NBER WORKING PAPER SERIES

WORKING FROM HOME, WORKER SORTING AND DEVELOPMENT

David Atkin
Antoinette Schoar
Sumit Shinde

Working Paper 31515
http://www.nber.org/papers/w31515

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2023, Revised December 2023

We would like to thank the field research team at LEAD at KREA University, especially the office managers Kothandam D and Christopher Glan. We would like to thank Juan Gambetta Rossi, Dahyeon Jeong, Pratibha Joshi, Lipika Kapoor, Mridulya Narasimhan, and Jiulei Zhu for their excellent research assistance. This project was made possible due Financial support from the DFID-ESRC Growth Programme and the International Growth Centre. We would like to thank Dave Donaldson, Daniel Keniston, and all the participants at IPA SME workshop, MIT, UCLA, the University of Bonn, and the 2022 Stanford Remote Work Conference for their valuable feedback. The experiment received IRB approval from MIT, UCLA, and IFMR. All errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Working from Home, Worker Sorting and Development
David Atkin, Antoinette Schoar, and Sumit Shinde
NBER Working Paper No. 31515
July 2023, Revised December 2023
JEL No. J01,J10,J22,J24,J46

ABSTRACT

Differences in the productivity of home-based versus office-based work may arise due to a treatment effect of the office or from workers with different abilities sorting into these locations. We conduct an RCT in the data entry sector in India that exogenously allocates workers to work from home (WFH) or from the office. We first find that the productivity of workers randomly assigned to WFH is 18%lower than those in the office. Two-thirds of the effect manifests itself from the first day of work with the remainder due to quicker learning in the office. Second, there is a negative selection effect into the office, since workers who prefer home-based work are 12% faster and more accurate at baseline. Finally, we find negative selection on treatment effects: workers who prefer WFH are substantially less productive at home than the office (27% less compared to 13% less for workers who prefer the office). These negative selection effects are partially explained by subgroups that face bigger constraints on selecting into office work and additional demands on their attention when at home, such as those with children or other home care responsibilities.

David Atkin
Department of Economics, E52-550
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
atkin@mit.edu

Antoinette Schoar
MIT Sloan School of Management
100 Main Street, E62-638
Cambridge, MA 02142
and NBER
aschoar@mit.edu

Sumit Shinde
Department of Economics,
UCLA
8276 Bunche Hall
315 Portola Plaza
Los Angeles, CA 90095
sumitshinde@ucla.edu

1 Introduction

Several studies document that productivity is substantially lower in small household-based enterprises than in larger firms, see, for example, Hsieh and Klenow (2014), La Porta and Shleifer (2014) or McCaig and Pavcnik (2018). Casual observation also suggests a relationship between the rapid expansions in factory and office employment and high growth rates both in the Industrial Revolution and East Asian "miracle" economies. In the popular debate, these findings are often seen as causal evidence that reallocating workers from household enterprises into offices and factories may play a central role in improving productivity (rather than simply the result of the sorting of high-ability workers into these jobs).

The debate about the cost and benefits of work from home (WFH) has recently received heightened interest, even in developed countries, as a result of the COVID-19 pandemic, which forced many employers to shift to WFH, see Barrero et al. (2021) for a recent review. Yet we still know relatively little about what are the productivity effects of WFH, how much of the differences we observe are the result of selection effects, and whether workers are choosing the work locations that make them most productive.

Productivity may differ across these two types of work environments for several distinct reasons. Most obviously, if production is more efficient or if learning is faster when organized in an office setting, the observed productivity difference between office and home-based work may be driven purely by a treatment effect of office-based production methods (as explored in Bloom et al. (2015), Brynjolfsson et al. (2020), Bloom et al. (2022), Emanuel and Harrington (2023), and Gibbs et al. (2023)).

Alternatively, productivity may differ between WFH and office work due to the sorting of workers with different abilities and preferences. If workers with higher ability or lower cost of effort prefer working in a more structured office environment (or face lower personal costs or social pressures to work in an office), they might select into working in the office more readily than low-ability or low-effort workers—a selection on ability effect. In this case, the office may serve as a sorting device for productive workers and would bias upwards estimates of the productivity advantage of working from the office versus from home in non-experimental settings. The opposite selection effect would result if some high-ability workers have stronger preferences for WFH.

Furthermore, we might expect workers to choose the work location where they will be relatively more productive and thus earn more—selection on treatment effects. If high-ability workers possess skills that complement office work, such as the ability to learn from their peers, this selection on treatment effects could drive the selection of high ability workers into office. Selection would go in the other direction if lower-ability workers benefit more from the discipline of being in the office.

Of course, it is also possible that selection into office work, and work outside the home more generally, could be constrained by factors that are orthogonal to productivity gains or even negatively correlated with them. For example, some workers might have additional demands on their time when at home such as child or elderly care which both reduce productivity but also increase the desire to WFH.[1]

While the organizations literature has explored the impacts of the productivity-enhancing practices used in offices and factories, and more recently the benefits of WFH, the literature examining this second explanation for such productivity differences; that offices and factories may act as sorting devices, is far more limited and almost all non experimental.[2] And we are not aware of any work exploring whether workers select into the work location in which they are most productive.

In this paper, we aim to measure the productivity differences between home-based and office-based production as well as the sources of these differences. To study this question, we set up a randomized control trial in the data entry sector in the city of Chennai, India. The Indian data-entry sector provides an excellent setting to explore these hypotheses. First, it is a sector where working from home is particularly feasible since workers do not need to collaborate with others in the organization (in this sense, we may think of our estimates as a lower bound of the treatment effects of working in an office in more collaborative sectors). Second, it is possible to record productivity and effort in great detail via data entry software on workers' laptops. Third, data entry and business process outsourcing, more generally, is an important and growing sector in India, a country where a large share of production is home based.

Our research design allows us to separate the treatment effect of the office environment from the selection of high-ability and high-effort employees into these more formal work environments. At the same time, we can also test whether social and cultural constraints affect the ability of workers to sort into different jobs. We first established our own data entry operation with several hundred workers in order to control both work conditions and allocation to home and office work. Potential data entry workers were recruited through ads in leading local newspapers. Qualifying applicants were invited to an entry interview where they completed an initial application as well as some brief data entry tasks to ascertain ability at baseline (measured through data entry tests that record speed and error rates). Applicants were asked at this stage for their preferences between office and home work with similar conditions and identical equipment (with the choice incentivized by informing the applicant that the probability of allocation to their preferred group was greater than one half). All applicants were then randomized into either the office or home work treatment for an 8-week data entry job. Workers were also informed that they each had a chance of being recommended to a longer term job upon completion of the 8-week position. Minute-by-minute productivity as well as idle times were recorded through the data entry software.

Our analysis follows three steps: First, treatment effects of home versus office work are measured by comparing the performance of people randomized into home work to those randomized into the office-

[1] In calculating impacts on aggregate output, it is important to recognize that these workers might not have participated in the labor force at all without the flexibility of home-based work.

[2] Emanuel and Harrington (2023) find that after the closure of their office due to COVID-19, previously office-based call center workers were more productive than those who always worked from home. Using experiments, Ho et al. (2023) measure selection and treatment effects for phone-based gig work provided to Indian women who have mostly never worked for pay and Kim et al. (2022) explore whether those choosing part time or full time work are more productive.

based group, independent of their preferences for either work environment. Second, the importance of selection is measured by comparing how workers of different initial ability make incentivized choices between working from home or the office. Finally, our research design also allows us to answer an additional set of questions. Is there a complementarity between certain types of workers and office-based work, and do workers' choices reflect this complementarity? If high ability, high effort types benefit most (least) from office work, sorting into office-based work may magnify (compress) initial productivity differences. To address these questions we also explore selection on treatment effects (i.e., do those with higher returns to office work disproportionately select into it).

Turning to the results of the experiment, we first show that there is a significant negative treatment effect of WFH. The productivity of workers randomly assigned to work from home is 18% lower than that of the workers assigned to work from office. Two-thirds of this difference manifests itself immediately, starting from the first day of work. The remainder is a result of slower learning for the home group over the subsequent eight weeks. These results also hold when controlling for baseline ability as well as when we look at other output measures such as typing speed, the accuracy of data entry, or a measure of data entry speed aligned to worker compensation. The treatment effect of WFH is especially negative when workers are assigned to harder tasks. We find some but relatively limited heterogeneity in treatment effects across worker types. Older female workers, richer workers, and married workers exhibit the strongest treatment effects. For some other groups, treatment effects are indistinguishable from zero (e.g. poorer workers, workers preferring part-time work, and women with family care obligations).

Second, we find a negative selection on ability into office-based work. The workers who state that they prefer WFH are 12% faster, not slower, when their data entry ability is measured at baseline as part of the interview process. They also show higher accuracy of data entry and less idle time. Thus, the productivity advantages of the office derive from treatment effects rather than selection.

What lies behind these negative selection effects? One possible explanation is that low-initial-ability workers have higher treatment effects from working in the office—and therefore are relatively more likely to self select into that environment. For example, low-ability workers might know that they have more to gain from working in the office because they have self-control issues, or need more guidance. Similarly, we would find negative selection into the office if high ability workers believed they were immune to such self control issues or have little to learn from others and so might as well enjoy the convenience of working from home. Our selection-on-treatment estimates reject these explanations. Specifically, we find evidence of negative selection on treatment. Workers who prefer home-based work are 27% less productive when allocated to working from home compared to working from the office, while this gap is only 13% for workers who prefer office-based work. In other words, the workers who prefer working from home have a particularly large negative treatment effect of working from home.

Thus, the selection of high productivity workers to WFH is not because this group does not benefit much from being in the office.[3] In further support of this interpretation, workers who both chose and were assigned to WFH do not use the flexibility of WFH to work more outside regular work hours relative to people who did not choose WFH; nor is this group more productive during off-hours.

Instead, our results suggest that some subsets of workers are constrained from choosing the work location in which they would be most productive. We next turn to investigate the exact form of these constraints and the characteristics associated with more productive workers selecting into home-based work. For example, norms may prevent educated women or those with home-care responsibilities from working outside the house. Alternatively, working in an office may be a status good for low-ability workers even if it does not make them more productive. We find some limited support for these and other hypotheses by including controls for different sets of baseline characteristics and evaluating how much these additional controls attenuate the selection effect. Controls for low status as well as home pressures, responsibilities, and distractions have the most explanatory power. However, even after including all sets of controls, we still find a substantial negative selection effect unexplained by observable characteristics.

