section, upticks in consumer demand before the pandemic did not systematically lead to a greater productivity gap between remote and on-site workers. On days when workers averaged 1.2 more call per hour as they did in the six-weeks after the pandemic, the gap between on-site and remote workers only tended to be 0.034 calls per hour greater (95% CI = [-0.057, 0.12]). The treatment effect estimated from the difference-in-difference lies outside the bounds of this confidence interval, suggesting that the increases in call volumes alone did explain the rising gap in call volumes when on-site workers went remote.

To consider a case study, Figure A.7 depicts the call volumes of on-site workers (in the solid line) and remote workers (in the dashed line) during the tail end of the 2020 holiday rush. During this time, call volumes declined by more than one call per hour. Despite this substantial change in call volumes, the difference in productivity between remote and on-site workers (in maroon) stays

[56]Further, the remote and on-site locations had fairly similar exposures to the pandemic during this six-week window — with 17.2 deaths per hundred thousand in the counties of the formerly on-site workers and 18.9 deaths per hundred thousand in the counties of the already remote workers.

[57]When we consider heterogeneity in the treatment effects, the results suggest the effects of remote work do not vary by worker age. Women's productivity improved when they worked remotely but less so than men's: however, these results are merely suggestive given the limited number of men in our sample. Similarly, caregivers' productivity appears to improve from remote work but less so than the productivity of those without these responsibilities. These results again must be viewed with caution given the limited number of non-caregivers in our setting.

quite flat with no economically or statistically significant change.

Our third identifying assumption is that remote and on-site workers responded similarly to the spillovers from the changes in others' productivity ($\mathbb{E}[\sigma_r] = \mathbb{E}[\sigma_o]$). In the call-center context, when one worker answers a call, another worker cannot receive it. If calls do not arrive at the retailer faster than they can be taken, a productivity increase of one worker will cause delays for other workers. Since calls are randomly routed between on-site and remote workers, both treatment and control workers experience these delays, so $\sigma_r \approx \sigma_o$ up to a first-order. The impact of the spillover on a particular worker depends on how frequently she appears in the queue and how quickly a replacement call arrives. On-site workers answer more calls and thus could be exposed to more delays as they return to the queue more frequently. These differential spillovers would tend to bias the estimated treatment effect downwards since on-site workers would be more frequently delayed by the increased productivity of other on-site workers. However, given the high volume of calls during the pandemic, new calls arrived nearly as fast as they could be answered, mitigating any potential bias.

Our analysis indicates that there is a positive causal effect of remote work on worker productivity. This complements our findings in Section VI.B, which found sharp increases in worker productivity around individuals' voluntary transitions from on-site to remote work. Together, these analyses suggest that working remotely causes workers to become more productive in this setting, both for workers who choose to work remotely and those who do not. Further, the point estimates are quite similar across the two analyses offering suggestive evidence that workers who opt into remote work gain similarly from working remotely as those forced to work remotely.

A natural question is where the productivity gains come from. Particularly, do workers spend more of their time on the phone when they work remotely? Or do they answer calls more quickly? Difference-in-differences suggest that working remotely increases workers' time spent on the phone but does not cause workers to handle calls more quickly as illustrated in Figure A.8. The patterns are potentially consistent with remote workers spending more time multi-tasking, allowing them to spend time on the phone but making each minute marginally less efficient.[58] The increases in call quantity also did not appear to come at the expense of customer satisfaction as detailed in Table A.8.[59]

[58] Figure A.8 separately considers workers' time spent on the phone (in the top panel) and the average time they spent on each call (in the bottom panel). The red dashed line indicates the office closures on April 6, 2020. Before the offices closed, on-site and remote workers spent similar amount of time on the phone. In the top panel, the solid and dashed black lines lie on top of one another to the left of the red line. After the offices closed, on-site workers spent a greater share of their working time on the phone than remote workers. To the right of the red line, the solid line lies consistently above the dashed line. The difference in these differences suggests that working remotely increased the percent of time that formerly on-site workers spent on the phone by 7.2pp (95% CI = [5.3,9.2]). Before the offices closed, on-site workers answered more calls because they handled each one more quickly. To the left of the dashed line, the solid line lies consistently below the dashed line in the bottom panel. After the offices closed, on-site workers continued to answer calls more quickly but the gap in speed did not widen and instead, if anything, closed. The difference in these differences suggests that working remotely increased the duration of formerly on-site workers' calls by 13.2 seconds (95% CI = [-7.6, 34.5]).

[59] The point estimates suggest remote work led to between a 0.021 decrease in average customer reviews and an 0.032 increase in these reviews. This amounts to one fewer star in every 47 calls or one more star in every 31 calls.

IX PRICE VARIATION

The retailer advertised remote call-center jobs at a uniform \$14/hr wage on national job boards. The uniformity in this wage contrasted sharply with the heterogeneity in workers' local outside options. Together, this generated variation in the implicit price workers had to pay to work at this remote job rather than a local on-site alternative.

In cities like Dallas TX, where the retailer's \$14/hr wage was far below the going rate for similar on-site work, workers had to pay a high implicit price to take this remote job. By contrast, in places like Lufkin, TX, a couple hours from Dallas, the retailer's pay exceeded many more of the on-site alternatives: thus, even workers who preferred the office might take this remote job.

If workers' demand for remote work were elastic because workers' preferences were homogeneous, the retailer would hire many more workers from Lufkin than Dallas. If demand were inelastic because preferences were heterogeneous, the gap in recruitment would be smaller.

This logic suggests that the ability of the retailer to recruit workers from Dallas versus Lufkin is revealing about the elasticity of worker demand for remote work.

We assume that 1% of local customer service workers in each local labor market saw the retailer's advertisement. We then estimate a linear probability model of offer acceptance against the wage penalty or price, $P_{\text{Remote},\ell(i)}$, of taking a remote job at the retailer:[60]

$$\text{Accept Offer}_{i,\ell(i)} = \beta_0 + \beta_{\$} P_{\text{Remote},\ell(i)} + \epsilon_{i,\ell(i)}. \tag{14}$$

The coefficient of interest, $\beta_{\$}$, is unbiased if other determinants of recruitment are orthogonal to the relative wage of the retailer. Indeed, the recruiter advertised similarly across the country. Recruiters posted these jobs on national job boards. And they were indifferent about the home locations of new hires.[61] However, if remote work is more difficult in lower wage labor markets because internet access is less broad-based or workers are less tech-savvy, then this would tend to bias $\beta_{\$}$ downward. Similarly, measurement error in the price of working remotely remotely would attenuate $\beta_{\$}$ towards zero.

Figure 10 illustrates this empirical strategy. The x-axis plots the gap between the retailer's uniform \$14/hour wage and the entry-level local wage in customer service — or the implicit price the worker must pay to take this customer service job instead of a local alternative ($P_{\text{Remote},\ell(i)}$). The y-axis plots the percent of workers recruited by the retailer assuming 1% of local customer service

While these differences are small in absolute terms, they represent about a tenth of a standard deviation change, as the limited variation in customer satisfaction reviews makes it more challenging to draw firm conclusions about the effects of remote work on call quality.

[60]This is a linear approximation of the nonlinear relationship between price and quantity in our model. Future drafts will estimate our model, incorporating both the promotion and recruitment data.

[61]Our contact at the retailer described recruiters as "blind to the locations" when recruiting remote workers to the point that recruiters sometimes accidentally recruit workers from places where \$14/hr falls below the local minimum wage.

workers saw the advertisement.

To the left of Figure 10, the retailer's remote wage is competitive with local outside options. These places are like Lufkin, where the price that local workers must pay to take this remote job is low. As a result of this low price, the share of workers who take the job is relatively high. To the right of Figure 10, the retailer's remote wage is much lower than local outside options. These places are like Dallas, where the price that local workers must pay to take this remote job is high. As a result of this high price, the share of workers who take the job is relatively low. On average, in places where workers' local alternatives pay $1/hour more, the probability that the retailer recruits a given worker is 0.43pp lower. Relative to the average probability of 2.2% that a worker takes the retailer's job, this represents an 21.5% reduction in quantity for a $1/hour increase in price (95% CI = [-41.2%, -1.8%]).

Since the average worker must pay $1/hour to accept the retailer's remote job instead of an on-site alternative, this implies a 100% increase in the price translates into a 22% reduction in quantity demanded, yielding a semi-elasticity of demand of 0.22. This suggest workers have an inelastic demand for remote work.

It's useful to consider how this compares to other estimates in the literature. Mas and Pallais (2017) estimate workers' preferences for the *option* to work remotely through a choice experiment in a US call-center. Mas and Pallais (2017) find that workers at the 25th percentile of the preference distribution were essentially indifferent to the option to work remotely — they were willing to give up $0.20/hour (se = 0.50) for this option. By contrast, those at the 75th percentile were willing to give up $2.45/hour (se = 0.68) or 18% of their wage for this option. Taking a linear approximation of this demand curve would suggest a $2.25/hour increase in the price of remote work would lead to a 100% reduction in the quantity demanded. This yields a semi-elasticity of demand of 0.44.[62] Our estimate would then imply that tastes for full-time remote work are twice as variable as tastes for the *option* to work remotely.

It seems reasonable to suspect that full-time remote work would evoke stronger preferences than the option to work remotely. However, the truth may lie somewhere between these estimated semi-elasticities of 0.22 and 0.44. As noted above, if there are more challenges to working remotely in low wage labor markets, this would bias $\beta_{\$}$ down. So too would measurement error in workers' outside options. These biases would make demand appear less elastic than it truly was. A less elastic demand curve leads to smaller distortions in quantity for any given set of costs (as illustrated in panel (c) of Figure 4). Thus attenuation bias in $\beta_{\$}$ will lead to conservative estimates of the quantity distortion and DWL in the next section.

We can rearrange our estimated relationship from equation 14 to arrive at the implied demand

[62] Maestas et al. (2018) similarly finds heterogeneous preferences for remote work in the context of hypothetical choices, again suggesting an inelastic demand curve.

curve for remote work:

$$P_{\text{Remote}} = -\frac{\beta_0}{\beta_\$} + \frac{1}{N\,\Pr(\text{Seeing the Ad})}\frac{1}{\beta_\$}Q_{\text{Remote}} \tag{15}$$

where Q_{Remote} denotes the quantity of workers the retailer would recruit if it were to offer different remote wages. In addition to learning about worker preferences, this empirical strategy offers insight into how recruitment would differ in a counterfactual where the retailer had a higher remote wage. The retailer's greater ability to recruit in places like Lufkin, TX, where its remote wage is more competitive with local alternatives, offers a sense of the effect of increasing the remote wage would impact recruitment. This exercise consequently allows us to consider the counterfactual in which the retailer set the remote wage at the marginal rather than the average cost, as we investigate in the next section.

X WELFARE AND COVID-19 COUNTERFACTUALS

We can leverage our estimated demand curve to characterize worker surplus from remote work in the market equilibrium and a counterfactual in which the retailer set the remote wage at the marginal cost, as it would if remote work did not increase the informational frictions between the worker and the firm. Figure 11 takes our theoretical model (illustrated in Figure 3) to our data on call-center workers.

The green line reflects the efficient price of remote work. This is determined by the marginal cost or treatment effect of remote work in this setting. Our preferred estimate of the treatment effect is 7.5%, with estimates ranging from 6.9% to 9.8%. If the retailer could price remote work at the margin, it would pay a 7.5% wage premium for remote work. This would also be the premium it would set if its workforce were fixed: if the retailer's wages only impacted workers' actual jobs and not their offered jobs, then the retailer would set the price of remote work to reflect the positive treatment effect.

The orange line reflects the market price of remote work. This is determined by the average cost of remote work, which in turn reflects both the treatment and selection effect of remote work. Our preferred estimate of the average cost is 7.6%, with estimates ranging from 4.8% to 11.8%. This would imply a \$1.14/hour difference in wages, which aligns well with the \$1.10/hour difference in base wages at the retailer (row 11 of Table 1). Thus, despite the positive treatment effect of remote work in this context — on the order of 7.3% to 9.8% — remote workers were paid less than on-site workers.

The gap between the orange and green lines reflects adverse selection into remote work. Both within and across labor markets, we find that workers who were offered remote work were less productive than those who were offered on-site work. This is consistent with the prediction of our model that greater informational frictions in remote work will deter better workers who have more to lose from being overlooked for promotion.

Since the retailer's workforce is flexible rather than fixed, it must consider the different composition of remote and on-site workers when setting its wages. To offset the lower expected quality of remote workers, the retailer pays remote workers less than on-site workers. This gives all workers an incentive to go on-site to pool with more productive workers. These marginal workers — whose demand lies between the market price and the efficient price — only choose on-site work because of the externalities from others' choices. If the retailer could price remote work at the margin — which would reduce the price of remote work by $2.52/hour — newly remote workers would gain an average of $1.26/hour (or 8.4% of the average $15/hour wage). Using our estimated demand curve — which takes the number of people who see the job ad as given — the quantity of remote work at the retailer would rise from 19% of the retailer's hires to 29%. Together this suggests that surplus at the retailer would rise by $738,804 annually if the retailer could price remote work at the margin.

To put this number in perspective, it's useful to think about our back-of-the-envelope calculation of the surplus from remote work in the current equilibrium. Our estimates suggest that each remote worker gains $2.31/hour (or 15.4% of the average $15/hour wage) from working remotely.[63] This amounts to $2,573,499 annually in social surplus from remote work at the retailer even when the price of working remotely is set inefficiently high. Lowering this price to reflect the marginal cost of remote work would increase social surplus from remote work at the retailer by 29%.

We can consider a simple back-of-the-envelope calculation of how these estimates would extrapolate to the rest of the economy. In 2018, there were 3.5 million phone workers, using Mas and Pallais (2017)'s definition.[64] If these estimates applied equally well to all phone workers, then $917 million would be lost annually from this market failure.

The problem of adverse selection also depresses the wages of inframarginal remote workers. Pricing at the margin would transfer resources to these workers. This would be socially beneficial if those who choose remote work high high welfare weights and are difficult to transfer resources to through other means. Given our finding that remote work is a tag for lower productivity, standard models of optimal taxation would suggest that subsidizing remote work would be a more efficient means of transferring resources to remote workers than other perturbations of the tax system (Mirrlees, 1976; Atkinson and Stiglitz, 1976; Kaplow, 2006). Thus, there might be an additional public policy motive to subsidize remote work.[65]

As emphasized in the model, the selection effect of remote work need not be immutable. Instead, it is a function of tastes and technologies. Over the course of the pandemic, many workers may have updated their beliefs about their own tastes for remote work. Some may have learned just how much they had hated their long commuted while others realized they missed water cooler

[63]This is consistent with the significant surplus measured for other forms of workplace flexibility (Chen et al., 2019).
[64]They define phone workers as customer service workers, bill and account collectors, telemarketers, and interviewers except benefit eligibility and loan officers.
[65]These concerns might be especially acute given the prevalence of child-care responsibilities among remote workers, if policy makers worried that these losses would not be offset by other compensatory changes in the tax system.

chats more than they had anticipated. To assess this possibility, we are currently surveying the retailer's workers about their tastes and how their beliefs about their own tastes have changed because of the pandemic. A pilot of a hundred MTurk workers suggests that tastes have become 6.5% to 24.5% more variable, consistent with Bayesian updating.[66] This suggests that a lasting impact of the pandemic may have been to increase the variation in tastes, thereby weakening the link between ability and demand, closing the gap between the average and marginal costs of remote work, and permanently increasing the provision of remote work towards the socially optimal level as in panel (c) of Figure 4.

