“WORKING” REMOTELY?

SELECTION, TREATMENT, AND THE MARKET FOR REMOTE WORK

Natalia Emanuel · Emma Harrington[1]

June 22, 2021

Abstract: Why are firms reluctant to hire remote workers? One possibility is adverse selection: remote jobs may attract workers looking to hide their ability while office jobs attract those looking to demonstrate it. We test this theory in the call-centers of a Fortune 500 retailer. Introducing a remote-work program in 2018 attracted less productive workers. Likewise, during COVID-19's shutdown, workers who had originally chosen remote jobs were 14.9% less productive than those who had originally chosen on-site jobs. Thus, hiring remote workers was costly even though working remotely increased workers' productivity by 7-8% during the remote-work program and COVID-19's shutdown.

[1]Contact: Princeton University, 20 Washington Rd, Princeton, NJ 08544, emma.k.harrington4@gmail.com. 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 6% of Americans worked remotely all of the time.[1] A few months into the pandemic, the majority of Americans were doing so (Brynjolfsson et al., 2020).[2] Why were so few jobs remote before the pandemic and why might many return to the office? One possibility is moral hazard — workers may shirk when out of sight of managers. Another possibility is adverse selection — workers who want to hide low productivity may choose to be remote while those who want to reveal high productivity choose to be on-site.

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. 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 before the pandemic, a high willingness to pay suggests remote work is costly for firms. However, an RCT in a Chinese call-center reveal no such costs, with remote work increasing productivity by 14% (Bloom et al., 2015). Thus, we are left with a puzzle of why remote work was rare.

During the pandemic, the same puzzle emerged as most workers reported being happier and more productive working remotely but few firms advertised jobs that would be permanently remote (Barrero et al., 2020; Ovide, 2021).[3]

[1]In the 2019 American Community Survey (ACS), 5.6% of workers reported working from home, based on the authors' calculations (U.S. Census Bureau, 2021). 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 on the day of the survey (Brynjolfsson et al., 2020; Bureau of Labor Statistics, 2020a).

[2]These estimates are consistent with those from the Bureau of Labor Statistics, where 35% of workers report working remotely because of the pandemic (Bureau of Labor Statistics, 2020b). However, even during the pandemic, fewer workers reported working remotely over time.

[3]Surveys of workers find 32% to 45% want to remain fully remote after the pandemic (PwC, 2020; Morning Consult, 2020; Ovide, 2021). However, ZipRecruiter finds only 8 or 9% of jobs are

We argue the missing piece to this puzzle is adverse selection, which increases the cost to firms of hiring remote workers instead of on-site ones even for a job well-suited to remote work.

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 model that ties promotion penalties to the selection of workers who choose remote jobs. Our model predicts that differences in promotion will lead to differences in worker selection. We test this prediction empirically in 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 our analysis into three parts.

The first part offers stylized facts about promotion in remote jobs and develops a model where career concerns shape the market for remote work. At our retailer, workers who chose remote jobs had about half of the promotion chances as those who chose on-site jobs, consistent with remote work's impact on promotion in Bloom et al. (2015)'s RCT. Further, managers appear to be less certain about remote workers' productivity: managers' evaluations are less predictive of the future performance of remote workers than on-site workers. Thus, in the model, remote work reduces the probability that firms learn about workers' abilities. Latently low-ability workers consequently sort into remote jobs to hide their ability while latently high-ability workers sort into on-site jobs to reveal high ability.[4] The resulting adverse selection into remote work raises its average cost above its marginal cost, causing remote work to be under-provided.[5]

actually permanently remote, up just 6pp from before the pandemic (Ovide, 2021).

[4]Others have similarly argued that management practices and performance pay can induce better worker selection (Lazear, 2000; Bender et al., 2018; Brown and Andrabi, 2020).

[5]Our model is most similar to Einav et al. (2010) but also shares features of classical labor market

In the second part of the paper, we identify remote work's costs to the firm. On average, remote workers were 8% less productive than on-site workers before COVID-19. Natural experiments at the retailer allow us to decompose this difference into selection and treatment effects to test whether remote work leads to adverse selection or moral hazard.

Identifying remote work's selection effect has been challenging because workers who choose remote and on-site jobs are typically in different roles, often at different firms. Our setting is unusual in that the retailer hired workers into remote and on-site jobs and randomly routed calls between them.

During COVID-19's lockdown, everyone at the retailer worked remotely. Those who originally chose remote jobs answered 15% fewer calls per hour than those who originally chose on-site jobs, indicating adverse selection into remote jobs.[6]

The retailer's introduction of a remote work program in 2018 also identifies the selection effect of remote work. Among workers who ultimately took up opportunities to go remote, some were hired before the program's introduction and others were hired after. Only later hires could have chosen the job because of the offer of remote work. Accordingly, adverse selection only shows up for these later hires. Workers who went remote were 12.2% less productive than their on-site peers in later cohorts but 8.8% more productive in earlier cohorts. The difference-in-differences suggests that offering remote work attracted workers who were 21% less productive. This design complements Linos (2018), which finds similar patterns around the US Parent Office's introduction of a remote work program.

The natural experiments at the retailer also identify remote work's treatment ef-

models of adverse selection (Salop and Salop, 1976; Miyazaki, 1977; Weiss, 1995).

[6]This productivity difference is not driven by caregiving, suggesting career concerns rather than constraints at home drive adverse selection. Adams-Prassl (2020) finds that women working on MTurk who have an infant at home are more likely to have work interruptions. Our results suggest this may not be the case for all parents.

fect. We first leverage the quasi-random timing of workers' transitions to remote work in the 2018 remote work program. We find productivity rose by 7.6% when workers went remote with no sacrifice in customer satisfaction reviews. This estimate is consistent with the positive treatment effect estimated for those who opt into remote work in Chinese call-centers (Bloom et al., 2015) and the US Patent Office (Choudhury et al., 2020).

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 6.6% relative to that of their already remote peers. This treatment effect applies to workers who do not necessarily want to work remotely.[7]

Our setting does not allow us to speak to tasks that hinge on coordination (Battiston et al., 2017; Gibbs et al., 2021) or intense concentration (Künn et al., 2020), where less positive effects of remote work have been found.[8] For such tasks, both the treatment effect and selection effect of remote work may contribute to its rarity.

In the third part of the paper, we pull these estimates together to quantify the distortion from adverse selection. We estimate workers' demand for remote work using the retailer's policy of paying all remote workers the same wage nationally, which creates variation in the opportunity cost of taking the remote job given the wage variation in workers' local on-site alternatives. Using the estimated demand curve, we find adverse selection likely reduces the share of call-center workers working remotely from 17% to 6% nationally, leading to losses of $824 million

[7]This design 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).

[8]Battiston et al. (2017) find that emergency phone operators communicate more efficiently when together physically. Gibbs et al. (2021) find programmers worked more hours but were no better at meeting managers' targets when remote during the pandemic. Similarly, Künn et al. (2020) find that chess players made more errors when competing online due to the pandemic.

annually just among the 3.2 million American call-center workers. The losses may be even more acute in other remotable jobs, where career concerns loom larger and it's harder to monitor remote workers from afar.[9]

Our analysis suggests that the pandemic will attenuate but not eliminate adverse selection into remote work. On the worker side, surveys suggest that the retailer's workers have learned more about their tastes during the lockdown, causing more high-ability workers to choose remote jobs. This reduction in the average cost of remote work would increase its prevalence by 1.1pp. On the firm side, we find little evidence that the retailer's experience with remote work reduced the promotion penalty and the consequent incentive for workers to sort on ability.[10]

The rest of the paper proceeds as follows. Section II introduces our data and context. Section III offers descriptive evidence on the career ladders of remote and on-site workers that motivates a model where career concerns shape the market for remote work. Section IV uses natural experiments at the retailer to test the model and estimate the costs of remote work. Section V quantifies the distortion caused by adverse selection. Section VI concludes.

II Data

Our data come from the call-centers of a Fortune 500 online retailer between 2018 and 2020. During this time, we observe 4,440 call-center workers, of whom 87% were recruited into on-site jobs and 13% were recruited into remote ones. Table 1 provides summary statistics on call-center workers hired into entry-level jobs be-

[9] Our analysis builds on the literature that investigates the impact of selection on the provision of workplace amenities. Tô (2018) finds evidence that taking parental leave is a negative signal about a worker's subsequent productivity. Adverse selection has been stressed as a motivation for government mandated benefits broadly (Summers, 1989) and workers' compensation insurance (Gruber and Krueger, 1991) and maternity leave specifically (Gruber, 1994; Ruhm, 1998).

[10] Managerial strategies such as virtual watercoolers might change remote workers' promotion prospects (Bojinov et al., 2021). However, low take-up of such strategies is consistent with the historically slow responses of managerial practices to changes in technology (Juhász et al., 2020).

fore the retailer closed its on-site call-centers due to COVID-19.[11] The final four columns consider workers hired after July 2018 when the retailer began hiring workers directly into remote jobs.

After new hires finish three weeks of training, they handle low-stakes, incoming calls that might ask when an order will arrive or how to return a product. Entry-level workers receive calls randomly routed from the same pool, regardless of whether they work on-site or remotely.[12] After six months, some workers are promoted to specialized roles, handling high-value clients (e.g. businesses), high-value products (e.g. mattresses), or high-value transactions (e.g. refunds for damaged products). These high-stakes calls are not randomly assigned. To ensure fair comparisons, we focus on workers' first six months when comparing productivity.