Additionally, we conduct an analysis of heterogeneity in the selection on treatment effects along the different observable dimensions detailed above. We find that selection on treatment is particularly negative among five groups for which heterogeneity in constraints may be particularly acute: workers with family care responsibilities—especially women with such responsibilities—workers with low family income, workers with children, and older workers. This marked heterogeneity suggests that workers whose productivity (and compensation) falls relatively more when working from home, are more likely to chose WFH because they derive benefits from the activities that serve as distractions when working from home. For example, among women who have family care responsibilities, those that have no or poor alternative family care arrangements choose WFH and face distractions when they do so. Meanwhile those with good quality care arrangements choose office but also face fewer distractions when working from home.

A caveat in interpreting the magnitudes of the selection effects is that to implement the experiment, we restricted the sample of workers at the interview stage to those who would, in principle, be willing to work in either home or office locations. Thus, applicants with the most extreme preferences were filtered out. These workers would have dropped out from the experiment before starting work were they not allocated to the location of their choice, leading to selective attrition.[4] As we find that the size of the selection effect in the filtered sample is smaller than the full applicant sample, we conjecture that the selection on ability in the population might be larger in magnitude than the one we report.

Overall, our results suggest that, although there are substantial productivity benefits to working in an office, many workers choose to work from home—particularly those who have high ability and those

[3] Most directly, we find no significant interaction between treatment and initial ability.

[4] That said, our sample still included many with strong preferences: in the earliest waves of hiring, we observed substantial differential attrition from groups that did not receive the work location of their choice despite this filtering. Only after introducing a sizable retention bonus were we able to avoid this attrition. Thus, our sample contains applicants who may have strong preferences but can be incentivized to work in an environment not of their choosing.

who would gain the most in terms of productivity from being in the office. Of course, to know whether such choices are optimal from the worker's perspective, we need to better understand their preferences and know more about the nature of the constraints under which they are making their decisions. For example, these patterns are particularly pronounced among those with care responsibilities at home and those with children. Such findings may be rationalized by heterogeneity in family pressures to stay inside the home or heterogeneity in preferences to provide family care or other help around the home during the workday. Whatever their source, our results show that preferences and constraints on the optimal sorting into office- and home-based work result in a significant loss in the productivity of the workforce. These results also raise the possibility that policies that relax the heterogeneity of these constraints, such as providing universal child care, may have substantial effects on aggregate productivity.

1.1 Related Literature

This paper contributes to several literatures in economics. First, we are motivated by the literature that highlights large productivity differences between informal firms, particularly small household enterprises, and larger formal firms—a pattern that seems particularly prevalent in developing countries (see, for example, Hsieh and Klenow (2009) and Bartelsman et al. (2013)). Most relevant, McCaig and Pavcnik (2018) show substantial labor productivity differences between household enterprises and non-household firms in Vietnam.[5] We aim to shed more light on the origins of these differences as well explore the constraints on the optimal sorting of workers into different work environments.

A related literature considers the development of work structures that accompanied the industrial revolution. These papers argue that some of the expansion of the manufacturing sector, and the movement from manufacturing in homes to the factory system, was due to the fact that factory work mitigated worker self-control problems that plagued home-based work (see, for example, Clark (1994), Kaur et al. (2010), Hiller (2011, 2018), Forquesato (2016)). This of course relies on the productivity gains coming from factory work itself rather than worker selection, a hypothesis we test directly.

Finally, there is a fast-growing literature on the productivity effects of working from home. Bloom et al. (2015) find substantial productivity improvements from workers in a large Chinese travel agency who were allowed to WFH for 4 days a week. Bloom et al. (2022) find that hybrid WFH at the same firm reduced worker attrition and had small positive impacts on output. These experimental studies differ from ours in three ways. First, their workers were selected from the subset of workers already at the firm (and in Bloom et al. (2015), who had volunteered to WFH), thereby shutting down the selection channel which is at the center of our analysis. Our work is complementary as we set out to analyze the role that sorting plays in driving productivity differences between home- and office-based work settings. Second,

[5] Relatedly, the organizational economics literature documents the importance of management practices even within large formal workplaces (e.g., Bloom et al. (2013)). For example, Kaur et al. (2015) carry out a range of experimental innovations at a data entry firm in India and show that some workers have self-control issues.

and closely related, the employees in these studies had previously been working in an office environment in the same firm (e.g. there was a six month minimum tenure requirement in Bloom et al. (2015)) and thus might have already absorbed the productivity-enhancing work habits that office work may foster. In contrast, our study population is poorer and less educated, and many applicants had not previously worked in formal office environments, let alone an office-based data entry job. Third, only workers with a private room at home were eligible to participate in Bloom et al. (2015). About one third of the productivity improvement they find comes from a quieter work environment for making customer service calls. In our setting, the office location is likely to be more not less quiet than the home location given the distractions and close quarters of a typical Indian urban household, and the high levels of noise pollution in Urban India. More recently, Ho et al. (2023) conduct an experiment in India asking whether digital gig-work jobs can increase female labor force participation by offering households jobs of varying flexibility. In contrast, we compare full time office and home work in a more urban setting.

There is also a recent non-experimental literature showing that home-based work arrangements lead to lower productivity. For example, Gibbs et al. (2023) and Emanuel and Harrington (2023) both find similar negative productivity effects to ours when studying IT professionals and call center operators, respectively, using first differences and difference-in-difference methodologies (see also Yang et al. (2022), Monteiro et al. (2021), Künn et al. (2022), Morikawa (2023), Shen (2023)).[6] Additionally, in line with our findings, the flexibility allowed by these new work arrangements is utilized by workers and they report better work-life balance (e.g. Choudhury et al. (2022), Angelici and Profeta (2023)). As discussed in the introduction above, Emanuel and Harrington (2023) also explore selection effects, although not selection on treatment, finding in their case that those workers who had WFH jobs prior to COVID-19 continue to perform worse than office workers in the same firm whose jobs were moved to home due to the pandemic. Our experimental design allows us to deal with potential confounds inherent in the comparisons of productivity changes across different types of workers pre and post large shocks that change work assignments and/or preferences such as the COVID-19 pandemic.

2 Theoretical Motivation

Worker productivity may differ between home and office work for at least two distinct reasons—a *treatment effect* whereby WFH has a causal impact on productivity, and a *selection effect* wherby ex-ante more productive workers sort into home or office work environments. These two mechansims may also interact if there is *selection on treatment effects*. In this section, we lay out the theoretical foundation for each of these. Our experimental design serves to measure the size of these three forces. Furthermore, by exploring how these effects vary with worker characteristics, we hope to shed light on the role of constraints and preferences in shaping worker productivity and work location choices.

[6] In contrast, hybrid work arrangements appear to have some positive impact on productivity (see, for example, Choudhury et al. (2022), Boltz et al. (2023), and Angelici and Profeta (2023)). Monte et al. (2023) explore the implications for city structure.

Treatment effects As discussed in the literature cited above, the office may provide a more productive work environment, more opportunities for learning when surrounded by supervisors and peers, and stronger incentives due to better monitoring. In contrast, home work may be more productive if there are fewer distractions from colleagues, work can be done during more productive hours, or workers are less weary from a long commute. We call any differences in productivity between home and office work environments, holding fixed the characteristics of the workers, the *treatment effect* of WFH.

Selection on ability Personal preferences or societal forces may lead workers to sort based on ability. One reason for such sorting is that the office is more demanding given fixed schedules, stricter norms, and more peer pressure—demands less costly or unpleasant for more productive workers. There may also be long-term benefits (e.g. promotions) from office work due to greater interactions with supervisors—an attractive feature to ambitious types. In either scenario, higher-ability workers would sort into the office. The productivity impacts of such sorting are magnified if high-ability workers are complements with each other in production, either through peer learning dynamics or an O-ring production function.

Selection on ability may also occur due to preference-relevant characteristics that correlate with productivity. For example, in many conservative societies, women are not allowed to work outside the home to limit their interactions with men. Conversely, men who work at home may be stigmatized. If women are highly productive, as has been noted in light manufacturing and garment production, this may generate selection on ability. The strength of these social and cultural sanctions may also vary with household wealth and education, potentially generating selection within genders. Relatedly, office work might be a status good, particularly so for workers with low social status. In these scenarios, lower productivity workers may sort into office work.

Selection on treatment effects If workers experience heterogeneous impacts of home versus office work, those with relatively high returns to office work may be more likely to select into such environments (at least if pay responds to performance). This selection on treatment effects may drive selection on ability, with the sign depending on whether the productivity-enhancing features of the office complement or substitute for ability (e.g., do higher-ability workers learn more from those around them or do lower-ability workers gain more from being close to supervisors). Alternatively, the discipline offered by the office may attract those with self-control problems who expect to procrastinate when working from home.[7] Whether such sorting generates positive or negative selection on ability is unclear—higher-ability workers may be more conscious of their self-control problems, but such problems may plague lower-ability workers relatively more (e.g., because they face greater distraction or less privacy at home).

Selection on treatment effects is also of independent interest. Do workers chose the work location where they are relatively more productive and thus relatively better compensated? If not, what other constraints or preferences lie behind their decision?

[7] We can think of office work as providing delayed benefits (i.e., higher wages) with an upfront cost (i.e., more effort now).

3 Research Setting

3.1 Context and Implementation

This study implements a randomized control trial in the data entry industry in the South Indian city of Chennai. This sector provides a number of benefits for our analysis. First, this type of work is very widespread in India and hence well-known to potential applicants. Second, it has relatively low skill requirements. Third and fourth, due to the discrete nature of the tasks, work can easily be done from home without support from colleagues and using the same technology as in an office setting. This last feature is crucial to ensure that productivity differences are not driven by the use of different technologies. Fifth, it is straightforward to collect detailed productivity and output measures from data entry work (e.g., input per minute, errors, time working, etc.). Furthermore, this type of data collection is common in the industry, allowing us to avoid imposing an artificial monitoring system.

We established a data-entry operation with the option of both home- and office-based work.[8] The operation was managed by professional data entry supervisors who had previously worked in the data entry industry. We also worked closely with a data entry firm in Chennai so that the upfront training, technical help with equipment problems, and compensation schemes mimicked a typical data entry firm in the city. The workers in both the office and at home were provided with identical work assignments and identical laptops to complete the data entry tasks.[9] To ensure that the two environments were as comparable as possible, workers were required to work for 35 hours per week in both locations.

The type of work, the wage structure, the criterion for not being fired, weekly targets, and managers were also identical. In the office environment, we had up to 25 workers working from 9 am to 5 pm for five days a week. In the case of the home environment, workers came into the office every Monday morning to submit the work done and receive new assignments. Like office workers, workers in the home environment had to work 35 hours per week, but unlike office workers, home workers had flexibility regarding when to work (both within and across days). To ensure each worker at home completed their own data entry tasks and did not outsource them to somebody else, we implemented a monitoring system that involved the use of the inbuilt laptop camera to take low-resolution pictures of the person working on the laptop every 15 minutes.[10]

We purposely held the amount of supervisor support and supervision as similar as possible across both work locations, consistent with existing data entry operations. Specifically, there were weekly meetings between the worker and manager regarding their progress. Additionally, both workers in the office and at home could reach out to managers with any queries, in the latter case via a special telephone hotline we

[8] Pictures of the office and a few sample home settings can be found in Appendix Figure A.1.