XI CONCLUSION

We consider why so few Americans were working remotely prior to COVID-19 even in seemingly remotable jobs. In the call-center context we study, the rarity of remote work seemed especially puzzling since (1) workers expressed a high willingness to pay for remote work (Mas and Pallais, 2017) and (2) working remotely seemed to make workers more productive, as evidenced inour natural experiments and Bloom et al. (2015)'s RCT.

We argue that the missing piece to this puzzle is selection into remote work. In our context, remote workers were 20pp less likely to be promoted in their first year than on-site workers. This complements Bloom et al. (2015)'s RCT evidence that remote work halved workers' promotion probabilities. A simple model of career concerns would suggest that better workers would be warier of remote work given its career consequences. Workers consequently sort into remote and on-site jobs on the basis of private information about their ability, leading to adverse selection in remote jobs.

This theoretical prediction is born out empirically. Those offered remote jobs were 18% less productive than those offered on-site jobs, holding the actual job constant. Thus, even though working remotely increased worker productivity on the margin, it was costly for the firm on average. Consistent with our model, the retailer paid remote workers less and, as a result, a small minority of the retailer's workforce worked remotely. Some workers opted out of remote work simply because they did not want to pool with less productive types. In the absence of asymmetric information, the quantity of remote work at the retailer would rise by 10pp and social surplus from the provision of remote work would rise by 29%.

As emphasized in our model, adverse selection into remote work is a function of technologies and tastes. When firms learn new techniques for identifying remote workers who merit promotion, this reduces the incentive for better workers to choose on-site jobs. Similarly, when tastes are more heterogeneous — either because tastes change or workers learn about their own tastes — the choice to be remote becomes a weaker signal of ability. Thus, the massive experiment in remote

[66] Specifically, the survey asks workers to rate their happiness working on-site, their initial happiness working remotely, and their current happiness working remotely on a scale from 0 to 100. It then asks them whether they have liked remote work more or less than expected. Finally, it asks them to rate their expected happiness working remotely on a scale from 0 to 100.

work caused by the pandemic could reduce adverse selection into remote work through either of these mechanisms. In future iterations of the paper, we will play these scenarios out and bring more evidence to bear on these ramifications of COVID-19.

While we believe our paper makes a useful contribution to the nascent literature on remote work, it leaves a few important questions unanswered.

Our study speaks to a couple of prominent remote work policies: namely, hiring workers directly into remote jobs and allowing on-site workers to transition to fully remote jobs. Our setting does not allow us to speak directly to the remaining possibility of allowing workers to work remotely for part of each week. While we think that the same selection effects are likely to arise, it is possible that this strategy might mute advertise selection by giving workers at least some face-time with their bosses. As firms ask how many days in the office are necessary after the pandemic, the ramifications for worker selection will be important to quantify.

Our setting featured workers doing independent tasks. Thus, we cannot speak to the important question of how remote work affects collaboration in teams.[67] Other work suggests that coordination may be more difficult from afar: for example, Battiston et al. (2017) finds that emergency phone operators communicated more efficiently with one another when they were physically in the same space.[68] Throughout the workforce, managers express more concern about remote work when workers are on teams (Beham et al., 2015; Juhász et al., 2020). However, even for collaborative jobs, there are independent tasks. Thus, learning about workers' at-home versus on-site productivity will be informative for these jobs as well, especially as firms increasingly ask how many days in the office are necessary.

We have estimated the sufficient statistics of our model in the specific context of call-center work but have proposed a framework that could be applied in other settings. Painting a similar picture in other sectors of the economy would be a fruitful direction for future work.

[67]The independent nature of the jobs also means much learning comes from learning-by-doing. Thus, we cannot speak to on-the-job learning in settings where peer mentorship is more pivotal.

[68]Beyond productivity effects, one worker's decision to go remote may have negative externalities on others' on-the-job experiences. In the US Patent Office, Linos (2018) finds that when one worker goes remote their coworkers have greater absenteeism and attrition.

References

Adams-Prassl, Abi, "The Gender Wage Gap on an Online Labour Market: The Cost of Interruptions," *CEPR Discussion Paper DP14294*, 2020.

Akerlof, George A, "The market for "lemons": Quality uncertainty and the market mechanism," in "Uncertainty in economics," Elsevier, 1978, pp. 235–251.

Atkinson, Anthony Barnes and Joseph E Stiglitz, "The design of tax structure: direct versus indirect taxation," *Journal of public Economics*, 1976, *6* (1-2), 55–75.

Bailyn, Lotte, *Breaking the mold: Women, men, and time in the new corporate world.*, Free Press, 1993.

Baker, George, "Distortion and risk in optimal incentive contracts," *Journal of human resources*, 2002, pp. 728–751.

Barbanchon, Thomas Le, Roland Rathelot, and Alexandra Roulet, "Gender differences in job search: Trading off commute against wage," *The Quarterly Journal of Economics*, 2021, *136* (1), 381–426.

Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis, "Why Working From Home Will Stick," *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2020, (2020-174).

Battiston, Diego, Jordi Blanes i Vidal, and Tom Kirchmaier, "Is distance dead? Face-to-face communication and productivity in teams," 2017.

Beham, Barbara, Andreas Baierl, and Steven Poelmans, "Managerial telework allowance decisions–a vignette study among German managers," *The International Journal of Human Resource Management*, 2015, *26* (11), 1385–1406.

Benabou, Roland and Jean Tirole, "Intrinsic and extrinsic motivation," *The review of economic studies*, 2003, *70* (3), 489–520.

Bénabou, Roland and Jean Tirole, "Incentives and prosocial behavior," *American economic review*, 2006, *96* (5), 1652–1678.

Bender, Stefan, Nicholas Bloom, David Card, John Van Reenen, and Stefanie Wolter, "Management practices, workforce selection, and productivity," *Journal of Labor Economics*, 2018, *36* (S1), S371–S409.

Bloom, Nicholas, James Liang, John Roberts, and Zhichun Jenny Ying, "Does working from home work? Evidence from a Chinese experiment," *The Quarterly Journal of Economics*, 2015, *130* (1), 165–218.

Brown, Christina and Tahir Andrabi, "Inducing Positive Sorting through Performance Pay: Experimental Evidence from Pakistani Schools," 2020.

Brynjolfsson, Erik, John J Horton, Adam Ozimek, Daniel Rock, Garima Sharma, and Hong-Yi TuYe, "COVID-19 and remote work: an early look at US data," Technical Report, National Bureau of Economic Research 2020.

Cadsby, C Bram, Fei Song, Francis Tapon et al., "Are you paying your employees to cheat? An experimental investigation," *The BE Journal of Economic Analysis & Policy*, 2010, *10* (1), 1–32.

Caldwell, Sydnee and Oren Danieli, "Outside options in the labor market," *Unpublished manuscript*, 2018.

Chen, M Keith, Peter E Rossi, Judith A Chevalier, and Emily Oehlsen, "The value of flexible work: Evidence from uber drivers," *Journal of Political Economy*, 2019, *127* (6), 2735–2794.

Choudhury, Prithwiraj, Cirrus Foroughi, and Barbara Zepp Larson, "Work-from-anywhere: The productivity effects of geographic flexibility," in "Academy of Management Proceedings" number 1. In '2020.' Academy of Management Briarcliff Manor, NY 10510 2020, p. 21199.

DeFilippis, Evan, Stephen Michael Impink, Madison Singell, Jeffrey T Polzer, and Raffaella Sadun, "Collaborating during coronavirus: The impact of COVID-19 on the nature of work," *NBER Working Paper*, 2020, (w27612).

Dingel, Jonathan I and Brent Neiman, "How many jobs can be done at home?," *Journal of Public Economics*, 2020, *189*, 104235.

Einav, Liran, Amy Finkelstein, and Mark R Cullen, "Estimating welfare in insurance markets using variation in prices," *The quarterly journal of economics*, 2010, *125* (3), 877–921.

Elsbach, Kimberly D, Dan M Cable, and Jeffrey W Sherman, "How passive 'face time' affects perceptions of employees: Evidence of spontaneous trait inference," *human relations*, 2010, *63* (6), 735–760.

Ford, Denae, Margaret-Anne Storey, Thomas Zimmermann, Christian Bird, Sonia Jaffe, Chandra Maddila, Jenna L Butler, Brian Houck, and Nachiappan Nagappan, "A tale of two cities: Software developers working from home during the covid-19 pandemic," *arXiv preprint arXiv:2008.11147*, 2020.

Goldin, Claudia, Sari Pekkala Kerr, and Claudia Olivetti, "Why Firms Offer Paid Parental Leave: An Exploratory Study," Technical Report, National Bureau of Economic Research 2020.

Grip, Andries De and Jan Sauermann, "The effects of training on own and co-worker productivity: Evidence from a field experiment," *The Economic Journal*, 2012, *122* (560), 376–399.

Gruber, Jonathan, "The incidence of mandated maternity benefits," *The American economic review*, 1994, pp. 622–641.

_ **and Alan B Krueger**, "The incidence of mandated employer-provided insurance: Lessons from workers' compensation insurance," *Tax policy and the economy*, 1991, *5*, 111–143.

Holmstrom, Bengt and Paul Milgrom, "Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design," *JL Econ. & Org.*, 1991, *7*, 24.

Juhász, Réka, Mara Squicciarini, and Nico Voigtländer, "Away from Home and Back: Coordinating (Remote) Workers in 1800 and 2020," *NBER Working Paper*, 2020, (w28251).

Kaplow, Louis, "On the undesirability of commodity taxation even when income taxation is not optimal," *Journal of Public Economics*, 2006, *90* (6-7), 1235–1250.

Karlan, Dean and Jonathan Zinman, "Observing unobservables: Identifying information asymmetries with a consumer credit field experiment," *Econometrica*, 2009, *77* (6), 1993–2008.

Lazear, Edward P, "Output-based pay: incentives, retention or sorting," *Research in Labor Economics*, 2004, *23* (4), 1–25.

Linos, Elizabeth, "Does teleworking work for organizations?," Technical Report, Harvard University Working Paper 2018.

Maestas, Nicole, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger, "The value of working conditions in the united states and implications for the structure of wages," Technical Report, National Bureau of Economic Research 2018.

Manchester, Colleen Flaherty, "General human capital and employee mobility: How tuition reimbursement increases retention through sorting and participation," *ILR Review*, 2012, *65* (4), 951–974.

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

_ **and** _ , "Alternative work arrangements," Technical Report Forthcoming.

Mirrlees, James A, "Optimal tax theory: A synthesis," 1976.

Miyazaki, Hajime, "The rat race and internal labor markets," *The Bell Journal of Economics*, 1977, pp. 394–418.

Morning Consult, "The Future of Work: How the Pandemic Has Altered Expectations of Remote Work," Technical Report, Morning Consult 2020.

Oettinger, Gerald S, "The incidence and wage consequences of home-based work in the United States, 1980–2000," *Journal of Human Resources*, 2011, *46* (2), 237–260.

Ovide, Shira, "What Job Sites Reveal About the Economy," *The New York Times*, Mar 2021.

Oyer, Paul and Scott Schaefer, "Why do some firms give stock options to all employees?: An empirical examination of alternative theories," *Journal of financial Economics*, 2005, *76* (1), 99–133.

PwC, "When everyone can work from home, what's the office for? PwC's US Remote Work Survey," Technical Report, PwC USA 2020.

Ruhm, Christopher J, "The economic consequences of parental leave mandates: Lessons from Europe," *The quarterly journal of economics*, 1998, *113* (1), 285–317.

Salop, Joanne and Steven Salop, "Self-selection and turnover in the labor market," *The Quarterly Journal of Economics*, 1976, pp. 619–627.

Schulte, Brigid, "What Moms Always Knew ABout Working From Home," *New York Times*, 2020, pp. Op–Ed.

Summers, Lawrence H, "Some simple economics of mandated benefits," *The American Economic Review*, 1989, *79* (2), 177–183.

Tô, Linh T, "The Signaling Role of Parental Leave," *Mimeo*, 2018.

Weiss, Andrew, "Human capital vs. signalling explanations of wages," *Journal of Economic perspectives*, 1995, *9* (4), 133–154.

Yang, Longqi, Sonia Jaffe, David Holtz, Siddharth Suri, Shilpi Sinha, Jeffrey Weston, Connor Joyce, Neha Shah, Kevin Sherman, CJ Lee et al., "How Work From Home Affects Collaboration: A Large-Scale Study of Information Workers in a Natural Experiment During COVID-19," *arXiv preprint arXiv:2007.15584*, 2020.

Figure 1: Trajectory of Calls per Hour for Remote and On-site Workers

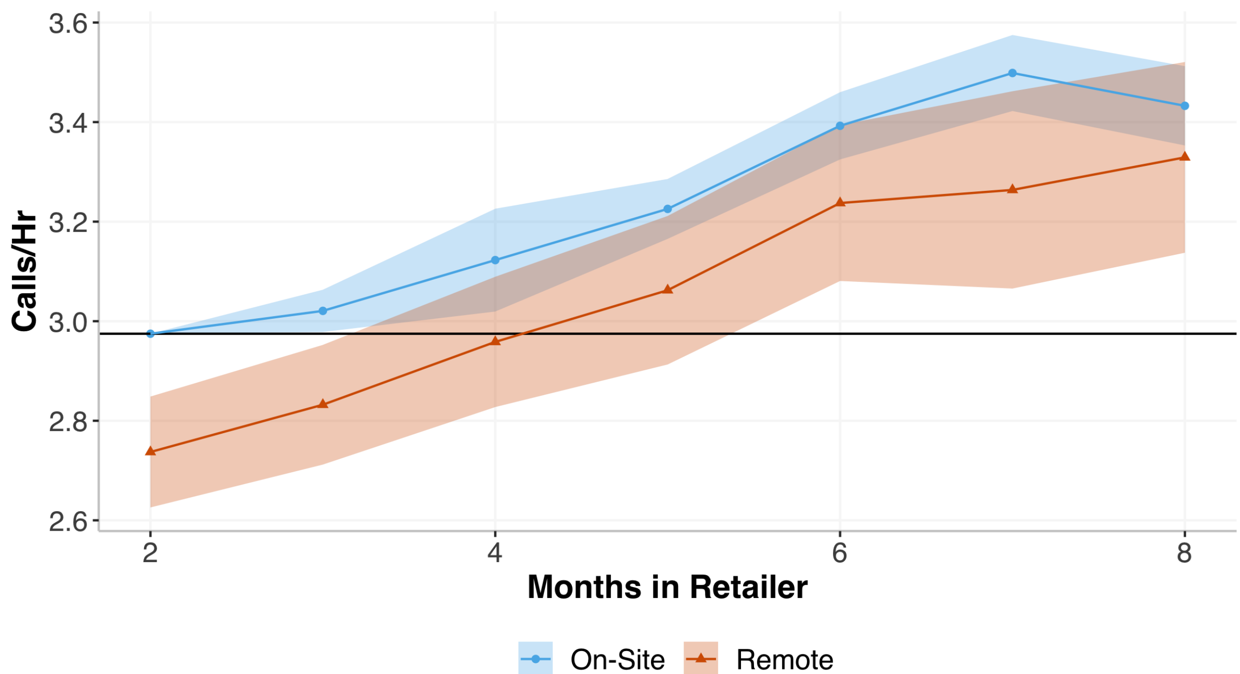

Notes: This figure plots the hourly call volumes of remote and on-site workers as a function of their time at the retailer. The x-axis represents the months in the retailer. The y-axis represents the hourly calls of on-site workers in blue circles and of remote workers in orange triangles. This analysis omits the first month of training when workers do not handle calls independently. These estimates are adjusted for date by time-zone fixed effects in order to absorb any shocks to consumer demand. The error ribbons reflect the 95% confidence interval with standard errors clustered at the worker level. The sample limits to workers hired between July 2018 — when the retailer began to hire remote workers — and April 2020 when the retailer closed its offices due to COVID-19.