During their first six months, workers hired after July of 2018 average 3.1 calls per hour in row 1 of Table 1 (standard deviation = 4.4). Each call averages 9.5 minutes, leaving workers with a half hour to do paperwork, take breaks, and wait for calls. Remote workers answer slightly fewer calls per hour than on-site workers, but this gap understates the productivity difference on any given day since remote workers were hired later when the retailer fielded more calls. Controlling for time, remote workers answer 0.33 or 10.7% fewer calls per hour (Appendix A.III).[13]

Workers can handle most calls by themselves, but occasionally ask managers for help in online chats or in-person conversations. At the retailer, remote teams consist entirely of remote workers overseen by remote managers.

In addition to tallying calls, the retailer tracks two proxies of call quality. One met-

[11]Table A.8 provides summary statistics on the subsamples used for our five core analyses.

[12]Calls were randomly routed to workers logged into the retailer's software at the same time. Shifts were determined by time-zone, so our productivity analyses control for time-zone.

[13]The gap in productivity persists as workers gain experience at the retailer. There are similar returns to experience for both remote and on-site workers as illustrated in Figure A.4, providing suggestive evidence that remote work did not preclude on-the-job learning at least for routine calls.

rics tracks customers' responses to surveys where they can rate their satisfaction from one to five stars (row 2 of Table 1). The other metric assumes questions were successfully answered if customers do not call back within two days (row 3).

Despite these metrics, the firm is uncertain whether workers did a good job. A quick call might be efficient, curt, or just lucky. A dissatisfied customer might leave a 5-star review to be polite (the mean review is 4.9 out of 5), while a satisfied customer fails to leave a review (the participation rate is 11%).[14] Similarly, some customers may call back because new questions crop up, while others do not call back because they lack confidence in the customer service agents.

The challenges in measuring call quality limit the retailer's use of performance pay. Since quantity is measured accurately while quality is measured noisily, performance pay can cause workers to optimize for speedier calls rather than more satisfied customers (Holmstrom and Milgrom, 1991). Given this unintended consequence, performance pay accounts for at most 17% of earnings at the retailer.[15]

The challenges of measuring call quality also cause managers to be uncertain about which workers to promote. Among workers hired after July of 2018, 9.8% were promoted to higher-stakes customer-service roles at some point in their time with the retailer (row 4 of Table 1). These promotions increase workers' wages by 13% and (b) give workers priority for avoiding night and weekend shifts. On-site workers were more than twice as likely to be promoted than remote workers — 10.9% versus 5.0% in row 4 of Table 1. Figure 1 illustrates these divergent career ladders. The promotion gap between remote and on-site workers persisted even as the retailer gained organizational experience with remote work (see Figure A.5).[16] This

[14]Each call is recorded so managers can do quality-assurance checks. However, managers are judged on their workers' performance, incentivizing them to give uniformly positive evaluations.

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

[16]Because of the changing nature of the retailer's business, the rate of promotion declined. Thus,

pattern is consistent with managers having more uncertainty about the ability of remote workers, which depresses promotion rates.[17]

In addition to promotions within the firm, workers value their managers' references for other jobs since about half of the workers leave the firm (row 6). Workers are rarely fired for poor performance at the retailer (row 7).

Consistent with their lower productivity, those hired into remote jobs were paid $1.14/hour less at the time of hire (row 8). The initial wages of on-site workers ranged from $14/hour to $16/hour and were set to reflect the pay of customer-service (CSR) jobs in the metropolitan statistical area (MSA) (row 9). By contrast, all remote hires were paid $14/hour regardless of where they lived.

The majority of the retailer's call-center workers were female, especially among their remote workforce — 88% of remote workers are female compared to 69% of on-site workers (row 11 of Table 1). Workers often come to this job with little prior experience, consistent with the average age of 32 (row 12). 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.

III Information Frictions in Remote Work

III.A Managers' Information About Remote and On-Site Workers

We investigate whether it is harder to assess remote workers' abilities.

the gap in promotions fell in levels but not in percentage terms. The proportion of promotions seems more broadly applicable to other firms with different business conditions.

[17]More uncertainty could also reduce firing rates. When the retailer fired more workers in 2018, workers who transitioned from on-site to remote work in their first month at the retailer were 8.7pp less likely to be fired (95% CI = [-13.8pp,-3.5pp]) than their peers hired at the same time and place. These workers were also 17.7pp less likely to be promoted (95% CI = [-36.0pp, 0.66pp]).

In December of 2018, the retailer's managers were asked to rate each worker's performance on a scale from one (unsatisfactory) to five (exceptional). These evaluations were an input into promotions to higher-stakes customer service roles.

When evaluating workers, managers could consider their impressions from listening to workers' calls with customers or conversations with coworkers. These impressions could come from passive observation or purposeful oversight.

If such impressions give managers' additional information about workers' productivity beyond workers' metrics, then managers should give better-than-expected evaluations to workers whose metrics understate their productivity. If so, highly-rated workers should have better metrics in subsequent months, conditional on their past metrics. If on-site managers have more information, their evaluations should better predict on-site workers' subsequent metrics.

We estimate how managers' evaluations in December of 2018 predicts subsequent customer satisfaction ratings. Letting i index the worker; t, the date; z, the time-zone; and j, the worker's job title, we estimate:

$$\begin{aligned}\text{Cust. Sat. after Dec 2018}_{i,t} = {} & \psi_1 \text{Positive Manager Evaluation in Dec 2018}_i \\ & + \psi_2 \text{Avg. Cust. Sat. July-Dec 2018}_i + \mu_{t,z(i),j(i)} + \epsilon_{i,t,z}. \quad (1)\end{aligned}$$

Column 1 of Table 2 estimates equation 1 for all workers. Managers appear to have additional information beyond the recorded metrics. However, managers' evaluations are only predictive for on-site workers. When an on-site manager gives an on-site worker a positive rating (of 4 or 5), the worker's subsequent customer satisfaction ratings tend to exceed predicted performance by 0.25 standard deviations (95% CI = [0.14,0.36] in column 2 of Table 2). By contrast, remote managers' evaluations are not significantly predictive of remote workers' subsequent performance. Together, on-site managers' evaluations are significantly more predictive (with a difference of 0.27 standard deviations, 95% CI = [0.043, 0.50]).[18] These patterns also apply for negative evaluations but are weaker and less precisely estimated, given the relative rarity of negative evaluations (see Table A.9).

This evidence suggests that managers can be less certain of their evaluations of remote workers and, thus, less confident about promoting them.[19]

III.B Model

This section models the market for remote work and shows how an informational friction can lead to an under-provision of remote work.[20]

In our two-period model, workers choose between remote and on-site jobs. Each job features two possible tasks — one low-skill and one high-skill. Workers vary in their tastes for remote work and their abilities. Firms post menus of jobs. All firms have the same, additive production function and operate in competitive markets.[21]

In period zero, each firm posts a menu of one-period contracts.[22] Each worker chooses a contract after privately learning the probability that she will be high-ability. During the first period of work, firms learn some workers are high-ability and some are low-ability, while remaining uncertain about others. Those revealed to be high-ability are promoted, while those revealed to be low-ability are de-

[18]The same pattern holds without controlling for workers' past customer satisfaction (difference = 0.30 standard deviations, 95% CI = [0.033,0.56]).

[19]There is limited evidence that the gap between remote and on-site workers narrowed as the retailer gained experience with remote work. When managers evaluated call-center workers in the summer of 2019, there was no significant difference in the predictive power of these evaluations (as reported in Table A.10). One caveat, however, is that the summer evaluation cycle featured a three-point scale rather than a five-point scale, reducing the signal in the evaluations and leading to significantly different evaluations for remote and on-site workers (see Figure A.6).

[20]More details of the derivations can be found in Appendix A.I.

[21]Our stylized model features two-periods and two rungs of the career ladder. This is a good approximation of our empirical context. Further, the insights are qualitatively similar for an infinite period problem with a continuous choice of what share of time to spend working remotely.

[22]We assume that firms cannot sort workers by varying the bonus for high ability. This constraint could reflect a lack of credible commitment, workers' fairness concerns, or workers' risk aversion.

moted. Firms are more likely to learn about on-site workers than remote ones.

III.B.1 The Firm's Problem

Each firm's production function is as follows. In the low-skill task ($T = L$), a low-ability worker ($\Theta_i = L$) produces y, while a high-ability worker ($\Theta_i = H$) produces $y + a$ where $a > 0$. When assigned the high-skill task, a high-ability worker's output increases by A and a low-ability worker's output decreases by C. Working remotely changes output by τ, the treatment effect of remote work. The per-period output Y of worker i in job $j \in \{\mathrm{r} \equiv \text{remote}, \mathrm{o} \equiv \text{on-site}\}$ and task T is:

$$Y_{ijT} = y + a \cdot \mathbb{1}[\Theta_i = \mathrm{H}] + \begin{cases} -C & \Theta_i = L,\ T = H \\ A & \Theta_i = H,\ T = H \end{cases} + \tau \cdot \mathbb{1}[j = \text{remote}], \qquad (2)$$

where C is assumed to be sufficiently high that the firm only assigns workers the high-skill task when they are known to be high-ability.