[9] A picture of the user interface for a sample data entry task can be found in Appendix Figure A.2.

[10] We first explained this monitoring system to all workers and then obtained informed written consent prior to the beginning of the work. The experiment received IRB approval for capturing pictures of workers.

set up. Thanks in part to the recent uptake in WFH, new management technologies have been developed to better support home-based workers. But these methods typically are aimed at higher level functions than the workers in our experiment.

To mimic a real job, all workers were offered a contract for 8 weeks of work. After the 8 weeks, workers were provided with references and training certificates and were matched to an employment agency to help find future employment in the industry.

We constructed the data entry tasks that the workers had to complete. Each data entry task consisted of four sections, each focusing on a different type of data entry, such as entering type-set text, entering strings of random alpha-numeric characters, etc. We had two levels of difficulty. The "difficult" tasks had an identical structure to the "easy" tasks, but the difficulty was increased. For instance, the type-set text was replaced by handwritten text, and strings of random alpha-numeric characters were replaced by strings of both alpha-numeric and special characters, which made typing difficult (examples can be found in Appendix Figure A.3). Workers were assigned easy tasks from weeks 1 to 3, harder tasks in weeks 4 to 6, and a random mix of both difficult and easy tasks in the last two weeks. Figure 1 presents the timeline from recruitment through week 8.

Figure 1: Worker Timeline

Recruitment	
Week -1	Applicants with relevant characteristics are invited to in-person screening via newspaper ads
Day 0	On-site evaluation of applicants, recorded home versus office preference, initial typing speed test
3 Day Training	
Day 1–3	Training and orientation, do incentivized and non-incentivized typing speed tests
Day 3	Workers are allocated to office or home work
Work Assignment	
Week 1–3	Easy data-entry tasks
Week 4–6	Difficult data-entry tasks
Week 7–8	Both easy and difficult data-entry tasks
Week 8	Job ends, training certificates and references provided

3.2 Recruitment and Sample Selection

To hire workers, entry-level data entry jobs were advertised in the jobs section of the main local newspapers. The objective was to reach potential employees aged 18–40 who lived in lower middle class localities and suburbs of the city (the target population for these types of jobs). Those interested in the job were asked to show up for an in-person 'walk-in' interview at the office location during the following week.

Two different types of newspaper ads were placed—one type advertising for home-based data entry jobs and another type advertising for office-based data entry jobs (see Appendix Figure A.4 for examples). We found limited heterogeneity based on the type of ad so our analysis combines the workers attracted by both samples, with results broken out by ad type relegated to the Appendix.

The interview process was designed to both elicit baseline worker characteristics and initial typing speed and accuracy. Applicants had to answer a number of interview questions as well as perform typing speed tests. Furthermore, we asked applicants to state their preference between home and office-based work. This question was incentivized as applicants were told that they would be more likely to get their first choice than their second but that it was not guaranteed.

We imposed two screening criteria on the applicants attending the walk-in interviews. First, the applicants had to be aged 18–40. Second, they had to confirm that they were willing to work in either a home or office environment if they were not allocated to their first-choice work location. Approximately half the applicants passed this screening and were invited to participate in three days of paid training at the office location. Ultimately, the non-pilot phases of the experiment recruited 235 workers from an applicant sample of 892 over a period of 15 months beginning in January 2017.[11] Workers were hired in batches as we were constrained by the office size which could only accommodate 25 workers.

3.3 Intervention

Once work location preferences were elicited, workers were randomly allocated a work location.[12] Four groups were formed through this process: Preferred home, allocated home (labeled HH); Preferred home, allocated office (HO); Preferred office, allocated home (OH); Preferred office, allocated office (OO).

The randomization allows us to estimate the treatment effect of being allocated to home or office independent of a worker's preference. Furthermore, the allocation to home or office work conditional on a worker's preference allows us to estimate selection on treatment effects. Specifically, we can compare the difference in productivity between the office and home (using the random assignment) for the group who preferred home work to the same productivity difference for those who preferred office work.

[11] Although around half of the 892 applicants were invited to the training, only 280 showed up and, of these, 45 dropped out prior to the beginning of the work. Hence the working sample consists of 235 workers.

[12] Due to an implementation error by the field team, workers were given their preference with $p = 0.5$ rather than $p = 0.55$ as instructed (as the workers were still informed that $p > 0.5$, they were incentivized to report truthfully despite the error).

3.4 Compensation Structure

The compensation structure provided to workers was designed to mimic a typical data-entry firm in the market. Both office- and home-based workers faced an identical structure. Additionally, both sets of workers were compensated for monetary travel costs incurred to reach the office (either to work every day for the office group or to submit and pick up assignments once a week for the home group).

Our compensation structure consisted of a fixed component and a performance-based variable component. The fixed component was equal to INR 8500 ($ 128.80)[13] per month which workers were eligible to receive on completing 35 hours per week and a target number of data entry tasks. These task targets increased each week to accommodate learning. If workers failed to meet either requirement for three weeks, their contracts were terminated. In addition, there was a performance-based component with an additional INR 65 ($ 1) for every task completed beyond the weekly target. A retention bonus of INR 2000 ($ 30.30) was paid after the completion of week 1 (Appendix Table A.1 provides further details).

To incentivize the accuracy of completed tasks, mistakes were penalized as follows. We first sorted all tasks completed by each worker during a week by their accuracy. Their most accurate tasks counted towards their weekly task target (18–26 tasks depending on the week). Additional completed tasks counted towards the variable component of the compensation. On these tasks, for a moderate amount of errors, the per-task variable component of INR 65 declined proportionately with the share of errors. For larger error rates, the penalty increased more than proportionately to 1.5 and then 2 times the error rate. Appendix Table A.2 presents the penalty structure and error-rate thresholds in full.

3.5 Outcome Measures

As part of the hiring process, we collected information on demographics, education, data entry and other work experience, employment status, job search, work preferences, and family care and other time commitments. During the training period, a baseline survey collected further details on these topics and covered additional domains such as household characteristics, income, social and economic status, and computer literacy. In addition, workers had to take an aptitude test, a personality test, a risk preference test, and a time preference test.

To gauge the baseline ability of applicants, three speed tests were carried out prior to the random allocation to home/office. As mentioned earlier, during the job interview all applicants were required to complete an hour-long typing test that could be done by a novice with no introduction to data entry. During the training, workers were required to complete both a cash-incentivized and non-incentivized typing speed test lasting 25 minutes. The incentivized test payed a reward based on the total number of correct characters. All three tests were conducted in an office where the interview and training took place.

A variety of data entry job outcomes were collected over the 8-week work period. We hired developers

[13] We use the average exchange rate between Indian Rupees and US Dollars during the experiment (INR 66 $\approx$ $ 1).

to create proprietary data entry software which kept detailed logs of data entry tasks, keystrokes typed, accuracy, and time spent working or idle for each worker. The measure of accuracy is defined to be the proportion of correct entries to total entries. The main productivity measure that we use is net typing speed which is defined as correct entries typed per minute. These records, as well as separate attendance records, reveal the hours worked each week and attrition for both home and office workers.

3.6 Attrition

In the first waves of the experiment, we had a simpler compensation structure and experienced substantial attrition in the first few days of work. That attrition was also highly heterogeneous across intervention groups, with workers not receiving their preferred location much more likely to attrit. This was particularly true for the 50 workers in these early waves who preferred home but were allocated office with 40 quitting immediately upon learning their assignment.

To address attrition and incentivize workers to stay longer, we adjusted the compensation structure in later waves. Most importantly, the retention bonus of INR 2000 ($ 30.30) was introduced, paid upon completing the first week of work.[14] This amount approximately equalled a worker's average weekly earnings. The changes substantially reduced attrition for all groups, and crucially there were no longer differences in attrition between those allocated their preferred choice and those not.[15] Appendix Figure A.5 and Table A.4 present this analysis. As differential attrition complicates the interpretation of treatment effects results, the analysis presented in the main text focuses on these later waves when this issue was addressed via the retention bonus. This leaves us with 280 workers, of whom 235 completed training and commenced work. We relegate results for the earlier (high-attrition) waves to Appendix Table A.7.[16]

4 Treatment, Selection, and Selection on Treatment Effects

4.1 Baseline Characteristics

We first check that our randomization led to balance on baseline characteristics for the groups of workers assigned to the home and office work locations. Columns (1)–(3) of Table 1 compare the 124 workers who were randomly assigned to work from home to the 111 workers who were randomly assigned to work in the office. The two groups are balanced in terms of our measures of baseline worker productivity, either

[14] Additionally, the initial filtering was strengthened. In earlier waves, our surveyors would ask whether the worker was willing to work in either environment and filter out those who were not. In later waves, the office managers would further probe whether the worker was sure of their answer. To limit experimental costs the job duration was also reduced from 12 to 8 weeks. Appendix Table A.3 presents a complete list of modifications.

[15] In the pre-bonus waves, 32%, 28%, 80%, and 12% of workers in the OO, OH, HO, and HH groups dropped out before the work began, respectively. These proportions dropped to 19%, 15%, 18% and 10% post bonus.

[16] The sample from the earliest (high attrition) waves also contains relatively more applicants who had been out of the labor market for extended periods and, thus, is less representative of the flow population that enters the job market. Eight advertisements over three months attracted 79 applicants per ad in the early waves while 33 ads over the subsequent 16 months only attracted 27 applicants per ad.