Figure 2: Percent Promoted By Tenure for Remote and On-Site Workers

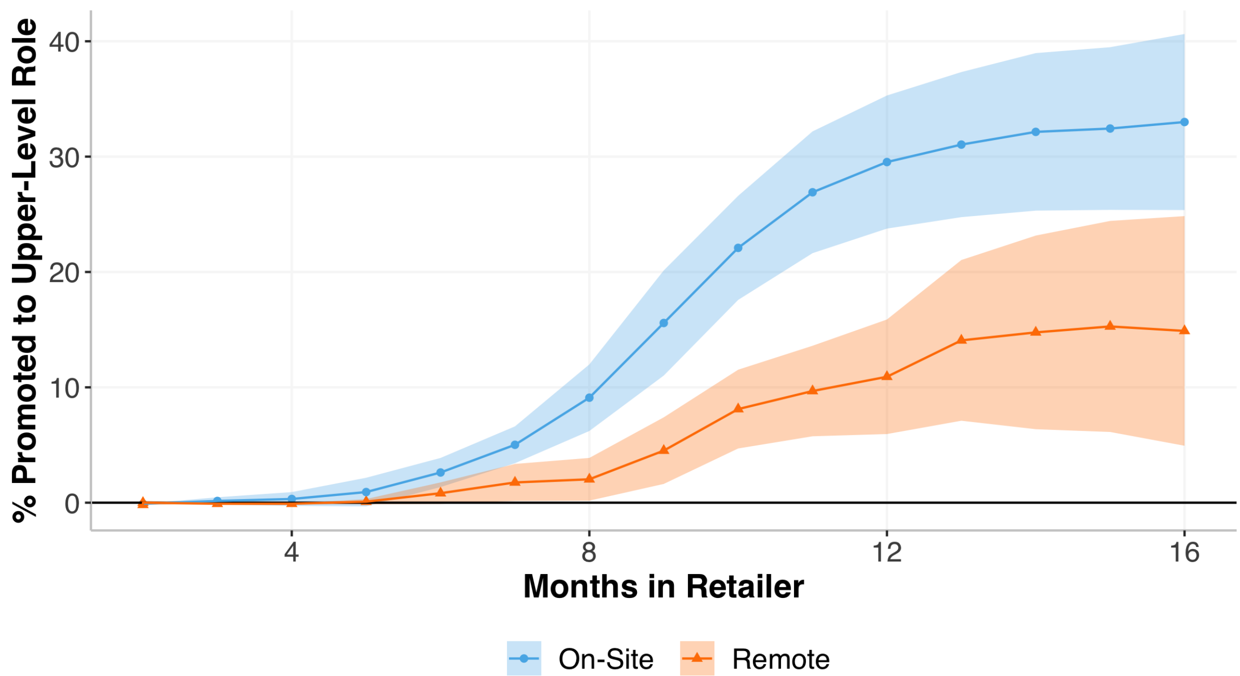

Notes: This figure plots the share of workers who have been promoted on the y-axis as a function of their tenure at the retailer on the x-axis. On-site workers are pictured in blue circles and remote workers in orange triangles. The error ribbons reflects 95% confidence intervals with standard errors clustered at the worker level.

Figure 3: Selection Market for Remote Work

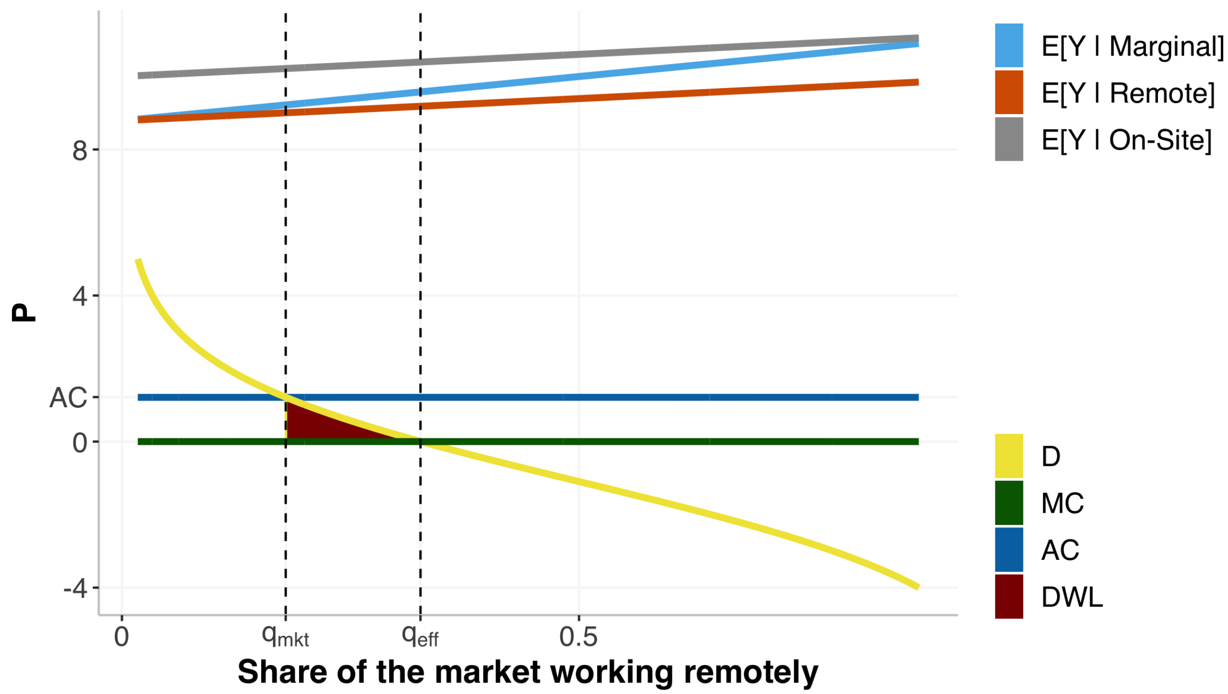

Notes: This figure plots the market for remote work under selection into on-site and remote jobs assuming there is no treatment effect of remote work on productivity. The x-axis represents the share of the market working remotely. The y-axis represents the price or wage penalty of remote work. The yellow curve plots the demand curve for remote work or the share of the market that would work remotely at any given price. Since the expected ability of workers on the margin of remote work, $\mathbb{E}[Y \mid \text{Marginal}]$, rises with the share of the market working remotely, the marginal product in remote work, drawn in light blue, is increasing. The average product in the remote job, $\mathbb{E}[Y \mid \text{Remote}]$, drawn in orange, integrates the light blue line from left to right to average over the output of marginal and inframarginal remote workers. The average product in the on-site job, $\mathbb{E}[Y \mid \text{On-Site}]$, drawn in grey, integrates the light blue line from right to left to average over marginal and inframarginal on-site workers. The differences in average product between the on-site workers (in grey) and the remote workers (in orange) produces the average cost, AC, of remote work to the firm in navy blue. This will be the equilibrium price of remote work in the market. The intersection with the demand curve in yellow will determine the equilibrium share of the market working remotely. By contrast, the efficient price of remote work would be zero, which would induce a higher share of the market to work remotely.

Figure 4: Comparative Statics in the Market for Remote Work

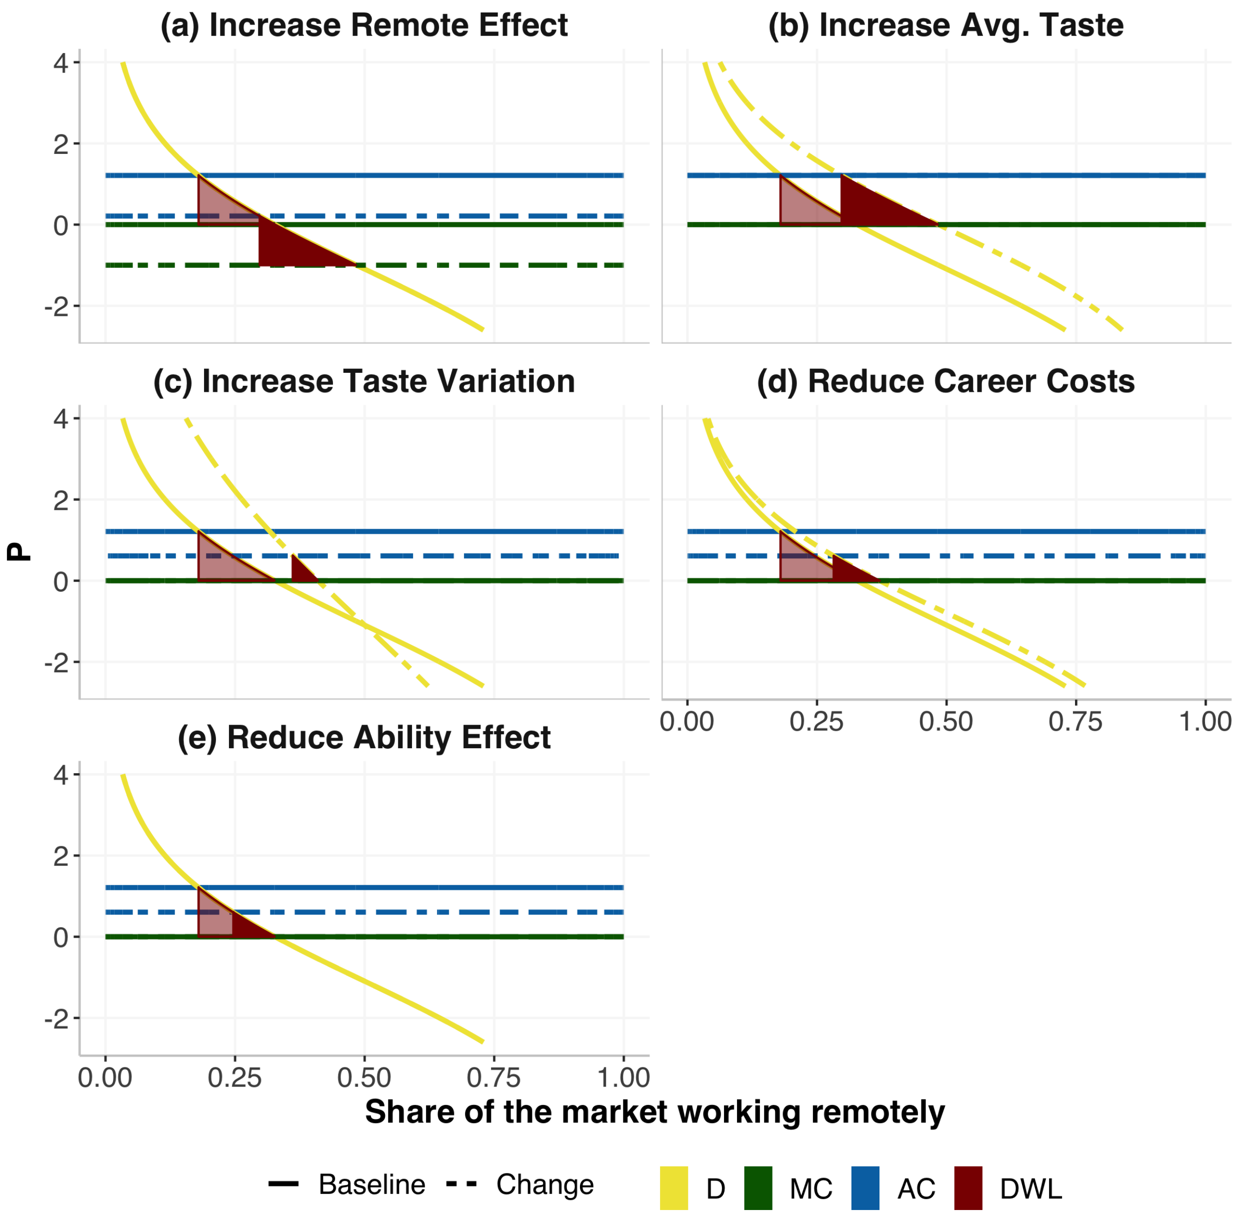

Notes: This figure considers the effects of perturbing the five forces that drive the market for remote work. In each panel, the x-axis plots the share of the market working remotely and the y-axis plots the price of remote work, or wage penalty. The yellow curves represent demand; the green lines represent the marginal costs of remote work; the navy lines represent the average costs. The solid lines plot the original demand and average cost curves. The dashed lines plot the demand and cost curves after each perturbation. The intersection of demand and marginal cost determines the efficient allocation while the intersection with average cost determines the market quantity. The red triangles highlight the deadweight losses from adverse selection. The five panels consider the five forces in this market: (a) the treatment effect of remote work on productivity (τ), (b) the average taste for remote work (μ), (c) the variation in these tastes (s), (d) the potential career costs of remote work ($\kappa = \frac{\delta}{1+\delta}(p_o - p_r)A$), and (e) the returns to ability in the low-skill task (a).

Figure 5: Intuition for Identification Strategy

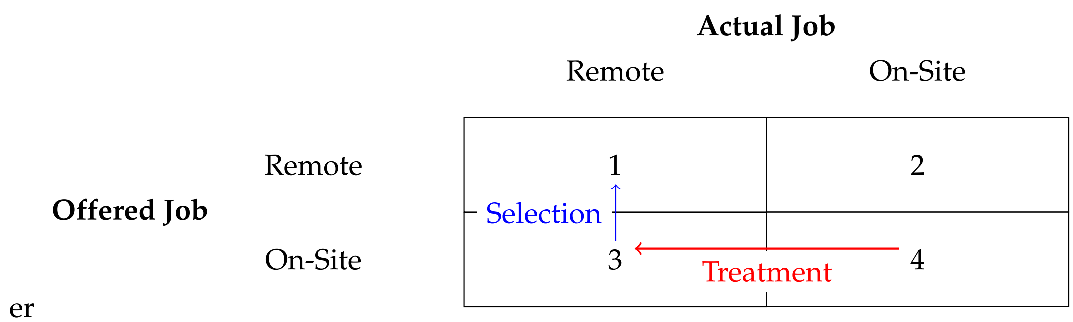

Notes: This figure illustrates the ideal thought experiment. The rows reflect the type of job that the worker was randomly offered at recruitment. The columns reflect the type of job that the worker was actually assigned. Under random assignment to these cells, the selection effect of remote work could be identified by comparing remote workers who were offered remote jobs (cell 1) to those who were initially offered on-site ones (cell 3). The treatment effect could be identified by comparing workers who were recruited into on-site jobs but then worked remotely (cell 3) to those who continued on-site (cell 4).