Initially, firms do not know individual workers' abilities and can only infer likely ability from workers' choices to be remote or on-site. Once workers' ability is revealed, workers are paid their marginal product since signals are public and markets are competitive. The average cost of hiring a remote worker instead of an on-site one equals the difference in average products in the first period:

$$\mathrm{AC} = \mathbb{E}_{\mathrm{o}}[Y_{ioL}] - \mathbb{E}_{\mathrm{r}}[Y_{irL}] = -\tau + a(\Pr(\Theta_i = H \mid o) - \Pr(\Theta_i = H \mid r)) \qquad (3)$$

The first term reflects the treatment effect of remote work; the second term reflects the self-selection of high-ability workers into on-site jobs.

III.B.2 The Worker's Problem

Workers vary in their abilities and tastes. Worker i's ability is either high or low, $\Theta_i \in \{H, L\}$. When choosing her first job, she privately knows her probability, $\theta_i \sim \text{Uniform}[0,1]$, of being high-ability. Each worker has an idiosyncratic taste for remote work, $\nu_i = \bar{\nu} + \jmath\epsilon_i$ where $\epsilon_i \sim \mathscr{L}(0,1)$ is logistic and orthogonal to ability.[23]

We assume that workers make fixed cost investments in their work arrangement that make switching prohibitively costly in the second period.[24]

Workers choose their job to maximize:

$$U(\theta_i, \nu_i) = \max_{j \in \{\mathrm{r,o}\}} \begin{cases} w_{\mathrm{r}} + (1+\delta)\nu_i + \delta\mathbb{E}[w \mid \theta_i, r] & \text{if remote} \\ w_{\mathrm{o}} + \delta\mathbb{E}[w \mid \theta_i, o] & \text{if on-site} \end{cases}, \tag{4}$$

yielding a threshold rule for choosing remote work of:

$$w_o - w_r \leq \nu_i(1+\delta) + \delta(\mathbb{E}[w \mid \theta_i, r] - \mathbb{E}[w \mid \theta_i, o]). \tag{5}$$

The worker weighs the first-period change in income against her tastes and her likely second-period income, which is discounted according to δ.[25]

When predicting her future income, the worker considers two possibilities. One, with probability, p_j, her ability is revealed and she earns her marginal product. This is more likely in on-site jobs than remote ones ($p_o > p_r$). Two, with probability,

[23] This might reflect, for example, the length of the worker's potential commute or her child-care responsibilities. Indeed, at the retailer, when a remote work program was introduced, those who went remote were 14.9pp more likely to live at least 15 miles from the office than those who remained on-site (se = 0.029).

[24] Workers might buy a car to commute or build a home office for working remotely.

[25] In reality, the gains from promotion may also include social validation and better amenities.

$1 - p_j$, her ability remains unknown and her wage remains constant, so:[26]

$$\mathbb{E}[w \mid \theta_i, j] = w_j + p_j(\mathbb{E}[\text{MP}_j \mid \theta_i] - w_j). \tag{6}$$

A worker who privately knows she is likely to be high-ability (high θ_i) expects her marginal product to exceed the pooled wage ($\mathbb{E}[\text{MP}_j \mid \theta_i] > w_j$). Thus, for her, working remotely is costly because it obscures her ability. By contrast, a worker who privately knows he is likely to be low-ability (low θ_i) expects his marginal product to fall short of the pooled wage ($\mathbb{E}[\text{MP}_j \mid \theta_i] < w_j$). Thus, for him, working remotely hides his low-ability and allows him to pool with more productive types.

Remote work's career consequences reduce the demand for remote work among workers who know they are likely high-ability. This downward shift is the source of the selection problem: at any given wage penalty — or price of remote work — a lower share of workers who are likely high-ability choose remote work.

III.B.3 The Market Equilibrium

Figure 2 illustrates the market for remote work. The x-axis plots the share of workers who are working remotely and the y-axis plots the wage penalty — or price of remote work. In equilibrium, the price of remote work equals the average cost of hiring a remote worker instead of an on-site one in the navy line. Even when the marginal cost of switching a given worker from on-site to remote work is zero as pictured in the green line, it can still be costly for a firm to hire a remote worker instead of an on-site one. Starting from equation 3, this cost can be shown to be:

$$\text{AC} \approx -\tau + a \frac{(p_\text{o} - p_\text{r}) \frac{\delta}{1+\delta} (A + a)}{\mathscr{s}} \text{Var}(\theta_i), \tag{7}$$

[26] The probability p_j is a feature of the job and not of the worker. Thus, nothing can be inferred about ability if it is not fully revealed.

where the first term is the marginal cost of remote work and the second term is the selection effect of remote work.

Workers' self-selection into jobs based on their private information about their ability drives a wedge between the marginal and average costs of remote work. The wedge is larger when there are greater returns to high-ability in the low-skill ability (a) and when more workers self-select into jobs based on their latent ability. Workers self-select more on ability when they have more private information about ability, $\text{Var}(\theta_i)$, and when remote work is more determinative of their second-period income. Remote work affects second period income more when (i) there is a greater gap in the probability that ability is revealed in the two jobs, $p_{\text{o}} - p_{\text{r}}$, and (ii) there is a greater discounted return to being observably high- rather than low-ability, $\frac{\delta}{1+\delta}(A + a)$. Workers self-select less on ability when there is more taste variation, $\mathscr{s}$, which can cause latently high-ability workers to go remote and latently low-ability workers to go on-site.

Since the average cost determines the equilibrium price of remote work, the market quantity, q_{mkt}, is found at its intersection with the demand curve in Figure 2.

There are two sources of inefficiencies in the market. First, firms price at the average rather than the marginal cost of remote work, leading to deadweight losses in the red Harberger triangle in Figure 2. Second, workers' demand for remote work also deviates from the marginal social benefit because the revelation of ability changes the attribution of credit as well as the assignment of tasks. These private gains lead to excessive sorting by ability and depress the demand for remote work around the equilibrium quantity. Thus, the Harberger triangle is a conservative estimate of the deadweight losses from asymmetric information.

This inefficient equilibrium, however, is not set in stone. Instead, it is a function of the technologies for evaluating remote workers, which determine $p_{\text{o}} - p_{\text{r}}$, and the distribution of tastes for remote work, which determines $\mathcal{s}$.

If firms become better able to evaluate remote workers, then the average cost of remote work will fall towards the marginal cost $\left(\frac{\partial \mathrm{AC}}{\partial (p_{\mathrm{r}} - p_{\mathrm{o}})} < 0\right)$. If firms have learned how to better assess the productivity of remote workers during the pandemic, COVID-19 could lead to a more efficient equilibrium.

If tastes become more variable, the average cost of remote work falls towards the marginal cost $\left(\frac{\partial \mathrm{AC}}{\partial \mathcal{s}} < 0\right)$. During COVID-19, tastes may have become more variable as many workers experienced full-time remote work for the first time. By forcing all workers to learn about their tastes, COVID-19 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.[27]

In the model, greater informational frictions in remote work make remote work (i) unattractive for latently high-ability workers who want their ability revealed and (ii) attractive for latently low-ability workers who want their ability hidden. Thus, the model's central empirical prediction is that remote workers will be adversely selected. Adverse selection leads to the model's central welfare implication that remote work will be under-provided.

IV Estimating the Costs of Remote Work

IV.A Identification Strategies: Intuition

The average cost of remote work depends 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

[27]COVID-19 may have also made remote work more attractive if workers bore fixed costs of setting up home offices or learning new technologies. These changes would increase both the efficient and market quantity of remote work so would not eliminate the market failure.

on private information about their ability. To build intuition for how we identify these two effects empirically, we consider the ideal experiment that would first randomly assign workers' offers at the time of hire and then randomly assign their actual jobs.[28] The resulting treatment cells are summarized in Schematic 1.

Schematic 1: Intuition for Identification Strategy

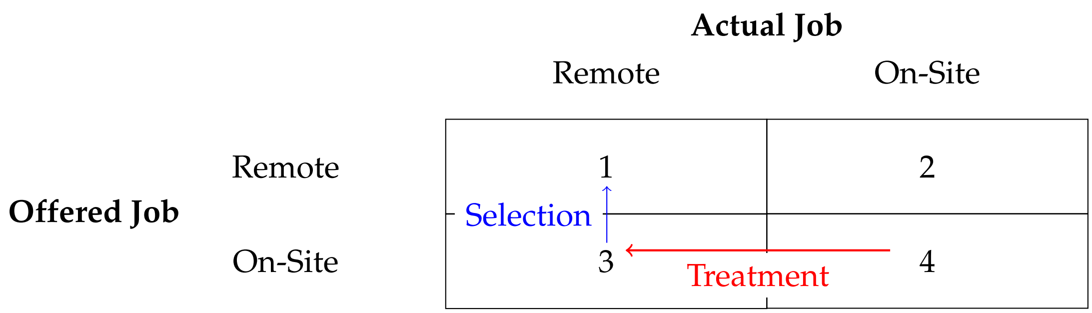

Note: This figure illustrates the ideal thought experiment. The rows reflect the type of job randomly offered at recruitment. The columns reflect the type of job actually assigned.

To identify selection into remote jobs, we would focus on the sample of workers who were working remotely (the first column of Schematic 1): within this sample, we would compare the productivity of those who had initially accepted a remote job (cell 1) to the productivity of those who had initially accepted an on-site one (cell 3). Comparing workers who were all working remotely but had been offered different initial positions would isolate the causal effect of offering remote work on the selection of workers who accept the offer.