Table 1: Baseline Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
	Assigned:			**Preferred:**		
	Home	Office	P-Value	Home	Office	P-Value
N	124	111		87	148	
Preferred home work (=1)	0.37	0.37	0.98			
Speed Tests						
Walk-In Speed	26.9	26.0	0.51	27.7	25.7	0.15
Cash Incentive Speed	33.1	33.4	0.78	35.9	31.7	0.00
No Incentive Speed	29.8	29.6	0.85	32.2	28.3	0.00
Demographic						
Female (=1)	0.58	0.43	0.02	0.49	0.52	0.70
Age (years)	24.7	25.3	0.38	26.1	24.3	0.00
Married (==1)	0.20	0.22	0.78	0.31	0.15	0.00
Number of Children	0.21	0.20	0.87	0.25	0.18	0.29
Has Family Care Responsibilities (=1)	0.04	0.10	0.07	0.11	0.04	0.03
Monthly Family Income (INR)	21,149	19,104	0.36	20,684	19,889	0.73
Commute Distance (KM)	13.0	12.5	0.68	12.3	13.0	0.55
Education						
Education (Years)	15.4	15.6	0.36	15.4	15.5	0.57
Used Computer Before (=1)	0.98	0.91	0.03	0.95	0.94	0.63
Typing Course—Self Reported (=1)	0.44	0.38	0.31	0.45	0.39	0.40
Typing Course—Showed Certification (=1)	0.21	0.11	0.03	0.14	0.18	0.45
Work						
Work Exp (Years)	2.1	2.6	0.24	3.3	1.8	0.00
Number of Previous Office Jobs	1.1	1.2	0.40	1.4	0.9	0.00
Unemployment Duration (Months)	3.0	3.1	0.41	3.0	3.1	0.54
Miscellaneous						
Never Leaves Things to Last Minute (Rank 1–6)	3.1	3.0	0.70	2.8	3.2	0.11
Estimated Time Discount Rate	0.98	0.95	0.19	0.99	0.95	0.05
Prefers Full-Time Job (=1)	0.94	0.98	0.13	0.93	0.98	0.06
Additional Study/Job Search Commitments (=1)	0.32	0.35	0.64	0.32	0.34	0.72

Notes: This table compares baseline characteristics between workers randomly assigned to work from home and from the office (columns (1)–(3)); and between workers who preferred to work from home and from the office (columns (4)–(6)). Columns (1) and (2) display mean values of characteristics for workers assigned home and office, respectively. Column (3) displays P-values for the test that there is no difference between means. Columns (4)–(6) repeat the exercise for workers who preferred home or office.

measured by the speed test conducted during their initial interview or the incentivized and unincentivized speed tests administered as part of training. We also find no differences in the proportion of workers who preferred WFH across the randomly assigned groups (37% of workers preferred WFH for both groups). Significant differences (at the 10% level) appear for only 4 out of 22 characteristics.[17]

[17] Of the workers who were assigned WFH, 58% are women whereas only 43% are women in the assigned-office group (significant at 2% level). The home group has 6% fewer workers with family care responsibilities, 7% more workers who have

The last three columns of Table 1 compare the characteristics of the 87 workers who preferred to work from home to the 148 workers who preferred to work from the office. Unlike columns (1)–(3) where we compare workers across randomized work environment allocations, preferences for workplace type are non-random and correlated with worker characteristics. In terms of demographics, workers preferring home are 1.8 years older on average, are 16% more likely to be married, and 6% more likely to have family responsibilities. They also have more years of work experience, held a higher number of office jobs previously and were less likely to prefer a full-time job. We explore differences in baseline productivity across these two groups when analyzing selection on ability in Section 4.3.

4.2 Treatment Effects

To estimate the impact of the random assignment to working from home on worker performance, Table 2 reports the results of running the following regression specification:

$$\text{Worker Performance}_{i,t} = \alpha \text{Alloc\_home}_i + \gamma X_{i,t} + \epsilon_{i,t} \tag{1}$$

Worker Performance$_{i,t}$ is one of our outcome measures described below, measured for each worker i and each task t. Alloc_home$_i$ is a binary variable that takes a value of one if worker i was randomly assigned to work from home; $X_{i,t}$ includes three sets of fixed effects that serve as controls: wave fixed effects picking up temporal differences in the quality of each cohort hired, week fixed effects capturing the week of employment the outcome is measured in (ranging from week 1 to week 8),[18] and section fixed effects capturing the type and difficulty of data entry task being performed (digitization of forms, surveys, numerical tables, or text, all crossed with difficulty level). Our unit of observation is the performance on a particular data entry task, e.g. an individual survey that takes about 2 hours to enter. However, the regression is essentially at the individual level as we re-weight observations such that each worker has a total weight of 1 over all his or her observations and standard errors are clustered at the individual level.

Our primary measure of worker performance is log(net speed) where net speed is defined as correct entries typed per minute. Column (1) of Table 2 shows that workers randomly allocated to work from home exhibit 18% lower net speed. This effect is statistically significant at the 1% level. Columns (2) and (3) report treatment effects for gross speed and accuracy, defined as total entries typed per minute and the ratio of correct entries to total entries typed, respectively.[19] Employees working from home have 12% lower gross speed and 2.48 % lower accuracy. Thus, the lower net speed is mostly due to lower gross speed.

The magnitude of the treatment effect is larger when we use alternative measures of worker performance. In column (4) of Table 2, we explore whether the treatment effect changes with the difficulty of the underlying data-entry task by limiting the sample of data-entry tasks only to hard tasks (which

used a computer before, and 10% more with a certificated typing course (significant at the 7%, 3%, and 3% level, respectively).

[18] We use week fixed effects instead of finer day ones because home workers had the freedom to work any day of the week.

[19] Net speed, gross speed, and accuracy are related as follows: net speed $=$ gross speed $*$ accuracy.

Table 2: Treatment Effects

	(1) Log(Net Speed)	(2) Log(Gross Speed)	(3) Accuracy (in %)	(4) Log(Net Speed) Hard Tasks	(5) Log(Net Speed) With Penalty	(6) Time on Data Entry (hrs/wk) All Weeks	(7) Time on Data Entry (hrs/wk) Complete Weeks	(8) Prop of Time 9–6, M–F	(9) Idle Time (in %)
Alloc_home	-0.18***	-0.12***	-2.48**	-0.30***	-0.24***	-2.71*	-0.041	-0.51***	2.46***
	(0.050)	(0.034)	(1.14)	(0.066)	(0.078)	(1.43)	(0.22)	(0.020)	(0.84)
Constant	3.45***	3.67***	81.7***	3.47***	3.08***	34.7***	33.7***	0.95***	15.6***
	(0.057)	(0.041)	(1.45)	(0.047)	(0.10)	(0.93)	(0.13)	(0.017)	(0.84)
Section, Week, and Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	138,646	138,646	138,646	72,625	138,646	1,880	1,128	1,451	138,646
R-squared	0.260	0.297	0.331	0.108	0.307	0.121	0.012	0.800	0.077

Notes: This table presents estimates of the treatment effects of randomly allocating workers to home- or office-based work environments. The specification given by Equation 1 regresses a worker performance outcome on the dummy variable Alloc_home_i that takes the value one if a worker is randomly assigned to work from home. The outcome in column (1) is the log of net speed, with net speed defined as the number of accurate characters typed per minute. Column (2) considers log gross speed, the log number of total characters typed per minute. Column 3 explores accuracy, the percentage of characters typed that are accurate. Column (4) again uses the log of net speed but only considers performance on hard data entry tasks. Column (5) examines the log number of remunerated characters typed per minute (the total characters typed minus an exponentially increasing penalty for incorrectly typed characters). The outcome in columns (6) and (7) is the time spent actively entering data (in hours per week). Column (6) considers all eight weeks for all workers (including those when workers received a warning, were fired or quit) while column (7) considers only those weeks where workers had met their 35 hours of work target (thus excluding attrition and warning weeks). In column (8), the outcome is the proportion of work done during office hours (i.e., between 9 am and 6 pm, Monday to Friday), restricting attention to weeks when workers had done some work (including the warning weeks but not weeks during which they quit mid-week). Column (9) explores idle time, the percentage of the total data entry time spent where there was no input on the mouse or keyboard. All regressions account for variation arising from the type and difficulty of the task being attempted, the week of employment, and the cohort of workers using section, week, and wave fixed effects, respectively. Except for columns (6)–(8), the unit of observation is the individual-survey task pair. In columns (6)–(8), the unit of observation is the individual-week pair. Despite observations being at the survey task or week level, all regressions are re-weighted to give a total weight of one to each worker across all observations. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

would require workers to concentrate harder and expend higher cognitive effort). We find that participants assigned to WFH display 30% lower net speed on hard tasks. To incentivize workers to make fewer errors, we imposed an exponentially increasing mistake penalty that followed industry norms. We find that the magnitude of the treatment effect is larger at -24% when measured by the remunerated speed that punishes errors more heavily than net speed and determines the variable pay component (column (5)).

One key benefit of working from home is the flexibility that it affords workers regarding their time use. We first explore how total time worked differs across home and office. Recall that, irrespective of work location, all employees were mandated to work 35 hours per week.[20] The 35 hours constituted two components—time spent working on data entry tasks and time spent on ancillary tasks pertaining to data entry (such as checking lists of completed and remaining data entry tasks for the week, checking performance and pay for prior weeks' work)

For the first measure of time spent on data entry, column (6) considers all worker-week pairs of observations including the weeks when the worker failed to complete the work hours targets and either received a warning, was fired, or quit. We find that workers randomly assigned to WFH worked 2.71 fewer hours (significant at the 10% level). Next, we exclude those worker-week pairs where the worker did not complete the mandated work hours. This allows us to focus on time spent on data entry rather than ancillary tasks, absent attrition effects. We find that employees across both locations spent 33.7 hours actively entering data with no significant difference across work locations (column (7)).[21]

We now turn to when the work was done. Individuals working in the office completed 97% of their work during office hours (i.e., between 9 am to 6 pm Monday to Friday) (column (8)). The remaining 3% came from employees being allowed to stay later in the office to compensate for public and personal holidays. On the other hand, only 46% of the work done by home employees was done during these office hours, indicating that home-based workers used the flexibility afforded to them (we explore the use of flexibility further in Section 4.2.3).

Finally, along with choosing when to work, WFH provides workers greater autonomy regarding breaks during working hours and potentially helps workers deal with moderate distractions. The software measured intervals of time when no action was performed by the worker using either the mouse or keyboard while logged in to the data entry system. We define the ratio of the total time spent in such intervals to the total time spent logged in as idle time—a measure of small breaks and distractions while working. Employees working from the office spent 14.6% of their time idle, and this rose by a significant 2.46% for those working from home (column (9))—although this additional idle time explains only a small fraction of the total productivity gap between home and office.

[20] Specifically, the software would not allow additional work once 35 hours had elapsed (workers could log out at any point and log back in with such a break not counting against their 35 hours).

[21] While our software had a feature indicating whether the worker had completed the mandated 35 hours each week, the data logs only saved measures of time spent working so we impute that the rest of the time was spent on ancillary tasks.

4.2.1 Robustness

In Table 3, we run several robustness checks to explore the sensitivity of the treatment effect estimates. Column (1) repeats our baseline estimate (column (1) of Table 2 above). Column (2) controls for workers' baseline speed during the cash-incentivized speed test. This control should increase precision and control for bias if, despite randomization, initial performance differences are driving lower productivity. The 18% lower net speed persists with a small decrease in the standard error of the coefficient.

Recall that we focus on the later waves where we resolved the issue of selective attrition via a retention bonus. Column (3) expands our sample to include the workers from these pre-bonus waves as well. The treatment effect remains unchanged at -18% and standard errors fall.