Figure 6: The Offer to go Remote and the Productivity of New Hires

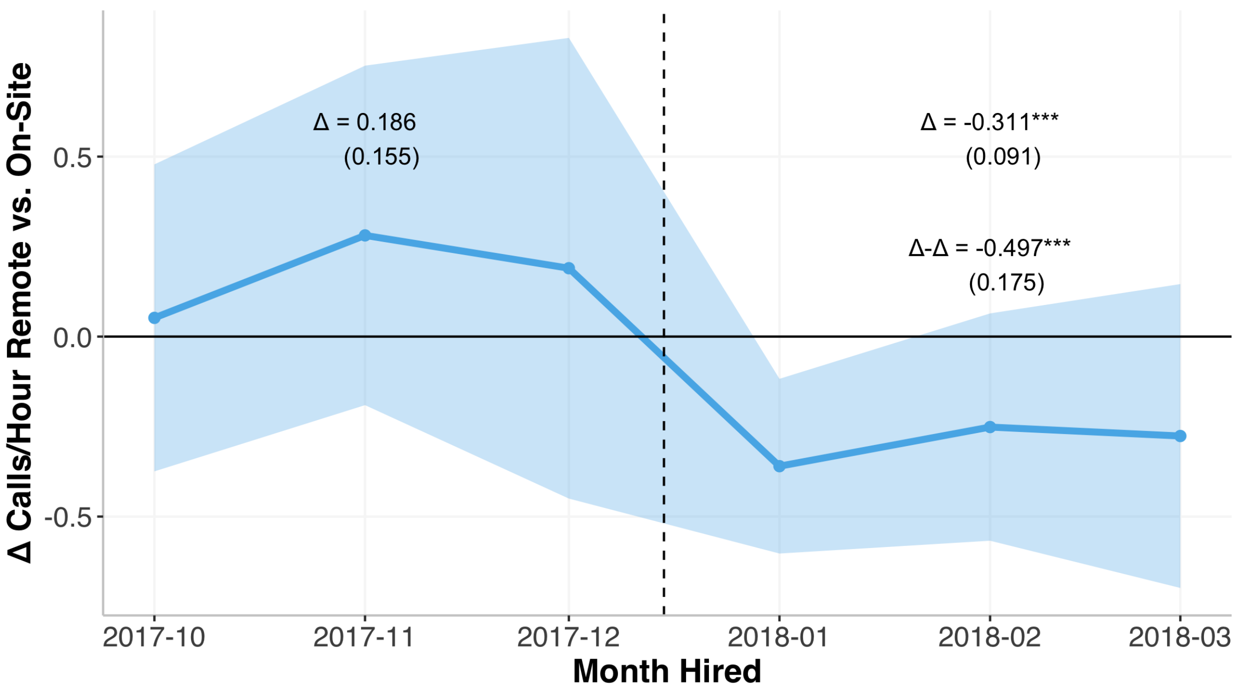

Notes: This figure analyzes the introduction of opportunities to go remote into the choice sets of new hires. The x-axis depicts the hiring month. The vertical dashed line indicates January 1, 2018 when the retailer introduced opportunities for on-site workers to go remote. The y-axis depicts the difference in calls per hour of remote and on-site workers from each cohort. The point estimates come from specification 7, estimated with date by hiring month by time-zone fixed effects. The ribbons reflect 95% confidence intervals with standard errors two-way clustered by date and worker. The annotations reflect estimates from the pooled specification 6, which are also reported in Table 3. The sample limits to workers who were hired into entry-level roles and were in their first 6 months at the retailer to exclude workers who have advanced to specialized roles with non-random routing of calls. The sample further excludes all workers' first three months because productivity data were not available in our data for the first three months of the cohort of workers hired in October 2017.

Figure 7: Event Study Around Individual Switches to Remote Work

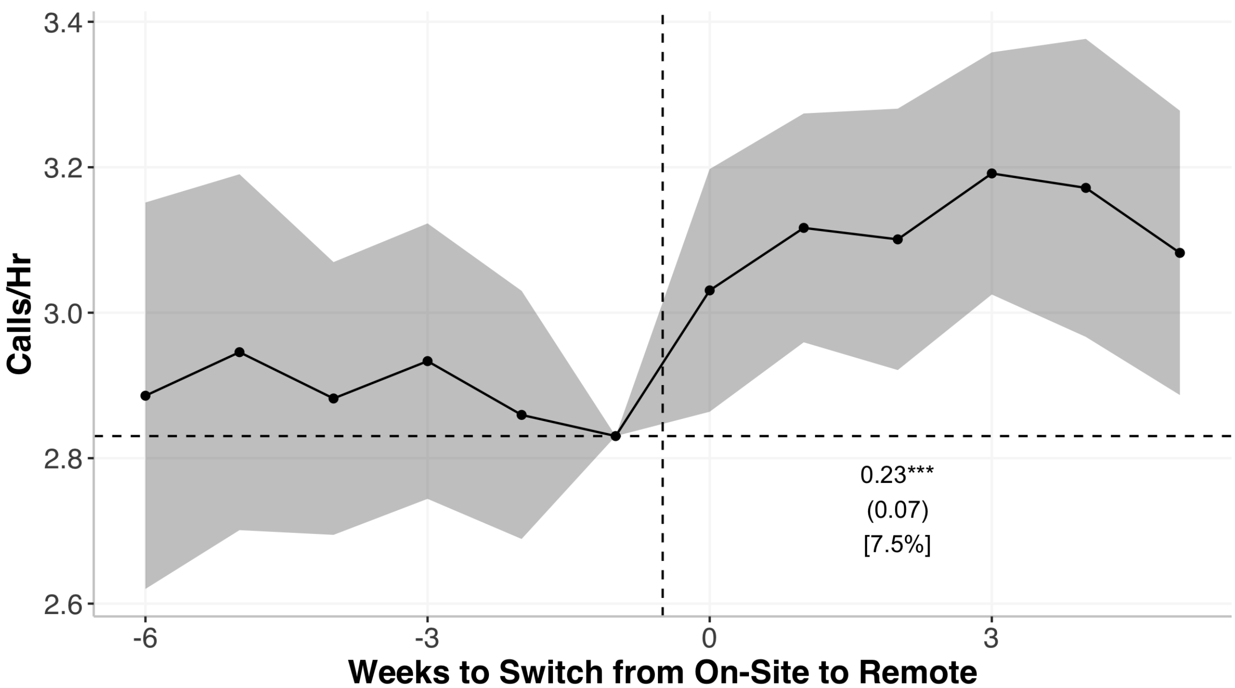

Note: This figure depicts the change in calls per hour of workers who switched from on-site to remote work when the retailer posted remote jobs that were open to on-site workers in 2018. The x-axis represents event time in weeks from the switch to remote work. The y-axis represents the worker's calls taken per hour. The dashed horizontal line reflects workers' average productivity in the index week before they transitioned to remote work. The error ribbon reflects a 95% confidence interval, which compares productivity in each week to that in the index week according to equation 8. Standard errors are clustered at the individual level. The annotated coefficient compares the full six weeks before and after individuals switch to remote work as in equation 9. The bracketed number represents the percent change in productivity. This analysis limits to a balanced panel of 120 entry-level workers who made a switch from on-site to remote work that was not proximate to a promotion or departure from the retailer.

Figure 8: The Distributions of Calls/Hour in Remote Work for Workers Initially Offered Remote and On-Site Jobs

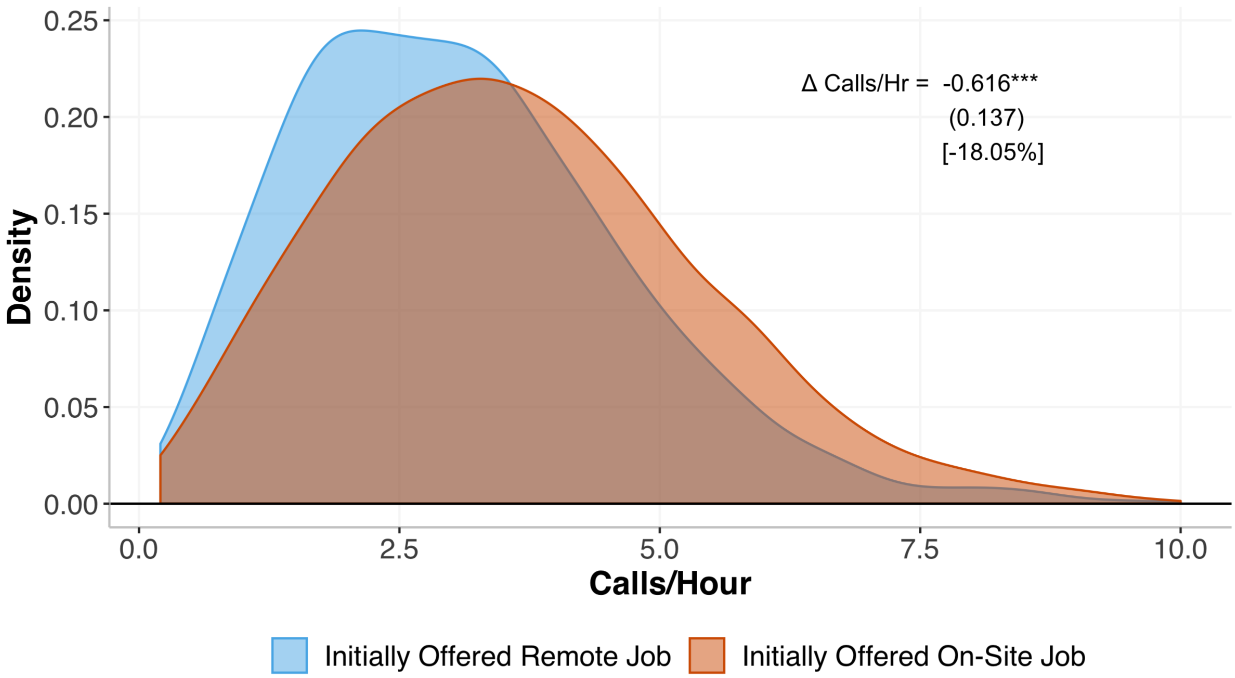

Notes: This figure illustrates the distributions of calls handled per hour for workers who were initially offered remote jobs (in blue) and workers who were initially offered on-site jobs (in orange) when all workers were working remotely due to COVID-19's lockdown. The x-axis represents the average calls taken per hour in a worker's day. The y-axis represents the density of each group at the call volume. The observations are at the day by worker level. The sample limits to remote workers and on-site workers at the retailer's $14/hour locations, who began in entry-level roles and were in their first six months at the retailer. The data ranges from April 6, 2020 when the retailer closed its on-site call-centers to August 6, 2020. The annotated coefficient estimates equation 11, with standard errors two-way clustered by worker and date.

Figure 9: Difference-in-Difference in Calls/Hour Around COVID-19 Office Closures

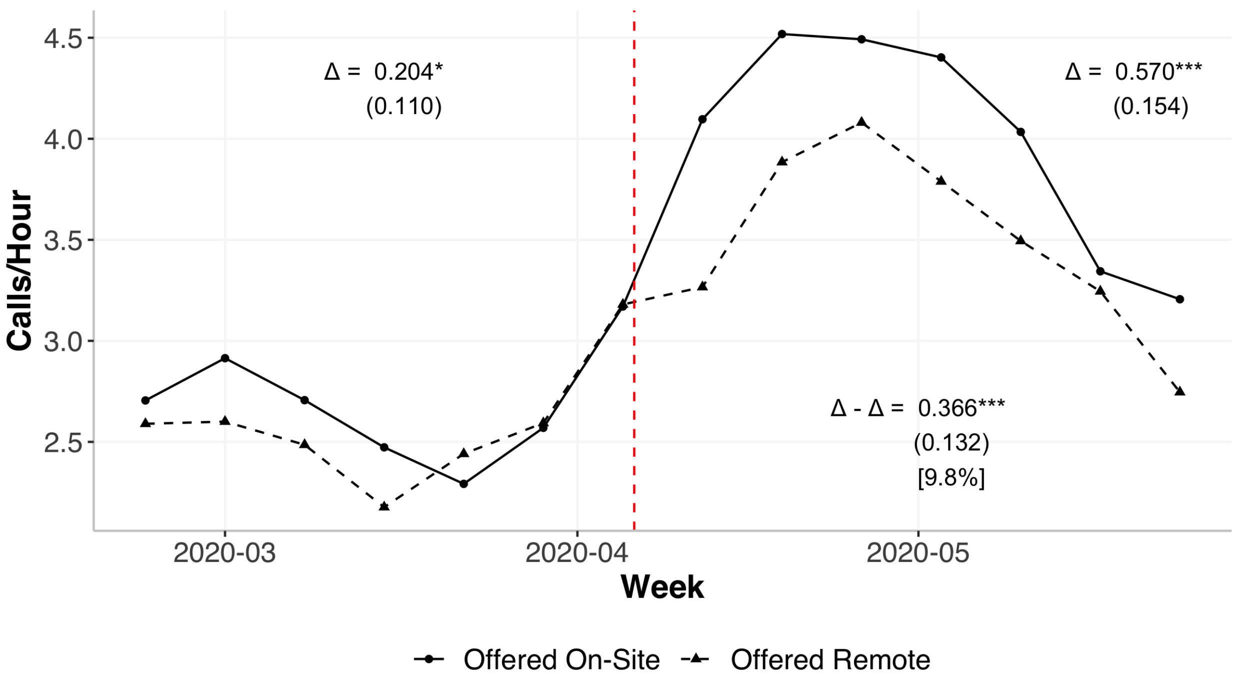

Notes: This figure illustrates the difference-in-difference in productivity between on-site workers who were transitioning into remote work (in the solid black line) and remote workers who were already working remotely (in the dashed line). The x-axis indicates the week in which calls were taken. The y-axis represents calls per hour. The vertical red dashed line indicates April 6, 2020, the date on which the retailer closed its on-site call-centers due to COVID-19. Each point is the average hourly calls of either formerly on-site workers (in circles and solid lines) or already remote workers (in triangles and dashed lines). The sample includes workers who were hired before the last week of February 2020 and had entry-level wages of $14/hour. The sample limits to workers' first six months at the retailer. During this time, only one of the workers in the sample left the retailer, mitigating concerns about selective attrition. The annotated coefficient on the left of the red line indicates the productivity difference between workers hired into on-site jobs and workers hired into remote jobs in the seven weeks prior to the office closures according to specification 10 with time-zone by date by hire month fixed effects. The annotated coefficient to the top right presents the same difference in the seven weeks after the office closures. The coefficient in the bottom right indicates the difference in these differences in levels and relative to call volumes after the COVID-19 closures. All standard errors are clustered by worker.

Figure 10: Recruitment as a Function of the Implicit Price of Remote Work

Notes: This figure analyzes the recruitment at the retailer as a function of the implicit price of accepting this remote job. The x-axis plots the price of remote work — the gap between the retailer's uniform $14/hour remote wage and the typical entry wage in local customer service jobs in 2018 and 2019, approximated by the average of the first two quartiles of the wage distribution. The y-axis plots the percent of local customer service workers who accepted the retailer's offer assuming 1% saw the advertisement. Each point reflects a decile of the price of remote work given the comparison between the retailer's uniform $14/hour wage and local outside options throughout the country. The sample excludes locations where the retailer has a physical call-center and, consequently, only hires on-site workers. The fit line estimates the linear regression in equation 14. The annotated coefficient reflects the slope of this line with the standard error clustered at the MSA level. The error ribbon reflects a 95% confidence interval around the predicted values. Information on customer service (SOC 5240) in local labor markets comes from Emsi, the statistical modeling company used by the retailer.