To identify the treatment effect of remote work, we would focus on the sample of workers who had been offered on-site jobs (the bottom row of Schematic 1): within this sample, we would compare the productivity of those who were working remotely to the productivity of those who were working on-site, while holding worker selection constant (cell 3 versus 4).

[28]Such an ideal design is similar to Karlan and Zinman (2009)'s experiment in the credit market.

We leverage two quasi-experiments 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 offers of new hires. Among workers who ultimately worked remotely, we compare 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 in Schematic 1).

This program also allows us to identify the treatment effect of remote work, by comparing worker productivity in remote and on-site work among those who took up the opportunity to go remote. 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, COVID-19 caused all of the retailer's workers to work remotely. Since July 2018, the retailer had hired workers directly into remote jobs. Thus, during the lockdown, we can compare the productivity of those who were initially offered remote jobs to the productivity of those who were initially offered on-site jobs to identify the selection effect of remote work (approximating cell 1 versus 3).

When the call-centers closed, workers who were initially offered on-site jobs transitioned from on-site to remote work. By comparing the change in their productivity to that of workers who were already remote, we can isolate the treatment effect of remote work (cell 3 versus 4) net of the common shocks of COVID-19.

IV.B Opportunities to Go Remote in 2018

In the beginning of 2018, the retailer started running out of desks in some call-centers so created remote job openings for existing on-site workers. Workers who opted into remote work continued to handle the same calls and earn the same wages.

IV.B.1 Estimating the Selection Effect

For remote workers hired before January of 2018, these remote possibilities were unknown when they were hired (akin to cell 1 in Figure 1). By contrast, for remote workers hired in the first few months of 2018, these remote possibilities were known and could have influenced their decision to accept the initial offer (akin to cell 3).[29] Thus, we can compare the productivity of remote workers who were effectively offered different jobs. Other than this difference in the offer, these workers had similar experiences. They were all paid the same wages on-site and remote, were all drawn from the same labor markets, and were all trained on-site.

To adjust for changes in consumer demand and worker experience, we use workers who chose to remain on-site as a control group.[30] For these workers, the new menu merely introduced an option they did not choose.[31] Our difference-in-difference design compares the gap in productivity between remote and on-site workers in later cohorts to that in earlier cohorts:

$$\text{Calls/Hour}_{i,t} = \beta_{\text{Selection},1} \cdot r(i) \cdot \mathbb{1}[h(i) \geq \text{Jan, 2018}] + \beta_r r(i) + \mu_{t,h(i),\ell(i)} + \epsilon_{i,t}. \quad (8)$$

where the fixed effects $\mu_{t,h(i),\ell(i)}$ capture (a) fluctuations in consumer demand across time, t, (b) fixed differences across workers who were hired in a particular month,

[29] These opportunities were widely known because (1) the jobs were posted on external job boards for workers with external experience and (2) fully 15% of workers transitioned to remote work in the beginning of 2018.

[30] Without a control group, we would either need to compare workers who had different experience or were working at different times when the retailer had different demand.

[31] One caveat is that the retailer did not accept all applications for remote positions, which might attenuate the selection estimate because (1) some control workers i in later cohorts may been selected because of their demand for remote work and (2) workers who were allowed to go remote may have been positively selected from a negatively selected group.

$h(i)$, and place, $\ell(i)$, and (c) the returns to experience that may be place-specific.[32]

Interpreting $\beta_{\text{Selection},1}$ as the selection effect of remote work requires that the productivity of on-site and remote workers would have proceeded in parallel had there been no change in the information about opportunities to go remote. Several facts make the parallel trends assumption plausible. First, calls are randomly routed to remote and on-site workers. Second, all workers came from the same labor markets. Finally, all workers were trained on-site by the same staff and subject to the same company policies.[33] Empirical support for parallel trends can also be seen in the stable differences in worker characteristics in Table A.11.

Figure 3 illustrates this design. The x-axis 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 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 were offered jobs that they knew could go remote.

In the three cohorts hired before the remote program began, those who went remote answered 0.24 or 8.8% more calls per hour than those who remained on-site (95% CI = [-2.0%, 20.3%]), suggestive of a positive treatment effect.

By contrast, in all three cohorts hired after the remote program began, those who went remote answered fewer calls than those who remained on-site, averaging 0.33 or 12.2% fewer calls per hour (95% CI = [-18.8%, -5.7%]). This estimates the average cost of remote work since it compares those who were offered remote jobs and worked remotely to those who preferred on-site jobs and worked on-site.

The difference in these within-cohort differences isolates the selection effect of re-

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

[33] Training is a key factor given its high returns in call-centers (De Grip and Sauermann, 2012).

mote work net of the treatment effect of working remotely. As reported in column 1 of Table 3, we estimate that being offered remote work reduced the productivity of new hires by 0.57 calls per hour or 21.1% (95% CI = [-33.9%, -8.2%]).

Column 2 of Table 3 shows the robustness of this selection estimate to controlling for the hiring call-center rather than the time-zone. The estimate is also robust to controlling for the time that workers spent on-site, which could affect the returns to experience (Table A.12). Further, these results are similar when controlling for available worker characteristics of age and gender, which hiring managers could conceivably (albeit illegally) use to screen workers (Table A.13).

As seen in columns 3 through 6 of Table 3, offering remote work did not change the quality of calls in an economically or statistically significant way.

Our results indicate that workers who accepted remote offers were less productive than those who accepted on-site offers, supporting our model's prediction that remote jobs attract latently less productive workers.

IV.B.2 Estimating the Treatment Effect

When the retailer introduced its remote work program in 2018, many on-site workers opted into remote work. Workers could choose whether to go remote but not when they did so, which depended on the time it took to process their applications and find them openings on remote teams.[34] Leveraging the quasi-random timing of these transitions identifies the treatment effect of working remotely for those who opt into remote work.[35]

The treatment effect captures any moral hazard from being managed from afar as

[34] Applications could be rejected, especially after the beginning of 2018 when remote openings were limited. As a result, those who were approved to go remote at any point were not always adversely selected, leading to higher promotion rates in Table A.8.

[35] This parallels the population in Bloom et al. (2015).

well as any productivity advantage from, for example, reduced distraction from coworkers.[36] Which effect dominates will determine whether productivity decreases or increases around individual transitions to remote work.

Let $e_{i,t}$ denote the event time at time t, or the number of weeks from worker i's switch to remote work. We estimate the productivity change around workers' transitions to remote work within one- and six-week bandwidths, according to:

$$\text{Calls/Hour}_{i,t} = \tau_1 \mathbb{1}[\mathrm{e}_{i,t} \geq 0] + \epsilon_{i,t} \quad \text{if } |e_{i,t}| \leq \text{Bandwidth}. \tag{9}$$

To hold the set of workers constant, we limit to a balanced panel of workers.[37] To hold the type of calls constant, we limit to workers who were in consistent roles.

Figure 4 considers the hourly calls of the 127 entry-level workers who transitioned from on-site to remote work while remaining in the same role. 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 their last six weeks on-site, workers' hourly calls were steady. In the first week of remote work, productivity increased by 0.21 calls per hour or 7.1%. During the subsequent five weeks, productivity hovered around this higher level, with all the confidence intervals lying above the reference level. Pooling across the six weeks before and after the transition to remote work, calls per hour increased by 0.23 or 7.6% (95% CI = [3.3%, 11.9%]), suggesting remote work allowed the same workers handling the same calls to answer more of them.

The estimate is robust to different choices of the bandwidth, since the positive ef-

[36]The treatment effect also captures any differences in managerial selection since workers transitioned from working under on-site managers to working under remote ones.

[37]Of the 336 entry-level workers who transitioned to remote work in 2018, 53 were not at the retailer for the full twelve weeks, 49 were initially in training, and 107 changed job roles.

fects do not fade with time (see Figure A.7). The results are also robust to accounting for the possibility that switchers would have become more productive had they remained on-site in a difference-in-difference design presented in Appendix A.II. Further, the changes in calls per hour were similar across workers who were offered remote and on-site jobs at the time of hire, suggesting the positive effects did not hinge on how workers selected into the firm.[38]

Workers answered more calls after going remote because they answered them 3.0% faster and spent 4.5% more time on the phone (see Table A.14). Remote work may have reduced distractions from coworkers, helped workers better time breaks around calls, or made it easier to juggle work with other responsibilities.[39]

The event studies reveal no statistically nor economically significant changes in customer satisfaction ratings nor in the the share of customers who called back within two days (columns 3−6 of Table 4). When drawing comparisons to workers who remained on-site in Appendix A.II, remote work appears to cause a small decrease in the share of calls with no callback. This decrease, however, is concentrated in the first two weeks of remote work when workers adjusted to working more independently and asking for help electronically (see panel b of Figure A.2).

This design suggests that call-center work is conducive to remote work. However, remote work attracts less productive workers to the firm, an effect that more than outweighs the positive treatment effect in our setting. Thus, on average, hiring remote workers was costly to the firm, even though, on the margin, it was more productive for a given worker to work remotely.

[38]Remote work increased the hourly calls of those initially offered remote opportunities by 0.24 (se=0.0803; N=77) and of those initially offered a purely on-site job by 0.21 (se=0.096, N=50). These results suggest similar treatment effects but cannot rule out large differences. We return to this question in the context of COVID-19's lockdown in Section IV.C.2.

[39]Such juggling is also consistent with reduced absent time (columns 4 and 5 of Table A.14).