Our baseline specification re-weights each worker-task observation such that each employee has equal weight. Thus, an individual data entry task receives lower weight for workers performing more tasks, either because they were faster or they attrited later. Instead, column (4) weights each worker-task observation equally. The table shows that the 18% lower productivity from working from home is almost unchanged, falling only slightly to 16%.

Table 3: Treatment Effects—Robustness Checks

	(1) Log(Net Speed) (Baseline)	(2) Log(Net Speed)	(3) Log(Net Speed) All Waves	(4) Log(Net Speed) Task Weights	(5) Log(Net Speed)
Alloc_home	-0.18***	-0.18***	-0.18***	-0.16***	-0.20***
	(0.050)	(0.043)	(0.042)	(0.054)	(0.049)
Initial Log(Net Speed)		0.75***			
		(0.14)			
Constant	3.45***	0.78	3.41***	3.62***	3.42***
	(0.057)	(0.50)	(0.080)	(0.040)	(0.095)
Characteristics Controls	No	No	No	No	Yes
Section, Week, Wave FE	Yes	Yes	Yes	Yes	Yes
Observations	138,646	138,646	213,859	138,646	138,646
R-squared	0.260	0.333	0.268	0.266	0.266

Notes: This table explores the robustness of our treatment effect estimates. In all columns, the dependent variable is the log of net speed and the independent variable is Alloc_home$_i$, a dummy variable that takes a value equal to one if the worker is randomly assigned to work from home. Column (1) repeats our baseline estimate from column (1) of Table 2. Column (2) controls for Initial Log(Net Speed), the log of net speed during the incentivized speed test conducted during training (prior to assignment). Column (3) includes observations from both pre and post-retention bonus waves. In column (4), each worker-task observation is weighted equally (in all other columns, each worker receives equal weight). Column (5) controls for four characteristics for which we observed baseline imbalance (gender, family care responsibilities, prior computer usage, and a typing certification). All regressions account for variation arising from the type and difficulty of the task, duration of employment, and cohort of workers using section, week, and wave fixed effects, respectively. For all specifications, the unit of observation is the worker-task pair. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Finally, recall that the two groups formed by randomly assigning work locations are not perfectly balanced. Column (5) controls for the four baseline characteristics for which there was imbalance: gender, family care responsibilities, prior computer usage, and typing certification. Treatment effect marginally increase to 20%. Taken together, there is strong evidence that workers are less productive when completing the same number of work hours in a home environment compared to an office environment.

4.2.2 Cumulative Learning

We now analyze how much of the positive treatment effect of the office is due to differential learning versus a static effect that is apparent from day one. Workers in both home and office locations experience an increase in productivity over time. The top-left panel of Figure 2 plots the average net speed of workers in both locations over the 8-week employment spell, with the drop in week 4 coming from the assignment of harder tasks in weeks 4 to 6 (with a mix in weeks 7 and 8). The top-right panel of Figure 2 separately plots the average net speed for each difficulty level. Finally, the bottom-left panel plots cumulative learning—the percentage increase in net speed relative to the speed on the first four tasks (about a day of work) of that difficulty level—against the number of weeks of experience the worker has with that type of task.

Figure 2: Learning Over Time

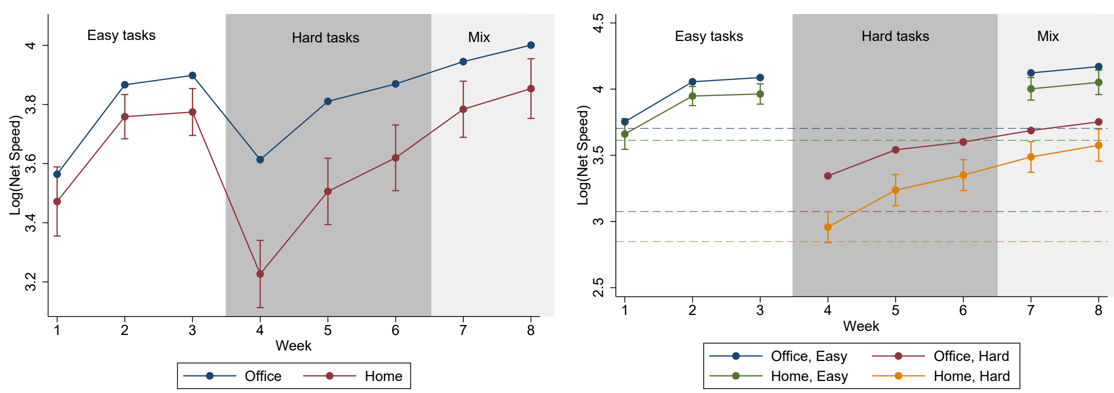

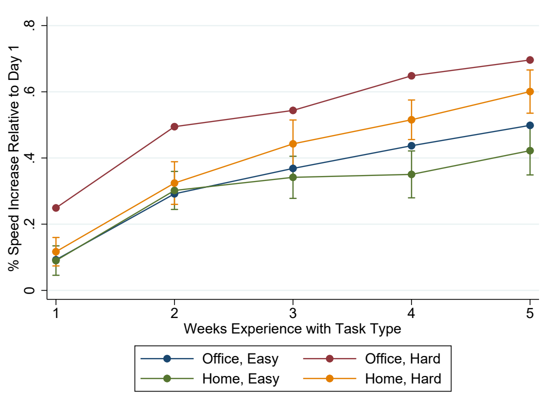

Notes: These figures plot net speed and learning over time. The top-left panel plots the average net speed for home- and office-based workers over 8 weeks. The top-right panel plots the average log net speed separated by task difficulty. The dashed lines represent the average log net speed on the first day for each difficulty-work location. The bottom-left panel plots average cumulative learning for home- and office-based workers for each difficulty. Cumulative learning is defined as the percentage increase in net speed compared to day-1 speed for that difficulty type (i.e. the first 4 tasks). To compare cumulative learning across difficulty levels and work locations, we plot the increase in speed against the total weeks of experience a worker has with that difficulty level.

Learning, in both locations and for both difficulty types, is high in the first few weeks a task is attempted with the rate of improvement slowing in later weeks. Office workers are always more productive, particularly so for hard tasks. And the learning in the first week is particularly substantial for office workers performing hard tasks (i.e. the increase in speed in week one relative to the first four tasks they did of that type, shown in the lower-left panel). In subsequent weeks, the gap between office and home workers narrows slightly for hard tasks while, if anything, widening for easy tasks.

To quantify how much of the total productivity advantage of the office is due to differential learning, Table 4 returns to the regression specification in Equation 1 but replaces the dependent variable with either the log of net speed on the first four tasks of that type, the log change in net speed relative to those first tasks, or the log of net speed excluding those first tasks. The sum of the initial and learning coefficients in columns (1) and (2) equal the total effect in column (3). Office workers are 13% more productive on day 1 of a new task type, and this difference rises another 7% over time (primarily in week 1 as seen in Figure 2) resulting in a total difference of 20%.[22] Columns (4)–(6) repeat the exercise only for easy tasks and (7)–(9) only for hard tasks, with all the learning occurring on harder tasks (for which a 19% advantage opens up on day 1, with learning accounting for a further 14% rise). The fact that learning is concentrated on hard tasks is consistent with learning curves that are steeper for more complex and difficult tasks.

4.2.3 Daily and Weekly Work Patterns

It is also interesting to explore how workers assigned to WFH take advantage of the additional flexibility and how productivity varies across the work week. In Figure 3 top left panel, we plot the proportion of work done per average weekday. For people working from home, the smallest share of work was done on Mondays, the day this group was required to visit the office to upload completed tasks and download new ones. The proportion of work done steadily rises as Monday approaches, with the highest share of work done on Sunday and Saturday. Unsurprisingly, office workers have a very stable work allocation Monday to Friday and do no work on the weekend. Looking at the allocation of work across the day in Figure 3 top-right panel, home workers start their workday a little later than office workers and spread most of it between 11 am and 10 pm. In contrast, office workers complete almost all their work between 9 am and 4 pm, with a dip around lunchtime.

How much does worker productivity vary depending on when the work is done? The bottom panels of Figure 3 plot log net speed by day of week and hour of day. The productivity of office-based workers steadily rises over the week. Home-based workers show a shallower slope between Tuesday-Saturday but are substantially less productive Sundays and Mondays. Across the workday, the productivity of office-based workers rises slightly upon arrival at the office and dips again in the afternoon. In contrast, the (lower) productivity of home-based workers remains essentially constant throughout the day, with a

[22] The -20% treatment effect is slightly different than the treatment effect reported in column (1) of Table 2, both because we exclude the first four surveys and because the observation weights are different. To ensure that columns (1) plus (2) equal column (3), here we re-weight observations such that each employee-task difficulty level has an equal weight (rather than each employee has an equal weight).

Table 4: Cumulative Learning

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Both Hard and Easy Tasks			Easy Tasks			Hard Tasks		
	Initial	Learning	Total	Initial	Learning	Total	Initial	Learning	Total
Alloc_home	-0.13**	-0.070***	-0.20***	-0.086*	-0.013	-0.10**	-0.19***	-0.14***	-0.33***
	(0.052)	(0.023)	(0.049)	(0.051)	(0.023)	(0.042)	(0.064)	(0.035)	(0.066)
Constant	3.29***	0.28***	3.58***	3.27***	0.26***	3.55***	3.40***	0.13***	3.54***
	(0.042)	(0.021)	(0.044)	(0.042)	(0.021)	(0.042)	(0.042)	(0.026)	(0.048)
Section and Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	131,923	131,923	131,923	62,251	62,251	62,251	69,672	69,672	69,672
R-squared	0.383	0.118	0.271	0.273	0.109	0.217	0.118	0.113	0.102

Notes: This table presents estimates of initial productivity differences and learning across home and office work environments. All columns regress a worker performance outcome on Alloc_home$_i$, a dummy variable that takes a value equal to one if the worker was randomly assigned to work from home. In columns (1), (4) and (7), the dependent variable is the log of initial net speed, which is defined as the average net speed for the first four surveys completed for a particular difficulty level by each worker (approximately one day's work). In columns (2), (5) and (8), the dependent variable is cumulative learning which is defined as the log change in net speed compared to the initial four surveys completed for a particular difficulty level. In columns (3), (6) and (9), the dependent variable is the log of net speed excluding the first four surveys for each difficulty level. Net speed is defined as the number of accurate characters typed per minute. Columns (1)–(3) consider both easy and hard survey tasks, whereas columns (4)–(6) and (7)–(9) consider only easy and only hard surveys, respectively. All regressions include section, week, and wave fixed effects. For all specifications, the unit of observation is a worker-task pair. All regressions are re-weighted to give a total weight of one to each worker across all observations. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Figure 3: Daily and Weekly Distribution of Work and Typing Speed