Figure 11: Status Quo and Efficient Welfare from Remote Work

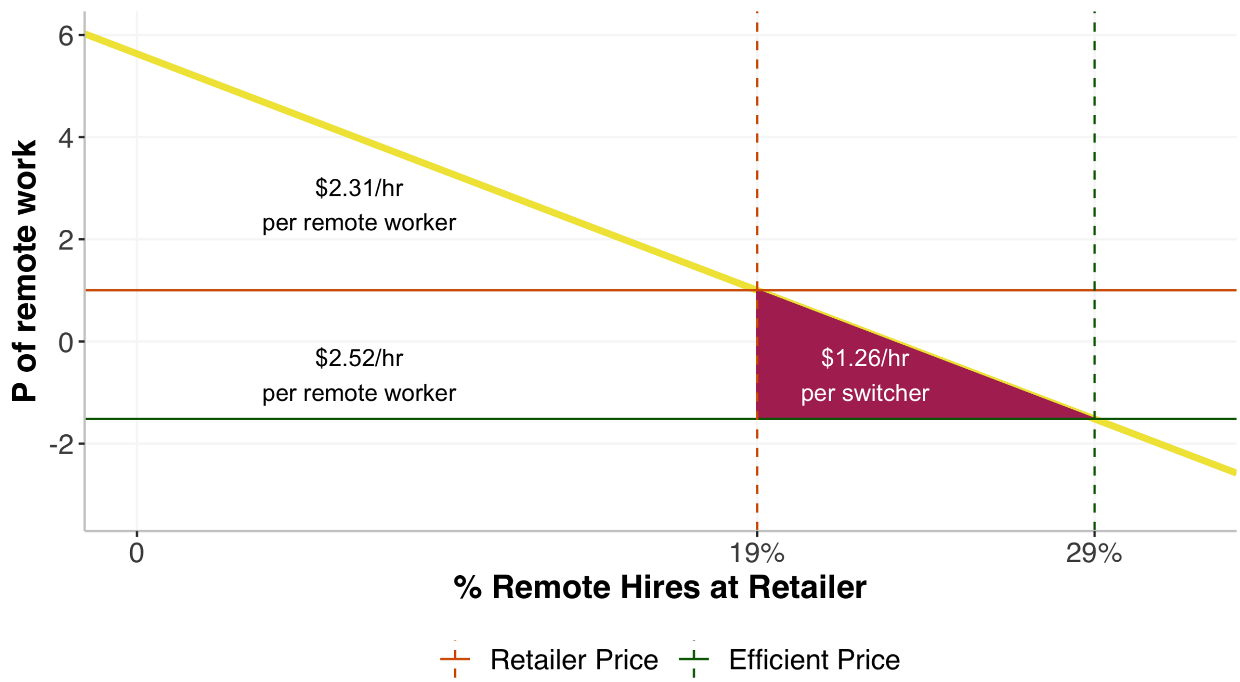

Notes: This figure interprets the model in Figure 4 in light of our empirical estimates for the retailer's call-center workers. The x-axis plots the percent of the retailer's post-2018 hires who were remote. The y-axis plots the implicit price of remote work, determined by the wage that workers must sacrifice to take this remote job instead of a local on-site option. The horizontal orange line reflects the current price workers must pay to take this remote job, set at the average cost of remote work. The horizontal green line reflects the price the retailer would charge for remote work if it could price at the marginal cost or treatment effect of remote work. The yellow line reflects the demand curve, found by pivoting the mapping between the implicit price of remote work and the number of hires at the retailer pictured in Figure 10.

Table 1: Summary Statistics: On-Site and Remote Workers

		All Workers	Hired On-Site	Hired Remote	Δ
1.	% Female	71.7	88.3	68.6	19.7***
2.	Age	32.3	35.6	31.6	3.9***
3.	% Age $\geq$40	22.1	29.5	20.7	8.8***
	Caregiving				
4.	% Any	61.4	73.8	58.7	15.1***
5.	% Child	46.1	57.7	43.7	14.0***
6.	% Disability	17.7	25.3	16.1	9.2**
7.	% Elder	13.5	15.6	13.0	2.5
8.	% Promoted	15.1	4.8	17.1	-12.3***
9.	Annual Turnover %	386.3	489.8	366.2	123.6
10.	% Terminations due to Fires	12.4	11.8	12.5	-0.7
11.	Wage	14.9	14.0	15.1	-1.1***
12.	MSA CSR wage	15.4	16.0	15.3	0.8***
13.	% Part-time	5.4	6.1	5.3	0.8
	# Workers	3440	558	2882	
	# Caregiving Respondents	854	149	705	
	# MSA	99	95	9	
	# Time-zone	4	4	4	

Notes: This table characterizes on-site and remote workers at the online retailer. The sample limits to those who were hired into entry level roles between 2018 and the beginning of the pandemic, taken to be March 15, 2020. The first column considers all workers; the second column, workers hired into remote jobs; the third column, workers hired into on-site jobs. The fourth column presents the difference between them with standard errors clustered at the employee level. Care-giving responsibilities were reported in a retailer-wide survey conducted in June 2020. Data on the median wage in customer-service (CSR) in the worker's metropolitan statistical area (MSA) comes from Emsi, an economic modeling company that the retailer uses to understand local labor markets. This company aggregates data from publicly available sources, particularly the Occupational Employment Statistics from the Bureau of Labor Statistics. Table A.1 reports the same comparisons in the American Community Survey.

Table 2: Predictive Power of Manager Evaluations Remote and On-Site

	Daily Customer Satisfaction in 2019 (in SD)			
	(1)	(2)	(3)	(4)
	All Workers	On-Site	Remote	All Workers
Manager Evaluation in Dec 2018	0.053**	0.069***	−0.026	0.066***
	(0.024)	(0.026)	(0.038)	(0.025)
Avg. Cust. Sat. July-Dec 2018	0.275***	0.253***	0.385***	0.282***
	(0.039)	(0.042)	(0.059)	(0.036)
Remote x Manager Eval				−0.080*
				(0.043)
Remote				−0.176
				(0.133)
Date x Time-Zone x Job Title FE	✓	✓	✓	✓
Mean Review	2.81	2.80	2.82	2.81
Std. Dev. Review	0.93	0.95	0.90	0.93
# Workers	899	646	237	899
# Days	108100	77742	28535	108100

Notes: This table analyzes whether managers have additional information about workers' skills beyond the recorded metrics and whether this additional information differs between remote and on-site workers. The first column estimates equation 1 for the full sample of workers who received performance reviews in December 2018. The second column limits this analysis to workers who were working on-site during the evaluation period from July to December 2018. The third column limits instead to workers who were working remotely during this time. The fourth column evaluates the interaction between a manager's evaluation and whether the worker was working remotely during the evaluation period. Standard errors are clustered at the worker level.

Table 3: The Offer to go Remote and the Productivity of New Hires

	Calls/Hour			Customer Satisfaction		
	(1)	(2)	(3)	(4)	(5)	(6)
Remote x Hired After Intro	-0.497^{***}	-0.518^{***}	-0.483^{**}	0.025	0.036	0.040
	(0.175)	(0.181)	(0.189)	(0.028)	(0.029)	(0.027)
Remote	0.186	0.133	-0.004	-0.023	-0.030	-0.044
	(0.150)	(0.150)	(0.306)	(0.025)	(0.025)	(0.033)
Months On-Site			0.037			0.004
			(0.069)			(0.008)
% Effect	-18.83%	-19.50%	-18.18%	0.50%	0.74%	0.81%
	(6.65)	(6.82)	(7.14)	(0.57)	(0.59)	(0.56)
t x Hire Month x Time-zone	✓			✓		
t x Hire Month x Call-center		✓	✓		✓	✓
Dependent Mean	2.64	2.65	2.65	4.89	4.89	4.89
Dependent Std. Dev.	1.43	1.42	1.42	0.41	0.41	0.41
# Workers	309	277	277	307	275	275
# Remote Workers	101	101	101	100	100	100
# On-site Workers	208	176	176	207	175	175
# Days	12940	11489	11489	9842	8717	8717

Notes: This table considers the productivity of remote workers who had different initial offers. Workers hired after the remote work program began in January 2018 knew that they could apply to go remote when they accepted the job offer. For earlier hires, these opportunities were unknown at the time they accepted the offer. Each column compares the productivity differences across these remote workers to that of on-site workers from the same cohort, according to specification 6. Columns one and four limit these comparisons to the same time-zone. The other columns further limits to those hired into the same call-center. Columns three and six include a control for months spent on-site to allow for different returns to experience on-site. Each observation is at the worker by day level. Columns one through three consider the average calls handled per working hour; the remaining columns consider the average customer satisfaction, which is only available on days where at least one customer left a review. Each analysis limits to workers hired into entry level roles between October 1, 2017 and April 1, 2018 who were in their first six months in the retailer, so were fielding calls randomly routed form the same pool data were available for all cohorts. The sample further excludes all workers' first three months because productivity data were not available in our data for the first three months of the cohort of workers hired in October 2017. Standard errors are two-way clustered at the day and worker level.

Table 4: Event Study Around Individual Switches to Remote Work

	Calls/Hr		Satisfaction Reviews		Unexcused Absent Min.	
	1 week	6 weeks	1 week	6 weeks	1 week	6 weeks
Post	0.202**	0.230***	0.003	0.011	−4.312	−1.897
	(0.084)	(0.068)	(0.027)	(0.011)	(2.638)	(1.493)
% Effect	6.92	7.55	0.06	0.22	-41.93	-18.95
	(2.87)	(2.23)	(0.55)	(0.22)	(25.66)	(14.92)
Dependent Mean	2.93	3.00	4.90	4.89	10.28	10.01
# Workers	120	120	120	120	120	120
# Days	988	5813	765	4562	988	5813

Note: This table reports the change in productivity around workers' switches from on-site to remote work during periods when the retailer posted remote job openings for on-site workers. The odd columns limit to one week on either side of the transition to remote work. The even columns consider the six weeks before and after the transition to remote work, as depicted graphically in the case of calls per hour in Figure 7. Standard errors are clustered at the individual level. The sample is limited to workers who switched to remote work while in entry-level roles and whose transition to remote work did not coincide with a promotion or departure from the retailer within the twelve-week time-span.

Table 5: Reduced-Form Differences: Call Quantity

	Calls/Hour			
	(1)	(2)	(3)	(4)
Remote	-0.316^{***}	-0.260^{***}	-0.234^{**}	-0.237^{**}
	(0.084)	(0.093)	(0.104)	(0.105)
MSA Median Wage in Customer Service		-0.034^{*}	-0.035^{*}	-0.035^{*}
		(0.019)	(0.019)	(0.019)
MSA % in Customer Service			-0.027	-0.022
			(0.057)	(0.057)
Age				-0.002
				(0.002)
Female				0.082
				(0.053)
% Difference	-10.21%	-8.40%	-7.56%	-7.65%
	(2.70)	(3.00)	(3.35)	(3.38)
t x Hire Month x Time-zone	✓	✓	✓	✓
Dependent Mean	3.10	3.10	3.10	3.10
Dependent Std. Dev.	1.50	1.50	1.50	1.50
# Workers	1276	1276	1276	1276
# Remote Workers	442	442	442	442
# On-site Workers	834	834	834	834
# Days	81717	81717	81717	81717

Notes: This table reports the reduced-form differences in calls per hour between remote and on-site workers and the implied differences in worker selection after adjusting for the treatment effect of remote work. Each column estimates equation 10. The sample only includes workers hired into remote jobs or into on-site jobs with an entry wage of $14/hour: Table A.2 considers robustness to including the full sample of on-site workers. The sample is limited to the period between July 2018, when the retailer started to hire directly into remote jobs, and April 2020, when the retailer closed its offices due to COVID-19. Information on the characteristics of the worker's metropolitan statistical area (MSA) come from Emsi, the statistical modeling company used by the retailer to determine call-center locations and wage-rates. When computing the share of the working age population in customer service, the retailer's own workers are excluded. Standard errors on the reduced-form differences are two-way clustered at the worker and date level. Table A.3 decomposes the differences into call duration and time spent on the phone, while Table A.4 considers differences in call quality.

Table 6: Lockdown Differences in Productivity: Call Quantity

	Calls/Hour				
	(1)	(2)	(3)	(4)	(5)
Remote	-0.616^{***}	-0.622^{***}	-0.618^{***}	-0.522^{***}	-0.607^{***}
	(0.151)	(0.155)	(0.154)	(0.175)	(0.186)
COVID-19 Deaths per 100K		-0.001	-0.001	-0.001	-0.001
		(0.001)	(0.001)	(0.001)	(0.001)
COVID-19 Cases per 10K		0.0002	-0.0003	-0.0003	-0.0003
		(0.001)	(0.001)	(0.001)	(0.001)
MSA Median Wage in Customer Service			0.067^{*}	0.075^{*}	0.080^{*}
			(0.039)	(0.042)	(0.041)
MSA % in Customer Service				-0.143	-0.140
				(0.139)	(0.140)
Age					0.001
					(0.006)
Female					0.282^{*}
					(0.148)
% Difference	-18.05%	-18.24%	-18.12%	-15.30%	-17.78%
	(4.42)	(4.53)	(4.52)	(5.14)	(5.46)
t x Hire Month x Time-zone	✓	✓	✓	✓	✓
Dependent Mean	3.41	3.41	3.41	3.41	3.41
Dependent Std. Dev.	1.70	1.70	1.70	1.70	1.70
# Workers	303	303	303	303	303
# Remote Workers	123	123	123	123	123
# On-site Workers	180	180	180	180	180
# Days	11313	11313	11313	11313	11313

Notes: This table analyzes the differences in calls handled per hour between workers who were initially offered remote jobs and those who were initially offered on-site jobs. The sample focuses on the period between April 6, 2020 — when the retailer closed its on-site call-centers — and August 6, 2020. The sample limits to remote workers and on-site workers at the retailer's $14/hour locations, who began in entry-level roles and were in their first six months at the retailer. Each column estimates 11 with standard errors two-way clustered by worker and date.

Table 7: Difference-in-Difference in Calls/Hour Around COVID-19 Office Closures

	Calls/Hr				
	(1)	(2)	(3)	(4)	(5)
Offered On-Site x After Office Closure	0.366***	0.645***	0.284*	0.298**	0.284
	(0.132)	(0.135)	(0.149)	(0.151)	(0.216)
Offered On-Site	0.204*	0.114	0.286**	0.286**	0.347*
	(0.110)	(0.108)	(0.131)	(0.131)	(0.204)
COVID-19 Deaths Per 100K				−0.002	−0.001
				(0.001)	(0.002)
COVID-19 Cases Per 100K				0.0001	0.001
				(0.001)	(0.002)
Child Caregiver					0.074
					(0.153)
Child Caregiver x After Office Closure					−0.134
					(0.167)
% Treatment Effect	9.78%	16.81%	7.57%	7.95%	7.3%
	(3.54%)	(3.52%)	(3.98%)	(4.02%)	(5.55%)
Bandwidth	6 weeks	3 weeks	Donut	Donut	Donut
Pre Mean Calls/Per Hour	2.56	2.43	2.67	2.67	2.79
Post Mean Calls/Per Hour	3.75	3.84	3.75	3.75	3.89
# Workers	368	368	368	368	168
# Formerly On-site, Treated Workers	209	209	209	209	87
# Already Remote, Control Workers	159	159	159	159	81
# Days	14616	7081	10880	10880	7452

Notes: This table evaluates how the productivity of on-site workers who transitioned to remote work compared to that of remote workers who were already working remotely prior to the office closures on April 6, 2020. Each specification estimates equation 13 in the sample of workers who were hired before the last week of February 2020 and had entry-level wages of $14/hour. The sample limits to workers' first six months at the retailer. During the twelve weeks around the office closures, only one of the workers in the sample left the retailer, mitigating concerns about selective attrition. The second column compares the three weeks before and after the office closures, during which times most schools were closed. The third through sixth columns compare the six weeks after the office closures to the three weeks at the end of February and beginning of March before the pandemic's effects on the conditions in the office had set in. Child-care responsibilities in the sixth column come from a caregiving survey that the retailer fielded in June of 2020. All standard errors are clustered by worker.