IV.C COVID-19 Lockdown: All Workers Work Remotely

On April 6, 2020, COVID-19 induced the retailer to close its call-centers, causing formerly on-site workers to start working remotely. At the same time, many of the retailer's workers were already working remotely, since the retailer had hired workers directly into remote jobs since July of 2018.[40]

IV.C.1 Estimating the Selection Effect

During COVID-19's lockdown, workers who had been offered on-site jobs worked remotely alongside those who had been offered remote jobs. Thus, during the lockdown, comparing the productivity of these workers captures selection into remote work but not the treatment effect of working remotely, as in:

$$\text{Calls/Hour}_{i,t} = \beta_{\text{Selection},2}\mathbb{1}[\text{Offered Remote Work}_i] + X'_{i,t}\alpha + \mu_{t,h(i),\ell(i)} + \epsilon_{i,t}. \quad (10)$$

We limit the analysis to on-site workers with the same base wage as remote workers to hold pay constant. For $\beta_{\text{Selection},2}$ to identify the selection effect of remote work, (1) remote and on-site workers must be drawn from conditionally similar labor markets and (2) remote and on-site workers must be similarly exposed to the productivity consequences of the pandemic.

In column 1 of Table 5, those who had been initially offered remote jobs answered 0.62 or 18.0% fewer calls per hour than those who had been initially offered on-site jobs during the lockdown (95% CI = [9.4%, 26.7%]). Remote workers did not compensate for lower quantity with significantly higher satisfaction (Table A.16) or lower callback rates (Table A.17).

The remaining columns of Table 5 consider the robustness of the results to the

[40] Appendix A.III uses this policy to estimate the average cost of remote work. The move towards hiring remote workers directly into remote jobs can be seen in Figure A.8.

inclusion of controls. Column 2 controls for local pandemic conditions; columns 3-4, for pre-pandemic, local labor market conditions; and column 5, for worker demographics.[41] The estimates of adverse selection range from 14.9% to 17.4%. Further, adjusting for the greater attrition of low-productivity, on-site workers only changes the estimate by 0.2% (see Appendix A.IV).

Within the sample who completed the retailer's June 2020 caregiving survey, including the control for caregiving has a barely detectable effect on the estimated selection effect (see Table A.15), since those with childcare responsibilities were equally as productive as other workers both before and during the lockdown.

The adverse selection effects identified from the shock of COVID-19 are similar to the estimated effects of 14.8% to 21.1% identified from the introduction of the 2018 remote work program in Section IV.B.1. These designs offer independent empirical support for our model's prediction of adverse selection into remote work.

IV.C.2 Estimating the Treatment Effect

The closures of the call-centers also shed light on the treatment effect of remote work. When the offices closed, the productivity changes of formerly on-site workers capture the treatment effect of remote work (τ_o) and the other spurious shocks of the pandemic (σ_o), so $\Delta \frac{\text{Calls}}{\text{Hour}}_o = \tau_o + \sigma_o$.[42] By contrast, already remote workers saw no change in their work arrangements, so $\Delta \frac{\text{Calls}}{\text{Hour}}_r = \sigma_r$. The difference-in-

[41] COVID-19 data are from The New York Times: `https://github.com/nytimes/covid-19-data`.

[42] The shock of COVID-19 identifies remote work's treatment effect for on-site workers who persist in remote jobs, which is the relevant population for firms considering whether to transition on-site workforces into remote work. To assess how our treatment effect might extrapolate to all on-site workers, we compare the effect for those who attrit soon after the office closures to that for the full sample. The similarity of the estimates (12.2% vs 10.2%) suggests that remote work did not have worse effects among those who exited.

differences is, $\Delta \frac{\text{Calls}}{\text{Hour}}_o - \Delta \frac{\text{Calls}}{\text{Hour}}_r = \tau_o + (\sigma_r - \sigma_o)$, which we estimate using:

$$\begin{aligned} \text{Calls/Hour}_{i,t} = \tau_{\text{DiD}} \cdot \mathbb{1}[\text{On-Site Offer}_i] \cdot \mathbb{1}[t \geq \text{April 6, 2020}] \\ + \psi \cdot \mathbb{1}[\text{On-Site Offer}_i] + X'_{i,t}\beta + \mu_{t,h(i),\ell(i)} + \epsilon_{i,t} \end{aligned} \tag{11}$$

where i denotes the worker; t, the date; $h(i)$, the month of hire; and $\ell(i)$, the time-zone. The identifying assumption is that remote and on-site workers experienced similar spurious shocks conditional on observables.

Figure 5 illustrates our difference-in-difference design. The x-axis represents the week of work. The y-axis represents calls taken per hour. The vertical line indicates the office closures on April 6, 2020. Around the closures, the productivity of already remote workers (in the dashed line) increased as the retailer saw a surge in consumer demand. However, the productivity of formerly on-site workers (in the solid line) increased by even more, rising by an additional 0.38 calls per hour or 10.2% as reported in column 1 of Table 6. Introducing worker fixed effects to address selective attrition in column 2 does not appreciably change the estimate.

The increased productivity gap between remote and on-site workers does not appear to be an artifact of the surge in consumer demand, since earlier surges in demand had parallel effects on remote and on-site workers.[43]

Now, consider the bandwidth of the estimator. The retailer closed its on-site call-centers three weeks after most American schools closed. Focusing on a three week bandwidth as in column 3 holds school policies constant. On the other hand, the onset of the pandemic may have made the office draining.[44] Indeed, in Figure 5,

[43]For example, during the 2020 holiday season, call volumes changed by more than one call per hour, but the difference in productivity between remote and on-site workers stayed flat. More generally, increases in hourly calls at the retailer were not significantly associated with the gap between on-site and remote workers. Using this cross-sectional relationship, an increase in hourly calls of 1.2 like that seen during the pandemic would only tend to lead to a 1.54% increase in the hourly calls of on-site workers relative to remote ones (95% CI = [-1.29%, 4.37%]).

[44]On-site workers may have feared the virus, taken time-consuming measures to avoid its

on-site workers became less productive relative to those already working remotely right before the closures.[45] In column 4, our donut design excludes the three weeks prior to the office closures, yielding an estimated treatment effect of remote work of 0.25 or 6.6% calls per hour (95% CI = [-0.62%, 13.8%]). This is our preferred estimate since it compares remote work to a more typical baseline in the office.[46]

Controls for local COVID-19 deaths and cases in column 5 do not appreciably change the estimate. Column 6 controls for workers' childcare responsibilities as reported in a June 2020 survey by the retailer, which yields a similar point estimate but wider confidence intervals given the smaller sample of survey participants.[47]

The increase in calls came from workers spending more time on the phone rather than handling calls more quickly as illustrated in Figure A.10. Thus, it is important to consider spillovers across workers. In the call-center context, when one worker answers a call, another worker cannot receive it. Since calls are randomly routed between on-site and remote workers, both treatment and control workers experience resulting delays, which net out to a first-order (Butts, 2021). Given the high volume of calls during the pandemic, new calls arrived nearly as fast as they could be answered, mitigating the impact of higher order terms.

Our analysis indicates that there is a positive causal effect of remote work on calls per hour with no sacrifice in call quality (as seen in Tables A.19 and A.20). This complements our findings in Section IV.B.2, which found sharp increases in worker productivity around individuals' voluntary transitions from on-site to re-

spread, and spent time preparing their home office for remote work.

[45] Between March 16, 2020 and the office closures, on-site workers took 5.3% more calls per hour than remote workers but, in the previous month, had taken 11.5% more calls per hour, a significantly greater gap in calls (p-value = 0.065).

[46] Considering different post-periods results in similar estimates as illustrated in Figure A.9.

[47] Similarly controlling for gender and age do not appreciably affect the estimated treatment effect in Table A.18. Women's productivity increased by less than men's when they went remote. Similarly, caregivers' productivity increased by less than non-caregivers'. However, these rely on a limited number of men and non-caregivers.

mote work. Together, these analyses suggest that working remotely causes workers to become 6.6% to 7.6% productive in this setting, both for workers who choose to work remotely and those who do not.

V WELFARE EFFECTS OF ADVERSE SELECTION

This section provides back-of-the-envelope estimates of the under-provision of remote work and the lost surplus due to adverse selection.

Let's return to the model pictured in Figure 2. In our setting, we find that working remotely increases productivity by 7% to 8%, so the green marginal cost line lies *below* zero. However, the workers who accept remote offers tend to be 15% to 21% less productive than those who accept on-site offers, driving a wedge between the average and marginal costs of at least \$2.25/hour. Since adverse selection more than outweighs the positive treatment effect, firms find it costly to hire remote workers instead of on-site ones in our setting. Thus, the navy average cost line lies *above* zero.[48] To offset this average cost, remote workers pay a price to work remotely in the form of 8% lower wages at firms like the retailer.

By driving a wedge between the marginal and average costs of remote work, adverse selection reduces the market provision of remote work from q_{eff} to q_{mkt}. The workers between these quantities choose on-site jobs because of the incentive to pool with more productive types. To estimate the magnitude of this quantity distortion, we leverage the retailer's policy of paying a uniform remote wage nationally, which contrasts with the variation in workers' local on-site wages. Workers in higher-wage labor markets face a higher opportunity cost of accepting this remote job and, as a result, must pay a higher implicit price to work remotely. Using the

[48]Based on our natural experiments, remote workers would be predicted to be 7% to 14% less productive, using the estimated selection effects of -15% or -21% and estimated treatment effects of $+7\%$ or $+8\%$. Consistent with these predictions, remote workers were 8% less productive than comparable on-site workers prior to the pandemic, as seen in Appendix A.III.

variation in the retailer's recruitment of remote workers in lower- and higher-wage labor markets, we find that the \$2.25/hour wedge created by adverse selection reduces the share of call-center workers working remotely from 17.1% (q_{eff}) to 6.2% (q_{mkt}), which leads to 350 thousand fewer remote call-center workers nationally (see Appendix A.V for details).