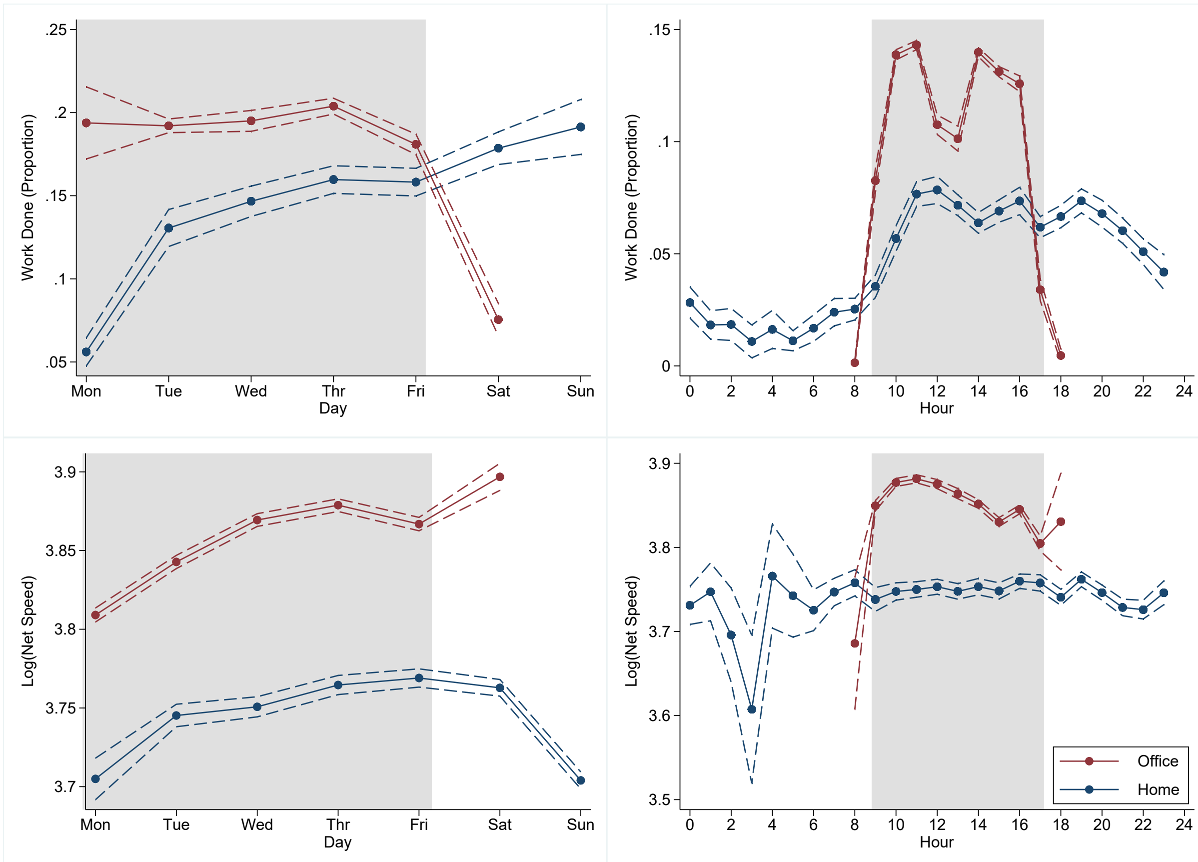

Notes: This figure plots the distribution of work done and typing speed by work location, both over days of the week and over hours of the day. The top-left and top right panels plot the proportion of work completed each day of the week and each hour of the day, respectively. The bottom-left and bottom-right panels plot the average log(net speed) over days of the week and hours of the day, respectively. The dashed lines display 95% confidence intervals.

considerable drop only observed in the middle of the night (2 am–4 am).

In addition, we check whether those who request WFH have different work patterns or are particularly productive when working outside regular office hours compared to those who do not request WFH? As reported in Appendix Table A.8, we find no support for either of these hypotheses. The time allocation of workers assigned to WFH is the same, irrelevant of whether they requested WFH or not. Similarly, those who requested WFH are no faster during evening or weekend hours compared to those who requested the office.

A related question of interest is, do workers shirk once they complete their weekly task target (at which point incentives weaken as they are guaranteed the fixed component of their salary), and does this shirking vary with workplace assignment? As shown in Appendix Table A.9, we find no evidence of such shirking, differential or not.[23]

[23] We show that productivity does not decline after reaching this target, either for those assigned to home or office locations.

4.3 Selection on Ability

Next, we turn to the question of whether workers sort into office versus home work on the basis of their innate ability. For example, if high-ability workers prefer office work because of lower costs of working in a more regulated environment, such sorting will magnify the treatment effects we found above. To investigate if higher ability workers select into office work, we regress initial worker performance on stated preferences for home work:

$$\text{Initial Worker Performance}_{i,n} = \beta \text{Pref\_home}_i + \gamma X_n + \epsilon_{i,n} \tag{2}$$

where Initial Worker Performance$_{i,n}$ is the log of net speed achieved by worker i on one of three different speed tests n that were conducted prior to beginning the job; Pref_home$_i$ is an indicator variable capturing the (incentivized) work location choice of the employee prior to being allocated to home or office and equals one if the worker preferred WFH; X_n contains wave fixed effects to account for temporal differences in cohort quality and dummies for each of the three speed tests we administered.

Table 5 presents the results of estimating Equation 2. Column (1) considers the sample of all 884 applicants who showed up for walk-in interviews and completed a speed test. Contrary to the hypothesis that there may be positive selection on ability into office work, the positive coefficient on Pref_home$_i$ indicates that applicants preferring WFH were 15% faster on the hour-long speed test conducted during the interview process (significant at the 1% level). In column (2), we restrict our sample to include only the 234 workers who moved forward to training, filtering out workers who were outside our age eligibility or were not willing to work in either location.[24] This filtering potentially removed those with the most extreme preferences for work locations. The selection effect persists in this restricted sample, although it is diminished to a 10% difference (significant at the 5 % level).

Column (3) presents our preferred selection specification that stacks the results from all three different speed tests conducted prior to the start of work. Thus, we both include the cash-incentivized test and increase precision (at the cost of focusing only on the sample of workers who started training). We find that workers preferring home are 12% faster than workers preferring office (significant at the 1% level). In sum, whether we look at the full sample of job applicants or those ultimately selected for work, we find that initially-more-productive workers are more likely to prefer WFH. This unexpected finding does not come from our filtering out of those who refused to work in one location—if anything our selection effect is stronger for that subgroup.

Columns (4) and (5) investigate whether the same selection effect is present in the performance of employees over the subsequent two months of employment, although now selection is coupled with learning on the job. To do so, we simply replace our initial worker performance measure in Equation 2 with our regular worker performance measures. Though the coefficient on Pref_home$_i$ shrinks slightly in magni-

[24] We have 235 workers in the post-retention bonus waves but are missing walk-in speed test results for one worker.

Table 5: Selection on Ability

	(1)	(2)	(3)	(4)	(5)
	Initial Log(Net Speed)			Log(Net Speed)	
	Applicants 1 Test	Workers 1 Test	Workers 3 Tests	Workers Work data	Workers Work Data
Pref_home	0.15***	0.10**	0.12***	0.084*	0.084*
	(0.025)	(0.049)	(0.033)	(0.050)	(0.048)
Alloc_home					-0.18***
					(0.049)
Constant	3.08***	3.13***	3.22***	3.32***	3.41***
	(0.023)	(0.037)	(0.032)	(0.057)	(0.062)
Speed Test FE	No	No	Yes	No	No
Wave FE	Yes	Yes	Yes	Yes	Yes
Section and Week FE	No	No	No	Yes	Yes
Observations	884	234	704	138,646	138,646
R-squared	0.089	0.040	0.148	0.248	0.264

Notes: This table contains estimates of the degree of selection based on initial ability. Columns (1)–(3) regress the log of net speed during speed tests conducted prior to the start of the employment spell on Pref_home$_i$, a dummy variable taking the value one if the worker chose home-based work when given an incentivized choice prior to the random workplace allocation (see Equation 2). Column (1) uses data from the 'walk-in' interview speed test attempted by all applicants. Column (2) additionally filters the sample of applicants to include only those who were subsequently hired and started training. Column (3) adds observations from two additional speed tests conducted during training and includes speed test fixed effects. All specifications control for wave fixed effects. Columns (4) and (5) replace initial speed with log net speed over the two months of employment and include section and week fixed effects, with column (5) further controlling for the workplace allocation Alloc_home$_i$. In columns (4) and (5), each observation is a worker-task pair and observations are re-weighted to give a weight of one to each worker. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

tude, we again find that those who prefer WFH perform better in whatever location they were assigned to (an 8.4% higher speed, significant at the 10% level). In column (5) of Table 5, we control for the allocated work location. Since the allocation of work location is randomized, it should be uncorrelated with preferences. Reassuringly, the coefficient on Pref_home$_i$ does not change.

We also explore the robustness of these selection effects. Appendix Table A.10 analyzes how the sorting on work location preference relates to other productivity measures. Overall, just like net speed, both applicants' and workers' samples show positive selection into WFH when considering gross speed, accuracy, and idle time (although the accuracy differences are small and insignificant).

Recall that, in some weeks, we posted ads highlighting home-based work opportunities and in others, office-based work opportunities. These ads may have attracted different worker types and added noise to our selection on ability results. In Appendix Table A.11, we run an identical set of specifications as Table 5 except we control for any selection effect due to the type of newspaper ad workers responded to.

On adding ad-type controls, the selection on ability effects change but not dramatically (shrinking by 3 percentage points in the full applicant sample but growing by 1–2 points in the training sample or using the work rather than speed test data).

The variation generated by the different newspaper ads also provides a second dimension of selection that we explore in Appendix Table A.11. In the applicant sample, we find that those responding to newspaper ads offering home work are 6.7% faster on the interview speed tests than those responding to office-based work ads (the same direction as the selection on self-reported WFH preferences).[25]

Taken together, we find robust evidence for negative selection effects of office work—i.e. initially-better workers are selecting into home work—not the positive selection effects that might explain the higher productivity of office- versus home-based production in observational data. In Section 5.1, we try to understand the origins of this selection by exploring how much the size of the effect attenuates with the addition of sets of observable worker characteristics that are likely to be correlated with the preferences or constraints that different groups face when selecting a work environment.