A Appendix

Figure A.1: Share of American Workers Working Primarily from Home

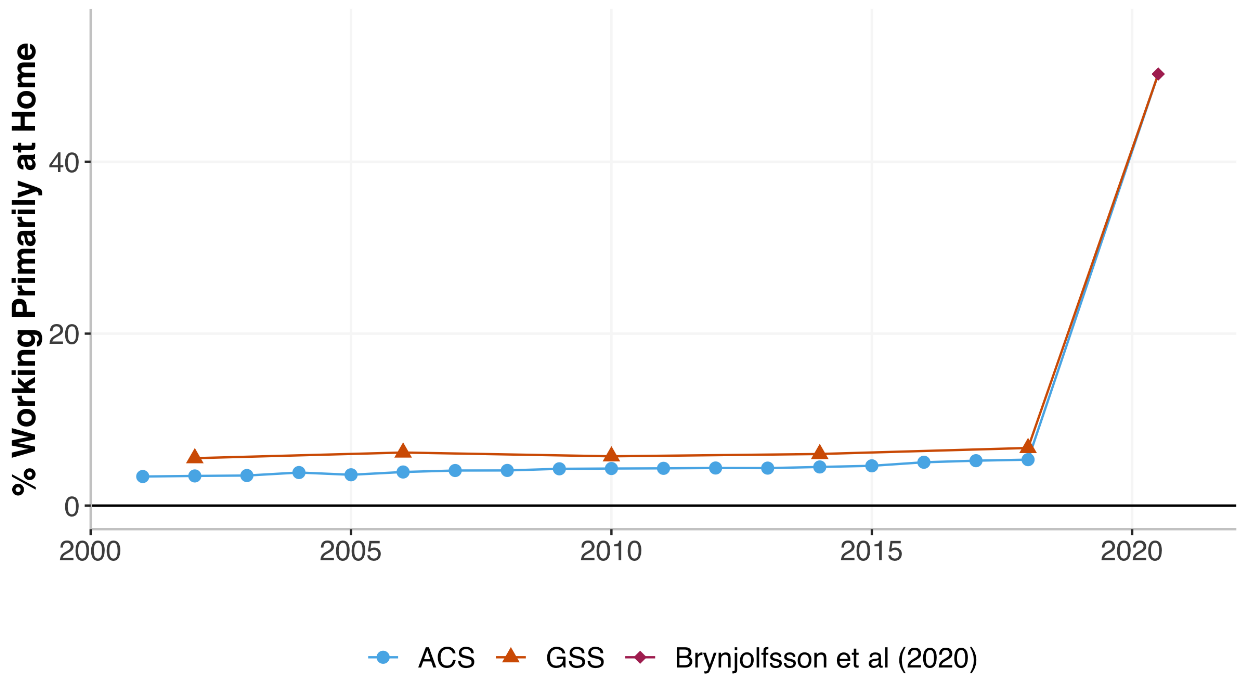

Notes: This figure depicts the share of American workers who report working primarily from home between 2001 and the late spring of 2020. The blue series is based on data from the American Community Survey (N = 21,456,652) which asks respondents about their primary mode of transportation to work in the previous week: this question includes a possible response of "worked at home". The orange series is based on data from the General Social Survey (N = 7,314) which asks employers workers how often they work at home as part of their job: the series plots the share of workers who report working primarily at home. The maroon diamond reports the results from Brynjolfsson et al. (2020) which surveyed fifty thousand workers in April and May 2020 (N=28,000 of whom were employed) to ask them whether they were working remotely. Fully 35% of workers reported being newly remote due to the pandemic, which matches the Bureau of Labor Statistics' estimate of the share of workers teleworking because of the pandemic. Appendix I.A reports the full text of these questions and the share of respondents who reported each answer.

Table A.1: On-Site and Remote Workers in the American Community Survey

	All Workers	On-Site Work	At Home Work	Δ
% Female	46.5	50.8	46.3	4.5***
% Age $\geq$40	50.4	62.4	49.8	12.7***
Age	40.0	43.6	39.9	3.7***
Caregiving				
% Child	42.8	46.7	42.6	4.1***
% Child $\leq$5	12.1	12.4	12.0	0.4***
Caregiving Among Women				
% Child Among Women	45.6	50.6	45.4	5.2***
% Child $\leq$5 Among Women	11.4	13.2	11.4	1.9***
Disability				
% Physical Disability	1.7	2.1	1.7	0.5***
% Cognitive Disability	1.7	1.8	1.7	0.0
% Part-time	23.3	29.1	23.0	6.1***
# Workers	4588693	222583	4366110	

Notes: This table characterizes on-site and remote workers in the 2015-2019 American Community Surveys. The sample limits to employed workers aged 18 and 64, who have less than four years of college. These restrictions make the sample more comparable to the call-center workers at the retailer. The first column considers all workers; the second column, workers who commuted to an on-site job in the previous week; the third column, workers who were working at home in the previous week. The fourth column presents the differences between on-site and at-home workers with standard errors clustered at the individual level. All statistics are computed with population weights from the Census. These statistics complement those in the retailer in Table 1.

Figure A.2: Distribution of performance evaluations remote and on-site in December 2018

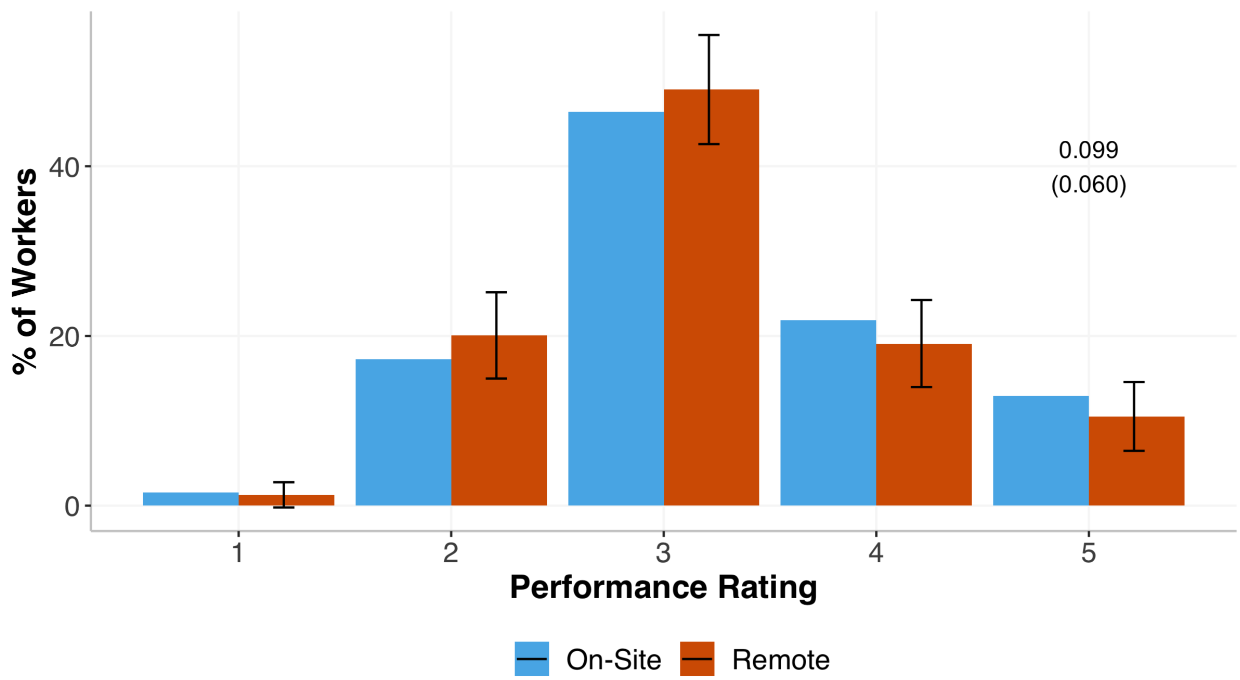

*Notes:*This figure plots the distribution of performance evaluations in December 2018 for on-site workers in blue and remote workers in orange. The x-axis indicates the performance rating from one to five — with one signifying unsatisfactory performance; two, inconsistent performance; three, meeting current expectations; four, highly effective performance; and five, exceptional performance. The y-axis indicates the percent of workers who received each review. The error bars represent the 95% confidence intervals of the comparison of the percent of remote and on-site workers who receive the given review. The annotated coefficient represents the average difference in review. Standard errors are all clustered at the worker level.

Figure A.3: Distribution of performance evaluations remote and on-site in June 2019

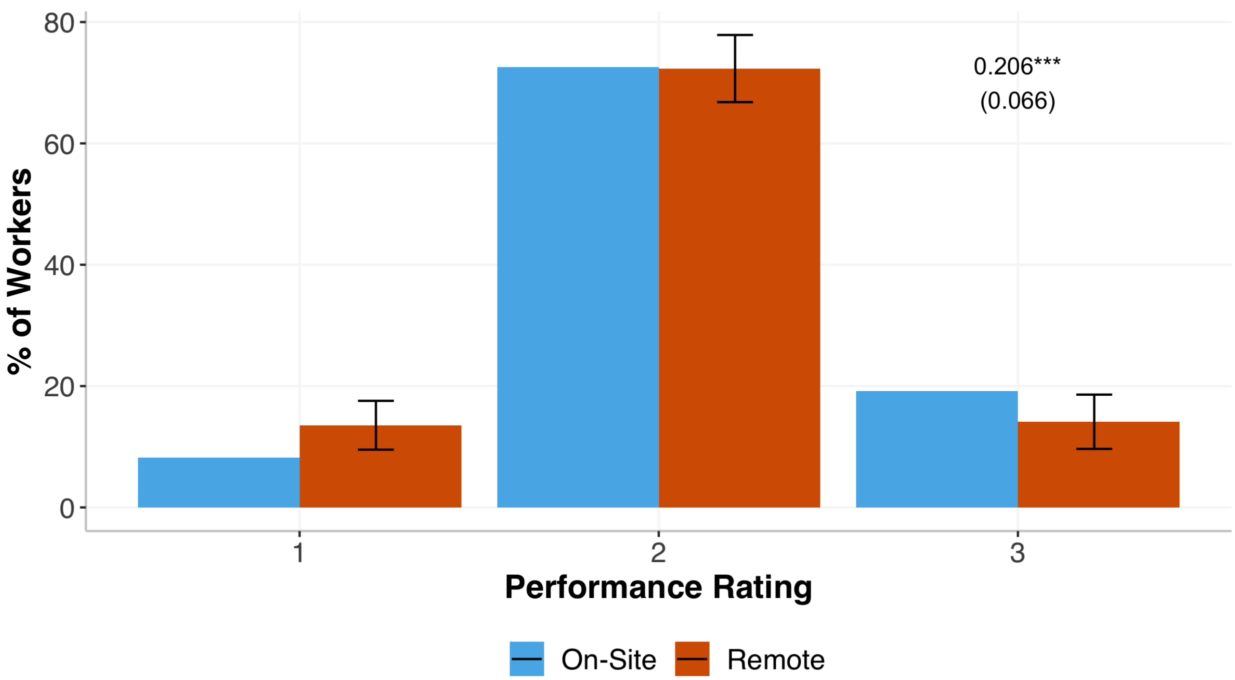

Notes: This figure plots the distribution of performance evaluations in June 2019 for on-site workers in blue and remote workers in orange. The x-axis indicates the performance rating from one to three — with one signifying unsatisfactory performance; two, meeting current expectations; and three, exceptional performance. The y-axis indicates the percent of workers who received each review. The error bars represent the 95% confidence intervals of the comparison of the percent of remote and on-site workers who receive the given review. The annotated coefficient represents the average difference in review. Standard errors are all clustered at the worker level.

Figure A.4: The Demand for Remote Work Among High- and Low-Ability Workers

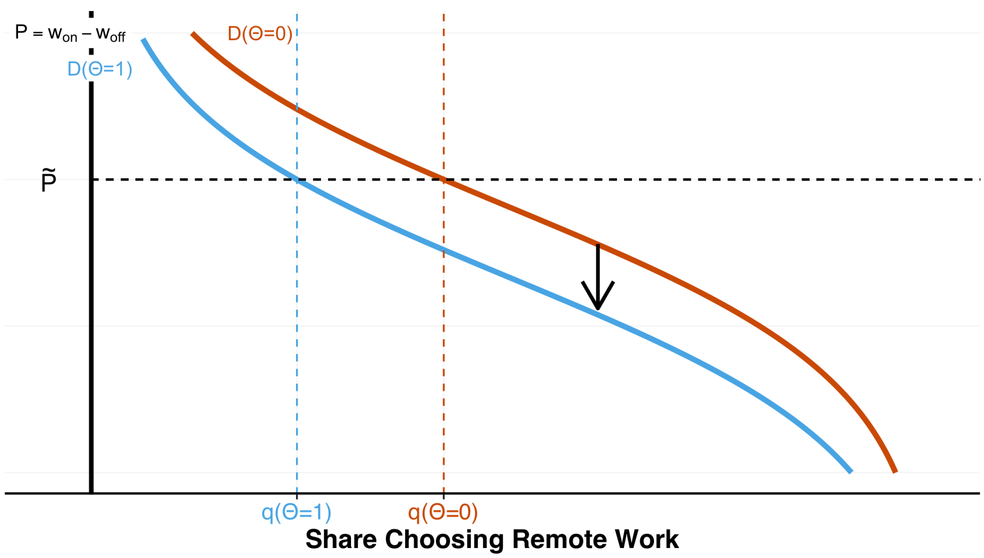

Notes: This figure illustrates worker choices in the model. The x-axis plots the share of workers with a given ability signal who choose remote work. The y-axis plots the price or wage penalty for remote work. The orange line graphs the demand for remote work among those who learn that they will not merit promotion, who we call "low ability". The blue line graphs the demand for those who learn they will certainly merit promotion, who we call "high ability". For low-ability workers, being overlooked has no career costs. For high-ability workers, being overlooked always means a missed promotion. These career costs shift down the demand of high-ability workers, who must pay a higher implicit price of remote work at any given wage penalty. This downward shift means that a high-ability worker who chooses remote work must have stronger tastes for remote work than a low-ability worker at any given price. Further, because of this downward shift, the share of low-ability workers who choose remote work will exceed the share of high-ability workers who choose remote work at any given price, e.g. $\tilde{P}$. The shape of each demand curve reflects the logistic distribution of preferences. At the extremes, price changes operate on the thin tails of the distribution and, thus, cannot lead to large changes in quantity, leading to an inelastic demand curve. In the center of the distribution, price changes operate on higher density leading to greater elasticity.