So, how does the distortion in workers' choices between q_{mkt} and q_{eff} affect welfare? First, consider workers on the margin of remote work at the market equilibrium, q_{mkt}. These workers choose on-site jobs solely to pool with more productive workers: both their preferences and productivity instead push towards remote work. Thus, their welfare losses reflect the adverse-selection wedge of \$2.25/hour. By contrast, workers on the margin of remote work at the efficient equilibrium, q_{eff}, prefer on-site work enough to offset remote work's productivity advantage. Thus, there are no welfare losses for these workers. Averaging between these extremes, adverse selection reduces social surplus by \$1.13/hour on average for workers between q_{mkt} and q_{eff}. For the 3.2 million call-center workers in the US, adverse selection reduces social surplus from remote work by \$824 million annually in the call-center sector alone.[49]

The selection problem in remote work is a function of tastes and technologies that could have changed during the pandemic. Consistent with Bayesian updating, we find that workers' tastes for remote work have become 12% more variable during the lockdown using a survey of the retailer's workers in April 2021 (described in Appendix A.VI). Such an increase in taste variation would reduce adverse selection into remote work as more workers would choose remote jobs because of

[49] As emphasized in the model, the Harberger triangle is an underestimate of the losses from asymmetric information, which also causes the demand curve to deviate from the marginal social benefit of remote work. Further, adverse selection into remote work also depresses the wages of inframarginal remote workers, who choose remote work (1) because of latently low-ability and (2) because of strong tastes, such as due to caregiving responsibilities. Transferring resources to remote workers tagged as low-ability would be more efficient than transferring resources to them through perturbations of the tax system (Mirrlees, 1976).

strong tastes rather than latently low-ability. As a result, the share of call-center workers working remotely would rise by 1.1pp, thereby marginally reducing the under-provision of remote work. The pandemic also gave firms experience evaluating remote workers. The persistent promotion gap between remote and on-site workers in our setting, however, suggests that experience alone does not close the opportunity gap (see Figure A.5). Taken together, our analysis suggests that the COVID-19 lockdowns will attenuate but not eliminate adverse selection into remote work, a problem which will continue to shape the provision of remote work.

VI Conclusion

We consider why so few Americans worked remotely prior to COVID-19. In our call-center context, the rarity of remote work seemed particularly puzzling since (1) workers expressed strong tastes for remote work (Mas and Pallais, 2017) and (2) working remotely seemed to make workers more productive (Bloom et al., 2015).

We argue that the missing piece to this puzzle is adverse selection into remote work. In our context, remote workers were half as likely to be promoted as on-site workers, consistent Bloom et al. (2015)'s RCT evidence. Our model of career concerns suggests that workers will sort into remote and on-site jobs on the basis of private information about their ability. Workers who want to hide low-ability will tend to choose remote jobs, while workers who want to reveal high-ability will tend to choose on-site ones, leading to adverse selection into remote work.

The theoretical prediction of adverse selection is born out empirically. Those offered remote jobs were at least 15% less productive than those offered on-site jobs, when all workers were working remotely. 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 than on-site workers, giving marginal workers an incentive to opt into on-site jobs simply to pool with more productive types. Thus, even in a job well-suited to remote work, a small minority of call-centers workers were remote before the pandemic. Our analysis suggests that adverse selection depresses the provision of remote work from 17% to 6% among call-center workers nationally. While the pandemic may increase the prevalence of remote work by attenuating the selection problem, our analysis suggests that adverse selection will continue to make remote work rarer than it would optimally be.

Our paper cannot speak to the important question how remote work affects work in teams, where others have found less positive effects of remote work (Battiston et al., 2017; Gibbs et al., 2021). Beyond productivity, we cannot assess how one worker's decision to go remote affect how enjoyable it is for others to go to the office.[50]

We estimated the sufficient statistics of our model in the specific context of call-center work. We proposed a mechanism for adverse selection into remote work that would apply to any setting where managers have more information about on-site workers than remote ones. Understanding the role of adverse selection in other contexts would help to quantify the broader contribution of this mechanism to the rarity of remote work in the past and its prevalence in the future.

[50]Indeed, 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.

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.

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.

Bojinov, Iavor, Prithwiraj Choudhury, and Jacqueline Lane, "Virtual Watercoolers: A Field Experiment on Virtual Synchronous Interactions and Performance of Organizational Newcomers," *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, 2021, (21-125).

Borusyak, Kirill and Xavier Jaravel, "Revisiting event study designs," *Available at SSRN 2826228*, 2017.

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.

Bureau of Labor Statistics, "2013-2017 American Time Use Survey Public Use Microdata Samples," 2020. data retrieved from IPUMS ATUS, `https://www.atusdata.org/atus/` (visited November 5, 2020).

_, "The Economics Daily, One-quarter of the employed teleworked in August 2020 because of COVID-19 pandemic," 2020. data retrieved from https://www.bls.gov/opub/ted/2020/one-quarter-of-the-employed-teleworked-in-august-2020-because-of-covid-19-pandemic.htm (visited June 21, 2021).

Butts, Kyle, "Difference-in-Differences Estimation with Spatial Spillovers," 2021.

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).

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.

Gibbs, Michael, Friederike Mengel, and Christoph Siemroth, "Work from Home & Productivity: Evidence from Personnel & Analytics Data on IT Professionals," *Chicago Booth Research Paper*, 2021, (21-13).

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).

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

Künn, Steffen, Christian Seel, and Dainis Zegners, "Cognitive Performance in the Home Office-Evidence from Professional Chess," 2020.

Lazear, Edward P, "Performance pay and productivity," *American Economic Review*, 2000, *90* (5), 1346–1361.

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

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

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.

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

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.

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.

U.S. Census Bureau, "2019 American Community Survey Public Use Microdata Samples," 2021. data retrieved from IPUMS USA, `https://usa.ipums.org/usa/` (visited June 19, 2021).

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: Promotion among Remote and On-Site Workers

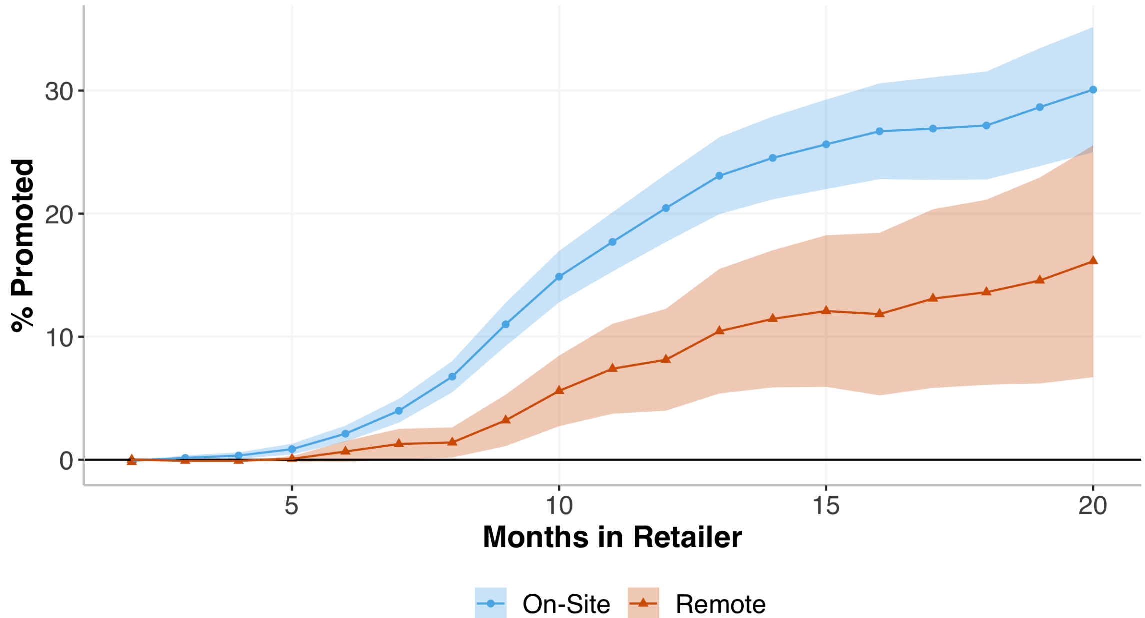

Note: This figure considers the promotion rates of remote workers (in blue circles) and on-site workers (in orange triangles). The x-axis plots the workers' tenure and the y-axis plots the percent who have been promoted among those who persist at the retailer. The error ribbons reflects 95% confidence intervals with standard errors clustered at the worker level. The sample limits to workers hired after July 2018 when the retailer began to hire workers directly into remote jobs and before April 2020 when on-site call-centers closed due to COVID-19.