4.4 Selection on Treatment Effects

We now turn to exploring selection on treatment effects. This serves two purposes. First, it is of independent interest to understand whether applicants choose the work environments where they are (relatively) more productive and hence (relatively) better remunerated. Second, positive selection on treatment effects may lie behind the negative selection on ability documented above. For example, low-ability workers might benefit more from having peers and supervisors around them in the office or face more distractions at home, and so be more likely to choose office work than high-ability workers. To test for this possibility, we ask whether those who chose WFH (using their incentivized preference) experience less sharp declines in productivity when working from home. Specifically, we examine how the WFH treatment effect interacts with the preference for home-based work via the following regression specification:

$$\text{Worker Performance}_{i,t} = \alpha \text{Alloc\_home}_i + \beta \text{Pref\_home}_i + \lambda \text{Alloc\_home}_i * \text{Pref\_home}_i + \gamma X_{i,t} + \epsilon_{i,t} \quad (3)$$

Worker Performance$_{i,t}$ for worker i and task t is again measured by log net speed, log gross speed, accuracy, and idle time; Alloc_home$_i$ and Pref_home$_i$ are indicator variables taking the value 1 for workers randomly allocated to WFH and for workers who preferred WFH, respectively; and $X_{i,t}$ capture week, section and wave fixed effects. The coefficient λ on the interaction between the allocation to and preference for WFH captures the selection on treatment effect. If $\lambda > 0$, those who prefer WFH see their productivity fall relatively less when at home rather than the office compared to those who prefer working from the office.

[25] We do not further consider selection driven by advertisement type in part because our filtering process to select workers from the applicant sample and subsequent attrition tampers this selection substantially. In the sample of those participating in training, workers who responded to home-based work adverts are 3.6% to 8.6% slower, not faster, compared to the workers selected from office-work ads. This selection on ad type does not affect our main findings. As shown in Appendix Table A.6, we find limited heterogeneity when breaking out our main specifications by ad-type.

Table 6: Selection on Treatment Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log(Net speed)		Log(Gross speed)		Accuracy		Idle time	
Alloc_home	-0.14**	-0.13**	-0.086*	-0.055*	-1.66	-1.59	1.76*	1.74*
	(0.066)	(0.055)	(0.044)	(0.030)	(1.53)	(1.55)	(1.05)	(1.04)
Pref_home	0.15**	0.068	0.063	0.029	3.85***	3.93***	-0.32	-0.34
	(0.064)	(0.058)	(0.049)	(0.033)	(1.37)	(1.37)	(1.04)	(1.03)
Alloc_home*Pref_home	-0.12	-0.14*	-0.081	-0.13***	-2.20	-2.29	1.88	1.99
	(0.095)	(0.082)	(0.067)	(0.049)	(2.11)	(2.12)	(1.68)	(1.68)
Initial Worker Performance		0.75***		0.57***		-0.032		0.042
		(0.14)		(0.067)		(0.043)		(0.11)
Constant	3.39***	0.73	3.64***	1.39***	80.2***	82.3***	15.8***	15.6***
	(0.066)	(0.49)	(0.047)	(0.26)	(1.62)	(3.14)	(0.92)	(1.07)
Section, Week and Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	138,646	138,646	138,646	138,646	138,646	138,646	138,646	138,646
R-squared	0.265	0.335	0.299	0.449	0.335	0.336	0.080	0.080

Notes: This table presents estimates of selection on treatment effects. The specification given by Equation 3 regresses worker performance measures on dummy variables for being allocated to home work (Alloc_home$_i$), choosing home work (Pref_home$_i$), and their interaction. In columns (1) and (2), the dependent variable is the log of net speed, the number of accurate characters typed per minute. In columns (3) and (4), the dependent variable is the log of gross speed, the number of total characters typed per minute. In columns (5) and (6), the dependent variable is accuracy, the percentage of characters typed that are accurate. In columns (7) and (8), the dependent variable is idle time, the percentage of the total data entry time where there was no input on the mouse or keyboard. Initial worker performance controls in even-numbered columns are the same performance measures as the dependent variable but calculated from the performance on the cash-incentivized speed test conducted during training. Regressions include section, week, and wave fixed effects. Each observation is a worker-task pair, and observations are re-weighted to give a weight of one to each worker. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 6 presents these results. In column (1), we see that in terms of net speed, workers randomly allocated to WFH are 14% slower than those in the office, while those who prefer WFH are 15% faster than those who prefer office. However, the interaction is equal to negative 12%. Put another way, while there is a negative WFH treatment effect of 14% for those who prefer office work, those who actually prefer WFH see their performance suffer even more, falling 26% when working from home compared to an office. This negative selection on treatment effects is sizeable but not significant at conventional levels. However, when we control for initial log net speed to increase precision as we did in Table 3, we find that the negative selection on treatment effects increases to -14% and becomes statistically significant at the 10% level (column (2)). Similar negative selection on treatment effects can be observed in the case of gross speed (significant at the 1% level when controlling for initial performance, column (4)). In addition, WFH treatment effects on accuracy are more negative, and those on idle time more positive among those preferring home, although both of these effects are smaller in magnitude and insignificant.

Appendix Table A.12 looks directly at whether high-ability individuals have larger WFH treatment effects. To do so, we replace the Pref_home_i dummy in Equation 3 with our measure of initial ability, Initial Log(Net Speed). We find that initially-faster workers see larger productivity drops from WFH—i.e., higher-ability people have more to gain from office work—although the interaction is insignificant. But as we saw in Section 4.3, this group is less likely, not more, to choose office-based work, consistent with the negative selection on treatment effects above.

In sum, people who choose home-based work see a large decrease in productivity and hence remuneration when they work from home compared to the office. People who prefer office-based work also experience lower productivity at home, but the treatment effect is only about half as large. Thus, there is negative selection on treatment effects. The hypothesis that workers positively sort on treatment effects is rejected and so such a mechanism cannot be behind the negative selection on ability found in the preceding section. Relatedly, we find no evidence that high-ability workers have more to gain from WFH. Instead, the results point to a different explanation for negative selection: that some workers might be constrained from choosing their most productive work location or face other costs and benefits from doing so. We now explore such possibilities by both conditioning on worker characteristics and exploring heterogeneity along these same dimensions.

5 Exploring the Origins of the Negative Selection Effects

The prior analysis shows that office work has large positive treatment effects compared to WFH, but initially-high-ability workers are less likely, not more, to sort into office jobs. Furthermore, the workers whose productivity suffers the most from WFH are the most likely to choose it. In this section, we investigate the origins of the negative selection on both ability and treatment effects.

5.1 Constraints to Selection on Ability

We first compile hypotheses for why high-ability workers might be particularly constrained from sorting into the more-productive office environment or have stronger preferences in favor of the home environment. Next, we obtain sets of worker characteristics from our baseline survey that proxy for the omitted variable driving the relationship between selection and ability implied by each mechanism. In the final step, we include these proxies as controls in our selection regression. If a particular hypothesized mechanism is important, the correlation between initial ability and WFH preferences should attenuate considerably when the relevant proxies are added.

We explore six hypotheses that can generate negative selection into office work based on initial ability [with the corresponding proxies in square brackets]:

1. High-ability workers tend to live further away from the office and so incur higher time and effort costs commuting. [Distance between home and the office location.[26]]
2. The office serves as a commitment device for low self-control/low-productivity workers. [Agreement with the statement "I never leave things to the last minute",[27] the time discount parameter.[28]]
3. Office work is a status good for low-ability workers. [Number of previous office jobs, total monthly household income, the interaction of the two.]
4. High-ability workers continue to study or search for a better job, both of which are easier to do with a more flexible schedule. [Prefer full-time work, have additional study or job search commitments.]
5. High-ability workers have more responsibilities at home (or low-ability workers anticipate more distractions at home and so choose office). [Family care responsibilities, married, has kids, age.]
6. High-ability women face greater social sanctions or pressures to work inside the home (or low-ability women anticipate more distractions at home and so choose office). [Female; interactions of female with family care responsibilities, married, has kids, age.]

Specifically, we run the following specification:

$$\text{Initial Worker Performance}_{i,n} = \beta \text{Pref\_home}_i + \sum_h \gamma_{1,h} \text{Characteristic}_{i,h} + \gamma_2 X_n + \epsilon_{i,n} \quad (4)$$

where Initial Worker Performance$_{i,n}$ is the log speed of worker i on each of the three initial speed tests indexed by n; Pref_home$_i$ takes the value one if worker i prefers WFH; $\{\text{Characteristic}_{i,h}\}_{h=1}^{H}$ denotes the set of characteristics that proxy for hypothesis h; and X_n are fixed effects for each of the three tests.

We summarize our results in Panel A of Table 7. The first row reports the coefficient and standard

[26] Here, we are implicitly assuming that non-monetary commute costs, including the cost of time, are proportional to the travel distance. Recall that workers were compensated for incurred monetary travel costs.

[27] Workers were asked to rank various positive attributes that best describe them in a personality test.

[28] The elicitation device is Andreoni et al. (2015)'s convex time budget (CTB). CTB uses variation in linear budget constraints over early and later income to identify long-run time discounting, present bias, and utility function curvature.

Table 7: Selection Effects—Controlling for Characteristics

Panel A: Selection on Ability

		(1)	(2)	(3)	(4)
		Controls for			
		All Characteristics		**1st Principal Component**	
Regression Specification		**Pref_home**	**(SE)**	**Pref_home**	**(SE)**
Baseline		0.116***	(0.033)		
Hypothesis controlled for	Costs	0.116***	(0.033)	0.116***	(0.033)
	Low Self-Control	0.102***	(0.032)	0.101***	(0.032)
	Status	0.089***	(0.033)	0.109***	(0.033)
	Outside Options	0.108***	(0.031)	0.116***	(0.032)
	Home Responsibilities	0.084**	(0.034)	0.109***	(0.034)
	Female Constraints	0.096***	(0.032)	0.118***	(0.033)
All hypotheses controlled for		0.051*	(0.031)	0.088***	(0.032)

Panel B: Selection on Treatment Effects

		Controls for			
		All Characteristics		**1st Principal Component**	
Regression Specification		**Alloc_home* Pref_home**	**(SE)**	**Alloc_home* Pref_home**	**(SE)**
Baseline		-0.14*	(0.08)		
Hypothesis controlled for	Costs	-0.14*	(0.08)	-0.14*	(0.08)
	Low Self-Control	-0.15*	(0.08)	-0.15*	(0.08)
	Status	-0.15*	(0.09)	-0.17**	(0.08)
	Outside Options	-0.16*	(0.08)	-0.16*	(0.08)
	Home Responsibilities	-0.17*	(0.09)	-0.13	(0.08)
	Female Constraints	-0.16*	(0.08)	-0.14*	(0.08)
All hypotheses controlled for		-0.18**	(0.08)	-0.17**	(0.08)