Figure A.5: Weeks On-Site for Remote Workers who were Offered Remote Work (Later Cohorts) Versus Workers who were not (Earlier Cohorts)

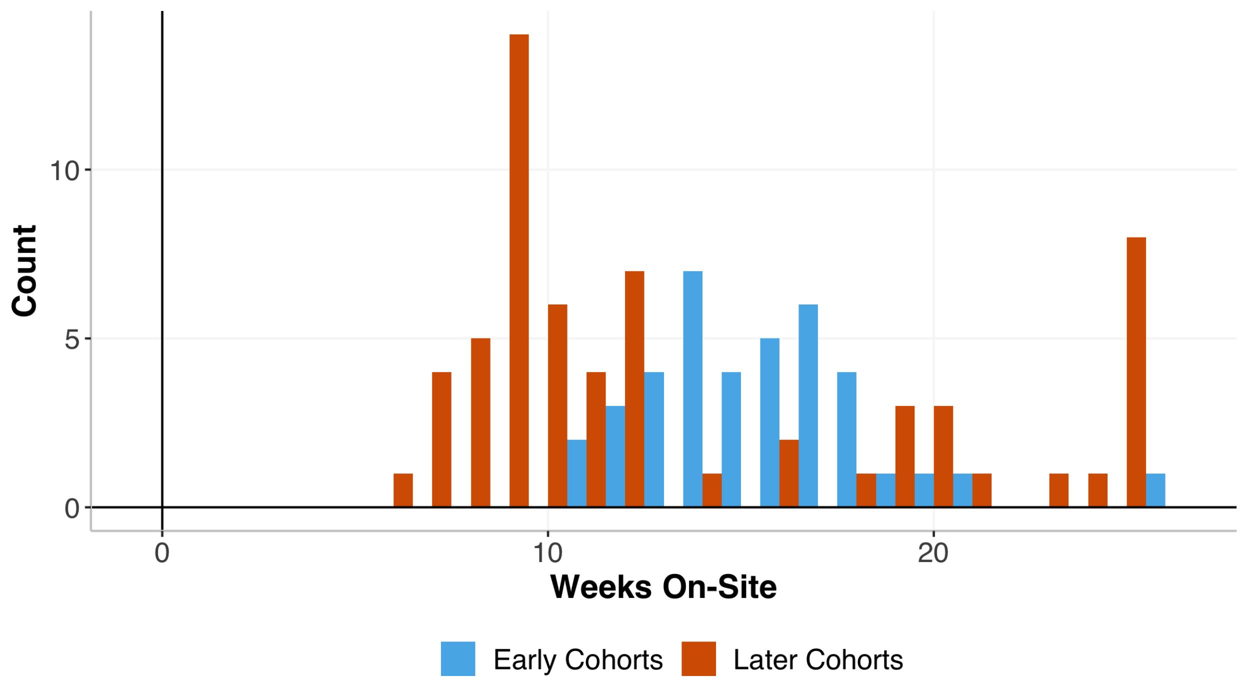

Notes: This figure analyzes the introduction of opportunities to go remote into the choice sets of new hires. The x-axis depicts the hiring month. The vertical dashed line highlights when the retailer introduced opportunities for on-site workers to go remote to open up desks in the call-centers. The y-axis depicts the difference in logged calls per hour of remote and on-site workers from each cohort. The point estimates come from specification 7, estimated with date by hiring month by time-zone fixed effects. The ribbons reflect 95% confidence intervals with standard errors two-way clustered by date and worker. The sample limits to workers who were hired into entry-level roles and are in their first 6 months at the retailer to exclude workers who have advanced to specialized roles with non-random routing of calls.

Figure A.6: The time-series of the share of hired into remote job at the retailer

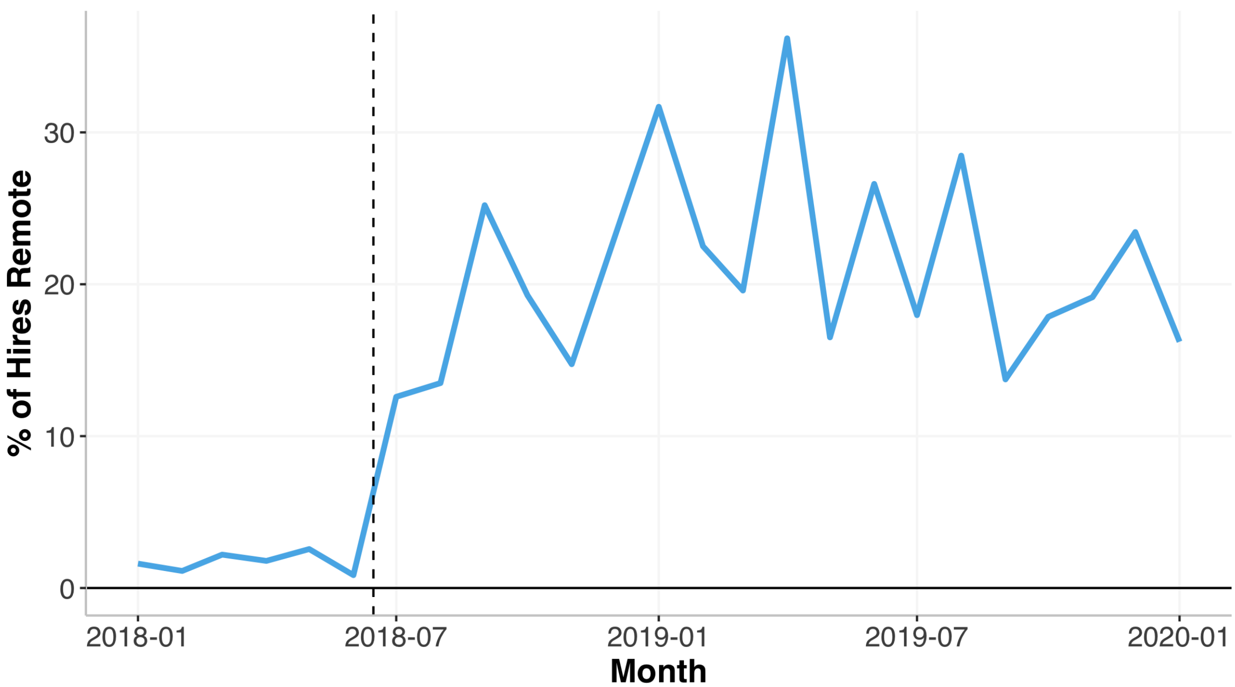

This figure depicts the time-series of the share of the retailer's new hires who started in remote jobs. Prior to July of 2018, workers could only start in remote jobs if they had experience outside the company. In July of 2018 — indicated in the vertical dashed line — the retailer started to hire entry-level workers directly into remote jobs if they lived far from one of the retailer's physical call-centers.

Table A.2: Reduced-Form Differences: Call Quantity in the Full Sample of Workers

	Calls/Hour			
	(1)	(2)	(3)	(4)
Remote	-0.288^{***}	-0.210^{***}	-0.158^{*}	-0.150^{*}
	(0.062)	(0.079)	(0.090)	(0.090)
MSA Median Wage in Customer Service		-0.036^{**}	-0.040^{**}	-0.038^{**}
		(0.018)	(0.018)	(0.018)
Retailer Entry $/hour		0.031	0.053	0.053
		(0.027)	(0.033)	(0.033)
MSA % in Customer Service			-0.048	-0.036
			(0.044)	(0.043)
Age				-0.004^{***}
				(0.001)
Female				0.056
				(0.034)
% Difference	-9.27%	-6.76%	-5.08%	-4.83%
	(1.98)	(2.53)	(2.88)	(2.91)
t x Hire Month x Time-zone	✓	✓	✓	✓
Dependent Mean	3.11	3.11	3.11	3.11
Dependent Std. Dev.	1.48	1.48	1.48	1.48
# Workers	2391	2391	2391	2391
# Remote Workers	442	442	442	442
# On-site Workers	1949	1949	1949	1949
# Days	146957	146957	146957	146957

Notes: This table complements the analysis in Table 5, by including the full sample of on-site workers at the retailer. Each column estimates the reduced-form difference in call quality from equation 10. The sample is limited to the period between July 2018, when the retailer started to hire directly into remote jobs, and April 2020, when the retailer closed its offices due to COVID-19. Information on the characteristics of the worker's metropolitan statistical area (MSA) come from Emsi. Standard errors on the reduced-form differences are two-way clustered at the worker and date level.

Table A.3: Decomposing Reduced-Form Differences in Call Quantity

	Min/Call	% Time on Phone	On Hold Min/Call
Remote	1.006^{**}	-0.007	-0.047
	(0.398)	(0.010)	(0.068)
MSA Median Wage in Customer Service	-0.035	-0.0005	-0.010
	(0.086)	(0.002)	(0.018)
% Difference	9.46%	-1.59%	-3.98%
	(3.74)	(2.17)	(5.73)
t x Hire Month x Time-zone	✓	✓	✓
Dependent Mean	10.63	0.46	1.19
Dependent Std. Dev.	5.09	0.19	2.34
# Workers	1209	1209	1209
# Remote Workers	420	420	420
# On-site Workers	789	789	789
# Days	65121	65121	65121

Notes: This table complements the analysis in Table 5, by decomposing the aggregate difference in calls per hour into minutes spent per call in column one and total time on the phone in column two. The third column considers the hold minutes per call, which accounts for 11% of the typical call (1.19/10.63). The sample only includes workers hired into remote jobs or into on-site jobs with an entry wage of $14/hour. The sample is limited to the period between January 2019, when detailed time-use data started to be recorded, and April 2020, when the retailer closed its offices due to COVID-19. The median wage in customer service in each worker's MSA comes from Emsi. Standard errors are two-way clustered at the worker and date level.

Table A.4: Reduced-Form Differences: Call-Quality

	Customer Satisfaction			
	(1)	(2)	(3)	(4)
Remote	0.004	0.010	0.003	0.002
	(0.007)	(0.007)	(0.008)	(0.008)
MSA Median Wage in Customer Service		-0.004^{**}	-0.003^{*}	-0.003^{**}
		(0.002)	(0.002)	(0.002)
MSA % in Customer Service			0.007	0.006
			(0.005)	(0.005)
Age				0.0005^{**}
				(0.0002)
Female				-0.004
				(0.004)
% Total Effect	0.08%	0.20%	0.06%	0.05%
	(0.14)	(0.14)	(0.17)	(0.17)
% Selection Effect	-0.47%	0.14%	0.00%	-0.01%
	(0.57)	(0.57)	(0.58)	(0.58)
t x Hire Month x Time-zone	✓	✓	✓	✓
Dependent Mean	4.89	4.89	4.89	4.89
Dependent Std. Dev.	0.38	0.38	0.38	0.38
# Workers	1222	1222	1222	1222
# Remote Workers	429	429	429	429
# On-site Workers	793	793	793	793
# Days	67688	67688	67688	67688

Notes: This table complements the analysis in Table 5, by considering call quality rather than quantity. Each column estimates the reduced-form difference in call quality according to a version of equation 10. The sample only includes workers hired into remote jobs or into on-site jobs with an entry wage of $14/hour. The sample is limited to the period between July 2018, when the retailer started to hire directly into remote jobs, and April 2020, when the retailer closed its offices due to COVID-19. Each observation is the average customer satisfaction of the worker during the day. Since most calls are not reviewed, workers do not receive reviews everyday. Each column is weighted by the number of customer satisfaction reviews of that worker on the day. Information on the characteristics of the worker's metropolitan statistical area (MSA) come from Emsi. Standard errors on the reduced-form differences are two-way clustered at the worker and date level.

Table A.5: Lockdown Differences in Productivity: Call Quantity

	Calls/Hour		
	(1)	(2)	(3)
Remote	-0.608^{***}	-0.768^{***}	-0.769^{***}
	(0.166)	(0.218)	(0.223)
COVID-19 Deaths per 100K	-0.001	-0.001	-0.001
	(0.001)	(0.001)	(0.001)
COVID-19 Cases per 10K	-0.0005	-0.0001	-0.0002
	(0.001)	(0.001)	(0.001)
MSA Median Wage in Customer Service	0.095^{**}	0.087^{*}	0.089^{*}
	(0.037)	(0.052)	(0.053)
MSA % in Customer Service	-0.149	-0.089	-0.088
	(0.128)	(0.169)	(0.169)
Age	0.003	0.007	0.007
	(0.006)	(0.007)	(0.007)
Female	0.215	0.200	0.200
	(0.133)	(0.188)	(0.186)
Child Caregiver			0.085
			(0.126)
% Selection Effect	-16.55%	-20.67%	-20.69%
	(4.53)	(5.87)%	(5.99)%
t x Hire Month x Time-zone	✓	✓	✓
Dependent Mean	3.67	3.71	3.71
Dependent Std. Dev.	1.67	0.20	0.20
# Workers	302	148	148
# Remote Workers	123	66	66
# On-site Workers	179	82	82
# Days	9661	5789	5789

Notes: This table complements Table 6 by controlling for child-care responsibilities when comparing productivity between workers who were initially offered remote jobs and those who were initially offered on-site ones during COVID-19's lockdown. Caregiving responsibilities come from a June 2020 survey conducted by the retailer. The participation rate was about half in this sample, with higher rates of participation among remote workers. Each column estimates 11 with standard errors two-way clustered by worker and date.

Table A.6: Lockdown Differences in Productivity: Call-Quality

	Customer Satisfaction			
	(1)	(2)	(3)	(4)
Remote	0.012	0.013	0.013	0.009
	(0.015)	(0.016)	(0.017)	(0.018)
MSA Median Wage in Customer Service		−0.003	−0.003	−0.004
		(0.003)	(0.003)	(0.003)
MSA % in Customer Service			0.00004	0.002
			(0.017)	(0.017)
Age				0.001*
				(0.001)
Female				−0.007
				(0.015)
% Selection Effect	0.25%	0.28%	0.28%	0.18%
	(0.32)	(0.32)	(0.35)	(0.37)
t x Hire Month x Time-zone	✓	✓	✓	✓
Dependent Mean	4.88	4.88	4.88	4.88
Dependent Std. Dev.	0.33	0.33	0.33	0.33
# Workers	284	284	284	284
# Remote Workers	110	110	110	110
# On-site Workers	174	174	174	174
# Days	4521	4521	4521	4521

Notes: This table complements Table 6 by comparing the call quality rather than call quantity of workers initially offered remote jobs versus those initially offered on-site ones. The sample focuses on the period between April 6, 2020 — when the retailer closed its on-site call-centers — and August 6, 2020. During this time, all workers were working remotely because of COVID-19; only one worker in the sample left the retailer, mitigating concerns about selective attrition. The sample limits to remote workers and on-site workers at the retailer's $14/hour locations, who began in entry-level roles and were in their first six months at the retailer. Each column estimates 11 with standard errors two-way clustered by worker and date.

Table A.7: Difference-in-Difference in Calls/Hour Around COVID-19 Office Closures: Additional Demographic Controls

	Calls/Hr			
	(1)	(2)	(3)	(4)
Offered On-Site x After Office Closure	0.284*	0.315**	0.325**	0.240
	(0.149)	(0.158)	(0.157)	(0.164)
Offered On-Site	0.286**	0.377***	0.375***	0.350**
	(0.131)	(0.134)	(0.133)	(0.142)
Worker Age		0.007	0.004	0.007
		(0.005)	(0.007)	(0.005)
Worker Age x After Office Closure		−0.002	0.006	−0.002
		(0.006)	(0.009)	(0.006)
Male		−0.317***	−0.318***	−0.469***
		(0.102)	(0.103)	(0.151)
Male x After Office Closure		0.006	0.008	−0.498**
		(0.148)	(0.149)	(0.243)
Offered On-Site x After Office Closure x Age			−0.012	
			(0.012)	
Offered On-Site x Age			0.005	
			(0.011)	
Offered On-Site x After Office Closure x Male				0.637**
				(0.299)
Offered On-Site x Male				0.200
				(0.196)
% Treatment Effect	7.57%	8.41%	8.67%	6.4%
	(3.98%)	(4.22%)	(4.2%)	(4.37%)
COVID-19 Controls	✓	✓	✓	✓
# Workers	368	368	368	368
# Days	10880	10880	10880	10880

Notes: This table complements the analysis in Table 7, with additional demographic controls and interactions. Each specification estimates equation 13 in the sample of workers who were hired before the last week of February 2020 and had entry-level wages of $14/hour. The sample limits to workers' first six months at the retailer. All standard errors are clustered by worker. Age is included as the difference from the mean age in the sample. The second column allows for different productivity changes by gender and age. The third and fourth columns interact the treatment effect with age (column three) and gender (column four). There are 209 treated workers, with 63 men, and 159 control workers, with 22 men.