Figure 2: Market for Remote Work

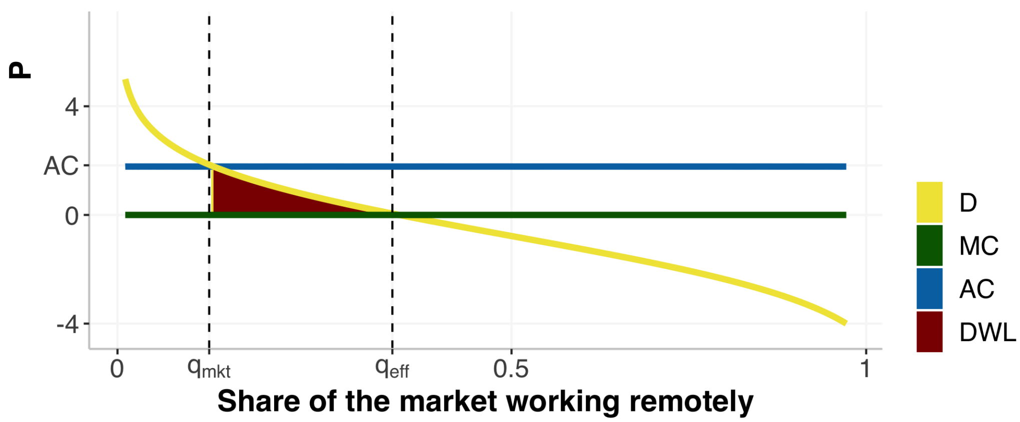

Note: This figure illustrates the market for remote work when there is no treatment effect of remote work on productivity (in green). The x-axis represents the share of the market working remotely. The y-axis represents the price or wage penalty of remote work. The yellow curve is the demand curve for remote work. The average cost (AC) of hiring a remote worker instead of an on-site one in orange reflects self-selection into these jobs. The intersection of demand with AC determines the equilibrium quantity (q_{mkt}), but the intersection with MC determines the efficient quantity (q_{eff}).

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

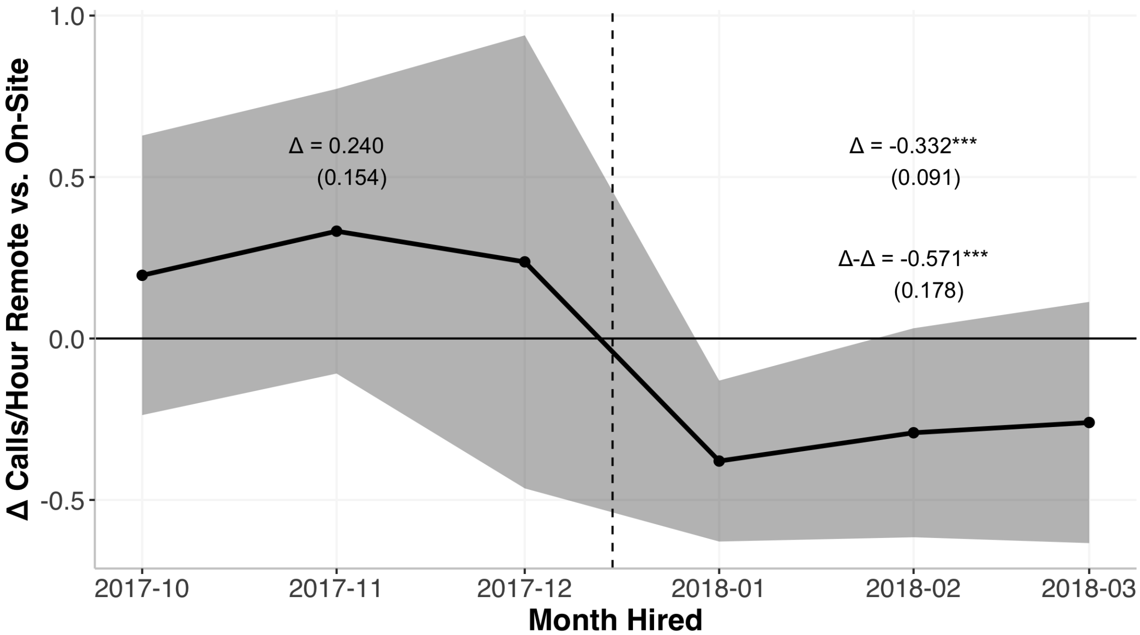

Note: This figure analyzes the introduction of opportunities to go remote into the choice sets of new hires. The x-axis depicts the hiring month. The vertical dashed line indicates January 1, 2018 when the retailer introduced opportunities for on-site workers to go remote. The y-axis depicts the difference in calls per hour of remote and on-site workers from each cohort. The point estimates come from a dynamic version of equation 8, 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 8, 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 could have advanced to specialized roles with non-random routing of calls. The sample further excludes all workers' first three months because productivity data were no longer available for the first three months of the cohort of workers hired in October 2017.

Figure 4: Switches to Remote Work in the 2018 Remote Program

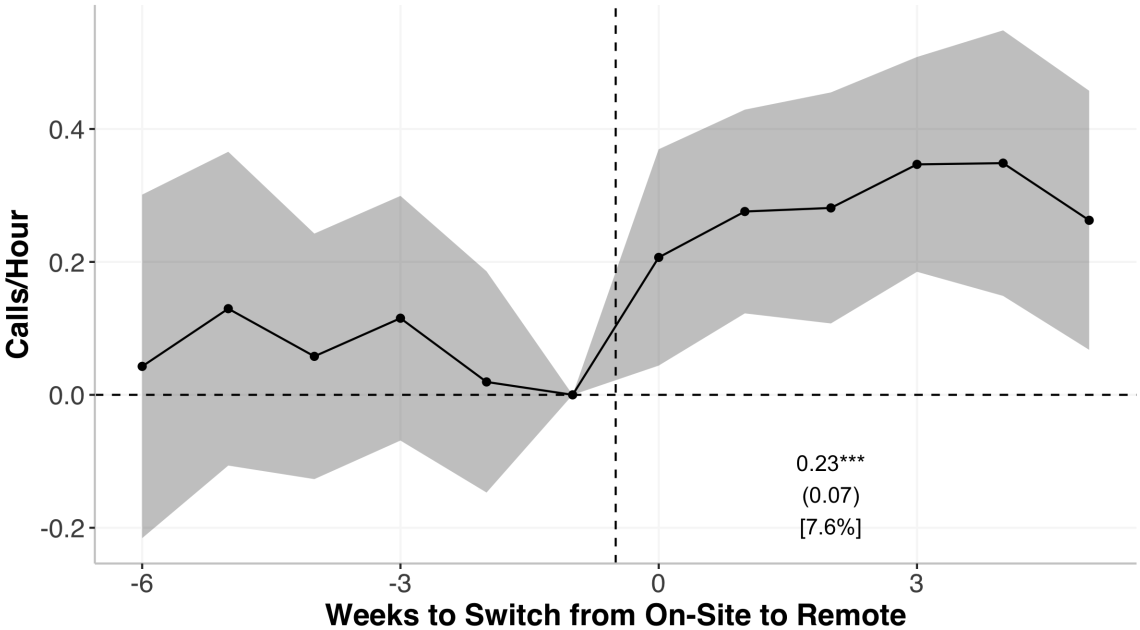

Note: This figure depicts the change in calls per hour of workers who switched from on-site to remote work when the retailer posted remote jobs that were open to on-site workers in 2018. The x-axis represents event time in weeks from the switch to remote work. The y-axis represents the workers' calls taken per hour. The error ribbon reflects a 95% confidence interval, which compares the productivity in each week to that in the index week. Standard errors are clustered at the worker level. The annotated coefficient compares the 6 weeks before and after individuals' switches to remote work as in equation 9. The bracketed number represents the percent change in productivity. This analysis limits to a balanced panel of 127 entry-level workers who made a switch from on-site to remote work that was not proximate to a change in role or departure from the retailer.

Figure 5: Difference-in-Difference in Calls/Hour Around COVID-19 Office Closures

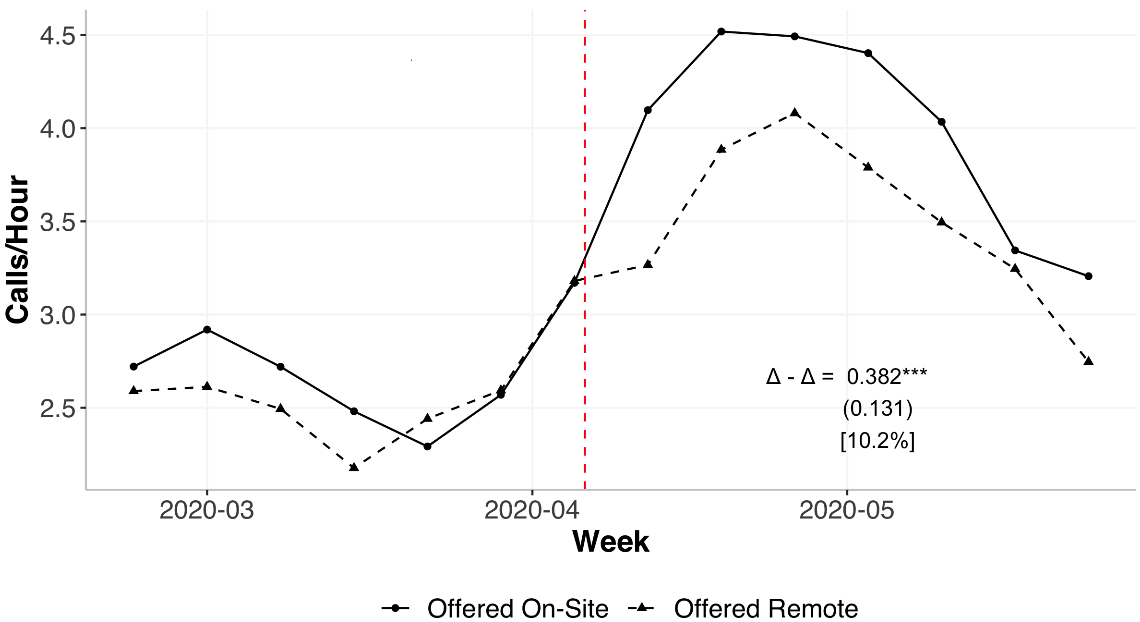

Note: This figure illustrates the difference-in-difference in productivity between on-site workers who were transitioning into remote work (in the solid black line) and remote workers who were already working remotely (in the dashed black line). The x-axis indicates the week in which calls were taken. The y-axis represents calls per hour. The vertical red dashed line indicates April 6, 2020 when 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. The annotated coefficient indicates the difference-in-difference estimated according to equation 11. All standard errors are clustered by worker.

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

			Hired After July 2018			
		All Workers	All Workers	Hired Remote	Hired On-Site	Δ
1.	Calls/Hour	3.0	3.1	3.1	3.2	-0.11**
2.	Customer Satisfaction	4.9	4.9	4.9	4.9	0.00
3.	% with No Call Back	84.2	84.7	84.1	84.8	-0.71***
4.	% Promoted	25.0	9.8	5.0	10.9	-5.91***
5.	% Promoted to Manager	2.3	0.4	0.2	0.5	-0.30
6.	% Turnover	48.9	43.2	45.9	42.5	3.37
7.	% Fired for Performance	2.7	1.7	1.6	1.7	-0.09
8.	Wage	15.0	14.9	14.0	15.1	-1.14***
9.	MSA CSR wage	15.2	15.6	16.0	15.5	0.55***
10.	% Part-time	9.1	6.6	7.4	6.4	1.02
11.	% Female	70.3	73.0	88.2	69.3	18.82***
12.	Age	32.3	32.6	35.6	31.9	3.77***
13.	% Any Caregiver	61.9	62.2	73.8	59.4	14.45***
14.	% Child Caregiver	46.9	46.5	57.7	43.8	13.97***
15.	% Disability Caregiver	17.9	17.8	25.3	16.0	9.36**
16.	% Elder Caregiver	14.4	14.5	15.6	14.2	1.39
	# Workers	4440	2860	560	2300	
	# Caregiving Respondents	1043	757	149	608	
	# MSA	99	99	95	9	
	# Time-zone	4	4	4	4	

Note: This table characterizes on-site and remote workers at the online retailer. Column 1 includes all workers who were hired into entry level roles before the retailer closed its on-site call-centers due to COVID-19 on April 6, 2020. The remaining columns further limit to workers hired after July 2018 when the retailer began hiring workers directly into remote jobs. Column 3 considers workers hired into remote jobs and column 4, workers hired into on-site jobs. Column 5 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. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

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

	Daily Customer Satisfaction After Dec 2018 (in SD)			
	(1)	(2)	(3)	(4)
	All Workers	On-Site	Remote	All Workers
Positive Manager Eval in Dec 2018	0.188***	0.253***	−0.023	0.257***
	(0.050)	(0.057)	(0.107)	(0.056)
Remote x Positive Manager Eval				−0.269**
				(0.116)
Avg. Cust. Sat. July-Dec 2018	0.310***	0.296***	0.368***	0.311***
	(0.037)	(0.046)	(0.051)	(0.037)
Remote				0.185***
				(0.064)
Share Positive Evaluation	0.32	0.33	0.28	0.32
# Workers	896	665	230	896
# Days	126771	93880	32779	126771

Note: 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. Column 1 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. A positive evaluation is defined as a rating of exceptional or highly-effective, the top two categories in the five-point rating system. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

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

	Calls/Hour		Customer Satisfaction		No Call Back	
	(1)	(2)	(3)	(4)	(5)	(6)
Remote x Hired After Intro	-0.571^{***}	-0.402^{**}	0.002	0.011	-0.007	-0.009
	(0.178)	(0.179)	(0.022)	(0.024)	(0.007)	(0.007)
Remote	0.240	-0.019	0.001	-0.002	0.002	0.002
	(0.153)	(0.148)	(0.018)	(0.020)	(0.005)	(0.005)
% Effect	-21.05%	-14.82%	0.04%	0.23%	-0.85%	-1.14%
	(6.54)	(6.61)	(0.45)	(0.49)	(0.80)	(0.89)
t x Hire Month x Time-zone	✓		✓		✓	
t x Hire Month x Hiring Call-center		✓		✓		✓
Dependent Mean	2.71	2.71	4.89	4.89	0.83	0.83
Dependent Std. Dev.	1.44	1.44	0.44	0.44	0.13	0.13
# Workers	465	386	462	383	456	378
# Remote Workers	138	138	136	136	133	133
# On-site Workers	327	248	326	247	323	245
# Days	19594	15969	14733	12121	19147	15584

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 8. The even columns 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. The first two columns consider the average calls handled per working hour; the next two columns consider the average customer satisfaction, which is only available on days where at least one customer left a review; the last two columns consider the share of a worker's customers who do not call the retailer back within two days, presumably because their question went unanswered. 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. 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. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

Table 4: Switches to Remote Work in the 2018 Remote Program

	Calls/Hour		Satisfaction (out of 5)		No Call Back	
	1 week	6 weeks	1 week	6 weeks	1 week	6 weeks
Post	0.207**	0.226***	-0.001	0.003	-0.009	-0.004
	(0.083)	(0.066)	(0.025)	(0.011)	(0.006)	(0.003)
% Effect	7.12%	7.59%	-0.03%	0.06%	-1.07%	-0.49%
	(2.86)	(2.21)	(0.51)	(0.23)	(0.78)	(0.41)
Dependent Mean	2.90	2.97	4.91	4.89	0.83	0.83
# Workers	127	127	127	127	127	127
# Days	1022	6019	806	4803	943	5705

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 4. Each specification estimates equation 9 with standard errors clustered at the worker 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 change in role or departure from the retailer within the twelve-week time-span. The robustness of the results of calls per hour to the choice of bandwidth is considered in Figure A.7. The robustness of the results to comparisons to workers who remained on-site is presented in Appendix A.II. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

Table 5: Lockdown Differences in Productivity: Call Quantity

	Calls/Hour				
	(1)	(2)	(3)	(4)	(5)
Remote	-0.616^{***}	-0.587^{***}	-0.593^{***}	-0.509^{***}	-0.590^{***}
	(0.151)	(0.153)	(0.153)	(0.175)	(0.185)
COVID-19 Deaths per 100K		-0.0004	-0.001	-0.001	-0.001
		(0.001)	(0.001)	(0.001)	(0.001)
COVID-19 Cases per 10K		-0.0004	-0.001	-0.001	-0.001
		(0.0004)	(0.0004)	(0.0004)	(0.0004)
MSA Median Wage in Customer Service			0.068^{*}	0.076^{*}	0.081^{**}
			(0.037)	(0.040)	(0.039)
MSA % in Customer Service				-0.128	-0.125
				(0.140)	(0.141)
Age					0.001
					(0.006)
Female					0.287^{*}
					(0.148)
% Difference	-18.04%	-17.20%	-17.37%	-14.91%	-17.29%
	(4.42)	(4.49)	(4.49)	(5.13)	(5.42)
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

Note: 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 10 with standard errors two-way clustered by worker and date. The dependent mean of calls taken per hour is high during the pandemic because the online retailer saw an uptick in consumer demand as consumers switched away from brick-and-mortar shopping. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

Table 6: Difference-in-Difference in Calls/Hour Around COVID-19 Office Closures

	Calls/Hour					
	(1)	(2)	(3)	(4)	(5)	(6)
Offered On-Site x After Office Closure	0.382***	0.379***	0.563***	0.246*	0.258*	0.302
	(0.131)	(0.116)	(0.126)	(0.138)	(0.141)	(0.204)
Offered On-Site	0.221**					
	(0.109)					
COVID-19 Deaths Per 100K					−0.001	−0.001
					(0.001)	(0.001)
COVID-19 Cases Per 100K					−0.001	0.001
					(0.001)	(0.001)
Child Caregiver x After Office Closure						−0.030
						(0.150)
% Treatment Effect	10.19%	10.13%	14.68%	6.59%	6.9%	7.74%
	(3.50)	(3.09)	(3.28)	(3.68)	(3.76)	(5.22)
Bandwidth	6 weeks	6 weeks	3 weeks	Donut	Donut	Donut
Worker FE		✓	✓	✓	✓	✓
Pre Mean Calls/Per Hour	2.56	2.55	2.43	2.68	2.68	2.79
Post Mean Calls/Per Hour	3.75	3.74	3.84	3.74	3.74	3.90
# Workers	369	280	267	270	270	136
# Formerly On-site, Treated Workers	210	158	158	158	158	70
# Already Remote, Control Workers	159	122	109	112	112	63
# Days	14647	13252	6571	9665	9665	5587

Note: This table evaluates how the productivity change 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 11 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. The second column introduces individual worker fixed effects to address selective attrition out of the sample. The third column compares the three weeks before and after the office closures, during which times most schools were closed. The fourth through seventh 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 seventh column come from a caregiving survey that the retailer fielded in June of 2020. All standard errors are clustered by worker. Figure A.9 considers robustness of the results to different durations of the post period. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.