Figure A.7: Change in Remote vs On-site Productivity In the Tail End of the 2020 Holiday Rush

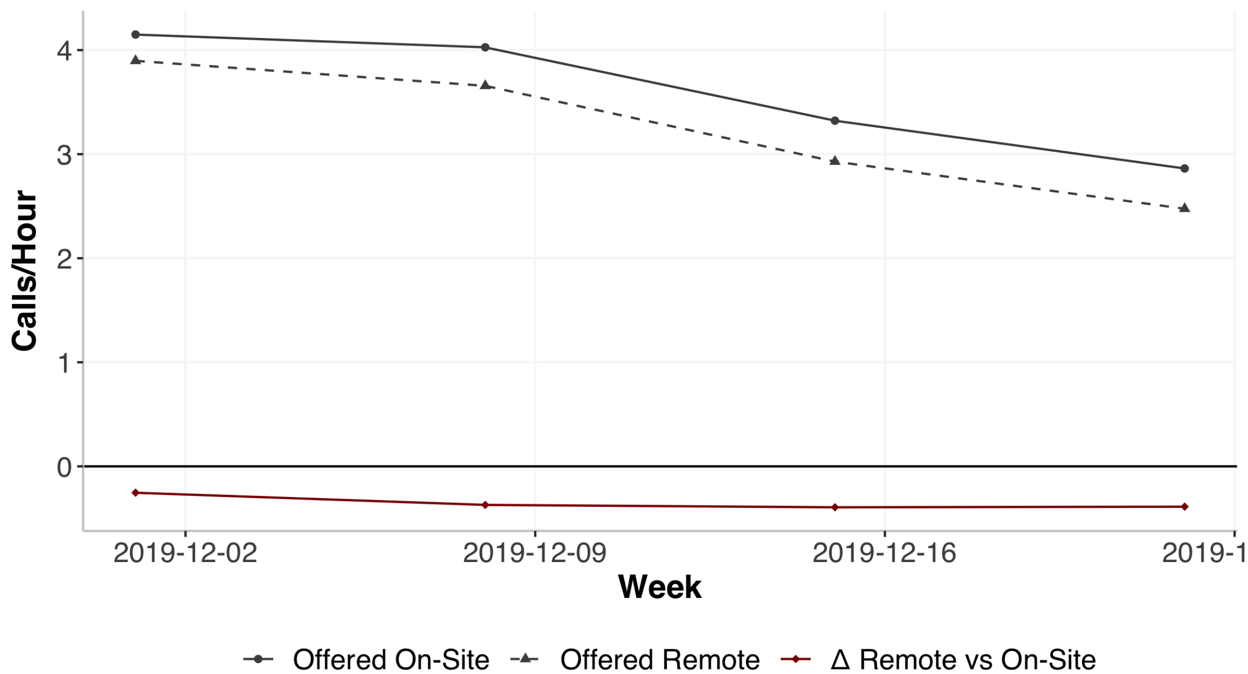

Notes: This figure illustrates the evolution of the gap in productivity between on-site workers (in the solid line) and remote workers (in the dashed line) during the tail end of the previous holiday rush. Since this period featured a large change in consumer demand but no change in remote work, this analysis serves as a placebo check for the analysis in Figure 9. The x-axis indicates the week in which calls were taken. The y-axis represents calls per hour. Each point is the average hourly calls of either formerly on-site workers (in circles and solid lines) or already remote workers (in triangles and dashed lines). The maroon line plots the difference in calls per hour between remote and on-site workers. The sample includes workers who were hired before the first week of December 2019 and had entry-level wages of $14/hour. The sample limits to workers' first six months at the retailer.

Table A.8: Difference-in-Difference in Customer Satisfaction Around COVID-19 Office Closures

	Customer Satisfaction (out of 5)				
	(1)	(2)	(3)	(4)	(5)
Offered On-Site x After Office Closure	-0.021	-0.005	-0.019	-0.024	0.032
	(0.021)	(0.027)	(0.025)	(0.025)	(0.038)
Offered On-Site	-0.005	-0.003	-0.007	-0.007	-0.059
	(0.018)	(0.020)	(0.023)	(0.023)	(0.039)
COVID-19 Deaths Per 100K				0.0003	0.00003
				(0.0002)	(0.0003)
COVID-19 Cases Per 100K				-0.0003^{*}	-0.0001
				(0.0002)	(0.0002)
Child Caregiver					-0.042^{*}
					(0.025)
Child Caregiver x After Office Closure					0.019
					(0.026)
% Treatment Effect	-0.43%	-0.11%	-0.39%	-0.49%	0.64%
	(0.43%)	(0.55%)	(0.51%)	(0.51%)	(0.77%)
Bandwidth	6 weeks	3 weeks	Donut	Donut	Donut
Pre Mean Calls/Per Hour	4.91	4.91	4.90	4.90	4.90
Post Mean Calls/Per Hour	4.88	4.88	4.88	4.88	4.89
# Workers	359	359	359	359	163
# Formerly On-site, Treated Workers	202	202	202	202	82
# Already Remote, Control Workers	157	157	157	157	81
# Days	10176	6007	7190	7190	5082

Notes: This table complements Table 9 by considering how call quality rather than call quantity changed for on-site workers transitioning to remote work compared to remote workers who were already working remotely. Each specification estimates equation 13 in the sample of workers who were hired before the last week of February 2020 and had entry-level wages of $14/hour. The sample limits to workers' first six months at the retailer. During the twelve weeks around the office closures, only one of the workers in the sample left the retailer, mitigating concerns about selective attrition. The second column compares the three weeks before and after the office closures, during which times most grade school were closed. The third through sixth columns compare the six weeks after the office closures to the three weeks at the end of February and beginning of March before the pandemic's effects on the conditions in the office had set in. Child-care responsibilities in the sixth column come from a caregiving survey that the retailer fielded in June of 2020. All standard errors are clustered by worker.

Figure A.8: Decomposing Difference-in-Difference in Calls/Hour Around COVID-19 Office Closures

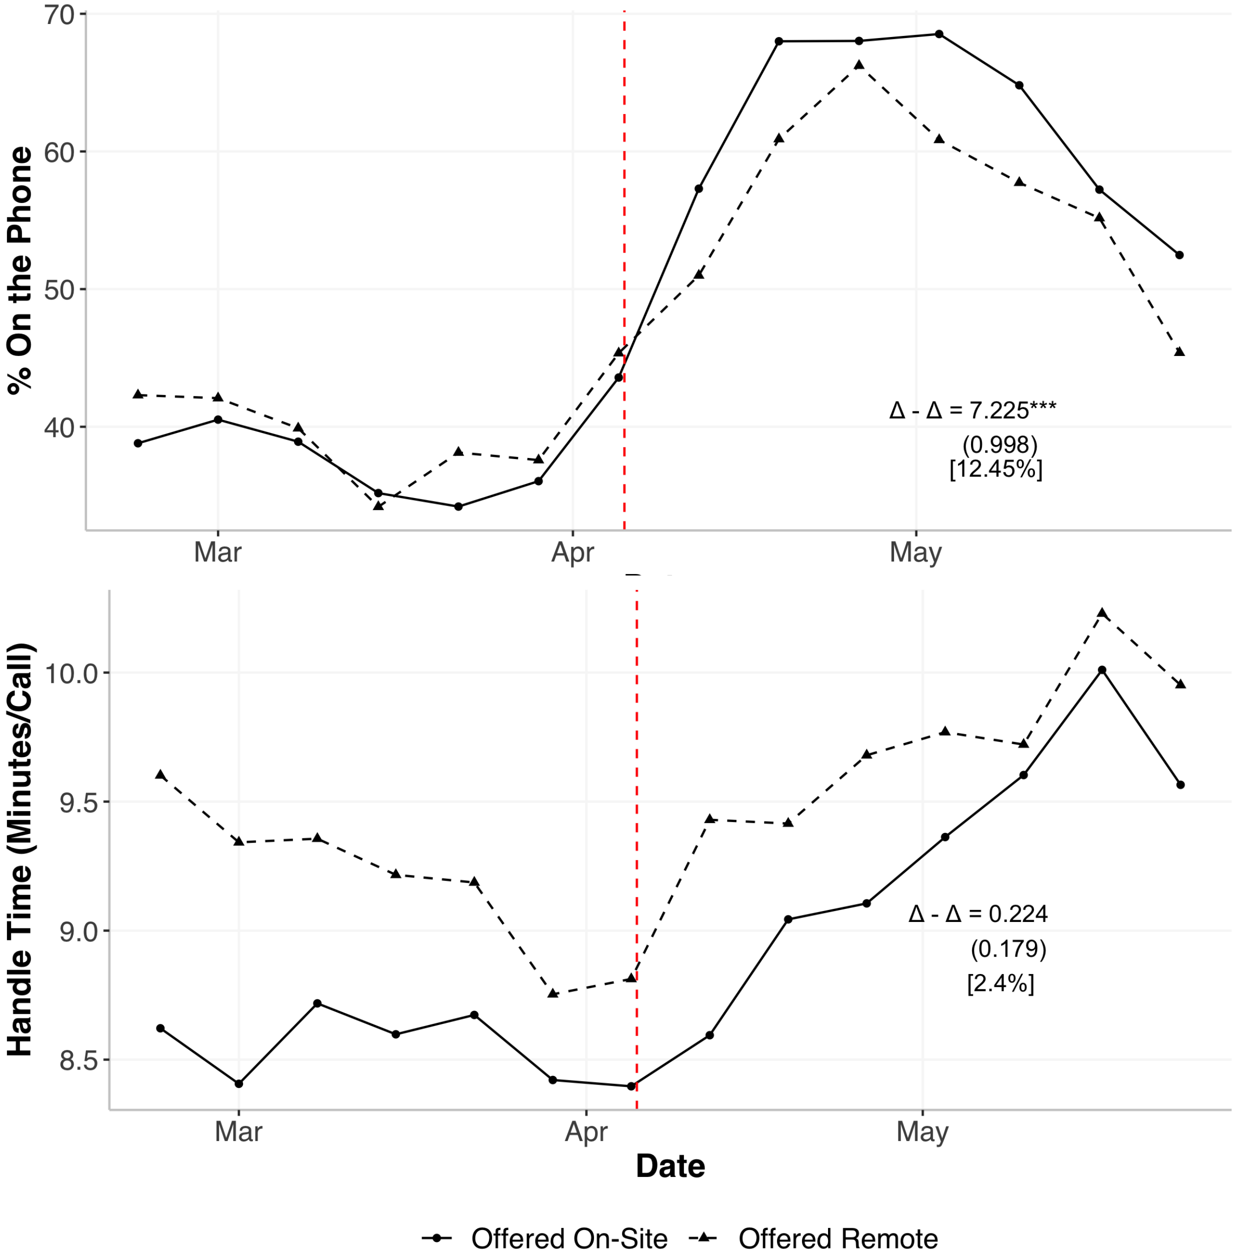

Notes: This figure decomposes the difference-in-difference in calls per hour in Figure 9 into the percentage of time workers were on the phone (in the top panel) and the speed with which they answered calls (in the bottom panel). The x-axis represents the week in which calls were taken. The y-axis represents calls per hour. The vertical red dashed line indicates April 6, 2020, the date on which the retailer closed its on-site call-centers due to COVID-19. Each point is the average hourly calls of either formerly on-site workers (in circles and solid lines) or already remote workers (in triangles and dashed lines). The sample includes workers who were hired before the last week of February 2020 and had entry-level wages of $14/hour. The sample limits to workers' first six months at the retailer. During this time, only one of the workers in the sample left the retailer, mitigating concerns about selective attrition. The annotated coefficient compares the difference between remote and on-site workers after the office closures to that before the office closures. The bracketed term compares difference-in-difference to the level after the COVID-19 closures. All standard errors are clustered by worker.

I.A Text of Questions about Remote Work

American Community Survey Question (2001-2018) [N = 21,456,652]: What was your primary means of transportation to work in the previous week?

Response	% of Workers	% of Phone Workers
Worked at home	4.26%	3.20%
Private motorized vehicle	86.6%	87.8%
Public transport	4.92%	5.44%
Taxi	0.14%	0.12%
Bicycle	0.51%	0.42%
Walked only	2.68%	2.34%
Other	0.91%	0.70%
N	21,456,652	394,353

The General Social Survey Question (2002, 2006, 2010, 2014, 2018) [N = 7,314]: How often do you work at home as part of your job?

Response	% of Workers	% of Phone Workers
Work mainly at home	6.10%	4.72%
Frequently (More than once a week/About once a week)	19.5%	6.3%
Rarely (About once a month/A few times a year)	14.2%	4.7%
Never	60.2%	81.1%
N	7,314	127

For those who report working at home some but not all the time, 41% say that they had to in order to keep up with their jobs; only 32% of these part-time work-from-homers said they worked from home because they wanted to; and another 27% report that they do so for another reason (e.g. a doctor's appointment during the day). A significant share of the time that most people were spending working at home before the pandemic was at times that they did not want to be working at home and were instead doing so to meet a deadline or accommodate another constraint. Among those who were mainly working from home, 47% were doing so because they wanted to, 48% were doing so for another reason, and 6% were doing so to keep up with their job.

Brynjolfsson et al. (2020) in April and May 2020 [N = 28,000]: Have you started to work from home in the last 4 weeks?

Response	% of Workers
Used to commute, now work from home	35.2%
Used to work from home and still do	15.0%
I continue to commute to work	37.1%
I have recently been furloughed or laid-off	10.1%
Used to work from home, but now I commute	2.5%
N	28,000

I.B Generalizing the interpretation of the model

We have focused on the possibility that on-site work increases transparency about worker ability, making it more desirable for more productive workers. Our model, however, generalizes to apply to other explanations for why the office might be complementary with ability. Three salient alternatives to information are (a) political capital, (b) human capital, and (c) worker preferences.

First, let's consider the political capital story. Suppose workers are only promoted if they (a) perform well and (b) develop political capital with their bosses. Remote workers may find it more difficult to build this political capital. If so, better workers will gravitate to the office because political capital is more pivotal for promotion. This mechanism will consequently generate patterns of selection similar to better information about on-site workers. Political capital and misinformation are difficult to distinguish empirically since one is about what managers know and the other is about what managers choose to do with their information. However, Section III offers suggestive evidence that mangers' reviews are more predictive of performance on-site than remote, suggesting a role for information and not just patronage.

Second, let's consider the role of human capital. Suppose workers are only promoted if they (a) perform well and (b) acquire skills that are more difficult to develop when working remotely. Since the returns to these new skills are higher for high ability workers, this will draw better workers into the office. It will thus generate patterns of selection similar to misinformation about remote workers. The parallel trends of learning in Figure 1 offers suggestive evidence that this is not the right story in our context. However, there may be aspects of worker skills that do not translate into their speed at answering calls but are necessary for promotion that these trajectories do not pick up.

Third, another complementarity between ability and office work could stem directly from preferences. In our baseline model, a component of a worker's demand for remote work reflects her idiosyncratic tastes for remote work and another reflects her career concerns that track her likely ability. Another force that creates a correlation between worker demand for remote work and worker ability will generate a similar dynamic as in Figure **??**. For example, suppose worker's drive contributed to both a worker's ability and her eagerness to go to the office. The productivity of workers on the margin of remote work would then be increasing, as those on the margin became increasingly keen to stay in the office and, thus, increasingly driven. Similarly, if workers seek promotions not only for their pecuniary value but also for personal validation and social approbation, then this would contribute to the component of demand that related to ability and strengthen the selection mechanism in the model for reasons that were not purely instrumental.