Notes: Panels A and B report estimates of selection on initial ability and selection on treatment effects after conditioning on worker characteristics. Each row presents the estimates from two separate regressions. The dependent variable in Panel A is the log of net speed (the number of accurate characters per minute) during three speed tests conducted prior to beginning work. The dependent variable in Panel B is the log of net speed during data entry tasks performed while working. Columns (1) and (3) present the coefficient estimates on Pref_home$_i$ from running the specification in Equation 4 (Panel A) and the coefficient estimates on Pref_home$_i$ $*$ Alloc_home$_i$ from running the specification in Equation 5 (Panel B). For both Panels, columns (2) and (4) report the standard error of the estimate in the previous column. Rows describe the controls included in the Characteristic$_{i,h}$ controls. The first row presents the baseline effect when no characteristics are controlled for. The following 6 rows control for 6 sets of characteristics each representing a single hypothesis. Section 5.1 describes the characteristic variables. Columns (1) and (2) include multiple controls within a set simultaneously, columns (3)-(4) include a single control per set, the first principal component of the full set of characteristics representing a particular hypothesis. Finally, the last row of each panel includes all controls for all hypotheses simultaneously (or all first principal components of all hypotheses in columns (3) and (4)). In Panel A, the unit of observation is the worker-speed test and all regressions control for the speed test type and wave fixed effects. In Panel B, the unit of observation is a worker-task pair and all regressions include section, week, and wave fixed effects. Panel B regressions are re-weighted to give equal weight to each worker. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

errors for our baseline selection estimation that includes no characteristic controls (Table 5 column (3)). Columns (1) and (2) of subsequent rows show regressions that include the sets of controls discussed above but the table only reports the size of the selection effect, the coefficient β in Equation 4 (for completeness, Appendix Table A.13 contains the full set of coefficients). The final row includes all sets of controls concurrently. As the number of characteristics representing each hypothesis varies, it is challenging to compare β attenuation across explanations. Thus, we also conduct a principal component analysis using the first component from each hypothesis' complete set of characteristics as the control for that hypothesis. The β coefficients from these regressions are reported in columns (3)–(4).

Although the selection effects attenuate with controls, the amount of β attenuation is relatively small. The coefficient on Pref_home$_i$ shrinks the most (from 0.116 to 0.084) when we control for the four measures of home responsibilities but remains significant at the 5% level. Proxies for office work being a status good, female constraints and low self-control also attenuate the coefficient but by smaller amounts. Comparing the attenuation only using the first principal component, controls for low self-control matter most but only reduce the coefficient from 0.116 to 0.101. Thus, no single hypothesis fully explains the negative sorting into office work although there is some support for the hypothesis that better (typically female) workers have larger home responsibilities, that low-ability workers choose the office as a status good, and that low-productivity low-self control types choose the office as a commitment device.

The final row of Panel A simultaneously controls for all six hypotheses. Unsurprisingly, the attenuation is greater than with any single set of controls with the coefficient falling to 0.051 (still significant at the 10% level). However, even when including all these 17 controls, there remains substantial selection into WFH by ability that remains unexplained.

5.2 Constraints to Selection on Treatment Effects

We can perform a similar exercise to shed light on the negative selection on treatment effects and ask whether our finding comes from comparisons between groups that might face different constraints when choosing an optimal work location. We run the following specification to explore whether accounting for the same six sets of characteristics explains the selection on treatment, with variables defined as above:

$$\text{Worker Performance}_{i,t} = \tau \text{Alloc\_home}_i + \delta \text{Pref\_home}_i + \lambda \text{Alloc\_home}_i * \text{Pref\_home}_i + \sum_h \gamma_{1,h} \text{Alloc\_home}_i * \text{Characteristic}_{i,h} + \sum_h \gamma_{2,h} \text{Characteristic}_{i,h} + \gamma X_{i,t} + \epsilon_{i,t} \quad (5)$$

To understand this specification, suppose that our negative selection on treatment effects is coming from women being both more likely to prefer WFH due to cultural constraints, and less productive at home compared to the office because of home responsibilities. In this scenario, after including the interaction of a female gender dummy with Alloc_home$_i$ (and the main effects), we would no longer observe selection on treatment effects (i.e. the negative sign of the coefficient λ would attenuate towards zero).

Panel B of Table 7 presents the λ coefficients after the inclusion of the controls. The first row repeats the -14% coefficient from our baseline regression. Each subsequent row reports regressions with either one set of controls or the first principal component of those controls. Rather than the inclusion of controls attenuating the coefficient, λ becomes more negative for all six sets of hypotheses above, growing to -18% when all controls are included. Thus, heterogeneity in treatment effects by worker characteristics coupled with correlations between characteristics and WFH preferences are not behind the negative selection on treatment effects—at least for the characteristics suggested by our six hypotheses.

Note that this result does not rule out the possibility that cultural constraints on certain groups, such as women, lie behind our finding. For example, suppose men choose whether to work from home or the office for idiosyncratic reasons uncorrelated with their relative productivity across environments. However, cultural norms mean that women's choices are determined by whether they have home responsibilities or not—and if they do, they are less productive at home than in the office. In this scenario, there may be little or no attenuation on the selection on treatment effects when $\text{Alloc\_home}_i * \text{Female}_i$ is included since the treatment-effect heterogeneity is not across genders per se but across WFH preferences themselves. Such an explanation generates an ancillary prediction, that selection on treatment effects should only occur within groups constrained in this way (females in this case).

To allow selection on treatment effects to differ within groups defined by characteristics, we create an indicator variable Sub_group_i by bisecting our sample into two subgroups based on whether the value is above or below the median for every characteristic control discussed above. For example, for gender we bisect our sample into male and female, with female taking the value 1. We then interact this indicator variable with alloc_home_i, pref_home_i, and the product of the two:

$$\begin{aligned}\text{Worker Performance}_{i,t} = {} & \tau \text{Alloc\_home}_i + \delta \text{Pref\_home}_i + \lambda \text{Alloc\_home}_i * \text{Pref\_home}_i + \\ & \tau' \text{Alloc\_home}_i * \text{Sub\_group}_i + \delta' \text{Pref\_home}_i * \text{Sub\_group}_i + \\ & \lambda' \text{Alloc\_home}_i * \text{Pref\_home}_i * \text{Sub\_group}_i + \theta \text{Sub\_group}_i + \gamma X_{i,t} + \epsilon_{i,t} \quad (6)\end{aligned}$$

Thus, we now allow each subgroup to have different selection on treatment effects (λ for the subgroup for which $\text{Sub\_group}_i = 0$ and $\lambda + \lambda'$ for the subgroup for which $\text{Sub\_group}_i = 1$).

Columns (5) and (6) of Table 8 present these two coefficients, λ and λ', with each row a separate regression, one for each of the characteristics discussed at the start of this section. (The earlier columns of Table 8 explore heterogeneity in treatment effects and selection effects that we discuss in Section 6.) For a number of characteristics, selection on treatment effects occur only within one of the two subgroups. They are particularly pronounced for households with low family income, those with family care obligations (particularly women), those with children, and for older workers. In all these cases, the negative selection on treatment effects is large and highly significant for this subgroup but close to zero and, if anything, positive for the workers not in the subgroup. Compared to the baseline negative selection on treatment effects of -14%, these subgroups exhibit additional negative selection of between 37% and 80%.

Table 8: Heterogeneity in Treatment Effects, Selection on Ability, and Selection on Treatment Effects

Panel A: Baseline Regression		(1)	(2)	(3)	(4)	(5)	(6)
		Alloc_home		Pref_home		Alloc_home* Pref_home	
Baseline		-0.18***		0.12***		-0.14*	
		(0.050)		(0.033)		(0.082)	
Panel B: Regressions with Heterogeneity							
Hypothesis	Characteristic	Alloc_home	Alloc_home* Sub_group	Pref_home	Pref_home* Sub_group	Alloc_home* Pref_home	Alloc_home* Pref_home* Sub_group
Costs	Distance to Office	-0.20***	0.05	0.15***	-0.07	-0.02	-0.21
		(0.07)	(0.10)	(0.05)	(0.07)	(0.11)	(0.17)
Low Self-Control	Never Leave Last Min	-0.18**	-0.01	0.11**	0.01	-0.17	0.04
		(0.07)	(0.10)	(0.05)	(0.07)	(0.11)	(0.17)
	High Discount Rate	-0.12**	-0.11	0.05	0.14**	-0.14	0.01
		(0.06)	(0.10)	(0.04)	(0.06)	(0.10)	(0.16)
Status	Low Family Income	-0.28***	0.18*	0.09*	0.05	0.09	-0.43**
		(0.08)	(0.10)	(0.05)	(0.07)	(0.12)	(0.17)
	No Prior Office Job	-0.16***	-0.05	0.10***	0.03	-0.18*	0.11
		(0.06)	(0.10)	(0.04)	(0.08)	(0.10)	(0.18)
Outside Options	Prefer Fulltime	-0.08	-0.12	0.24**	-0.14	-0.19	0.04
		(0.13)	(0.14)	(0.10)	(0.11)	(0.17)	(0.18)
	Additional Commit	-0.16***	-0.06	0.09**	0.08	-0.17*	0.07
		(0.06)	(0.11)	(0.04)	(0.07)	(0.10)	(0.19)
Home Responsi-bilities	Family Care Needs	-0.19***	0.14	0.10***	0.13	-0.10	-0.73***
		(0.05)	(0.16)	(0.03)	(0.12)	(0.09)	(0.15)
	Married	-0.17***	-0.09	0.08**	0.14	-0.11	-0.12
		(0.06)	(0.12)	(0.04)	(0.09)	(0.09)	(0.20)
	Has Kids	-0.18***	-0.07	0.10***	0.12	-0.07	-0.48**
		(0.05)	(0.14)	(0.03)	(0.10)	(0.09)	(0.21)
	Older	-0.16***	-0.07	0.06	0.12*	-0.01	-0.37**
		(0.06)	(0.10)	(0.04)	(0.07)	(0.10)	(0.18)
Female Constraints	Female	-0.16**	-0.07	0.11**	0.02	-0.22*	0.11
		(0.07)	(0.10)	(0.04)	(0.06)	(0.13)	(0.17)
	Female*Fam Care	-0.20***	0.31*	0.11***	0.13	-0.12	-0.80***
		(0.05)	(0.17)	(0.03)	(0.16)	(0.09)	(0.16)
	Female*Married	-0.18***	-0.02	0.10***	0.16	-0.14	-0.00
		(0.05)	(0.14)	(0.03)	(0.11)	(0.09)	(0.22)
	Female*Has Kids	-0.18***	-0.07	0.10***	0.16	-0.10	-0.31
		(0.05)	(0.15)	(0.03)	(0.14)	(0.09)	(0.22)
	Female Older	-0.15***	-0.23*	0.10***	0.12	-0.13	-0.02
		(0.05)	(0.13)	(0.03)	(0.10)	(0.10)	(0.22)

Notes: This table explores heterogeneity in the treatment effect (in columns (1) and (2)), in selection on ability (in columns (3) and (4)), and in selection on treatment effects (in columns (5) and (6)). Each pair of cells presents results from a single regression. For each pair of columns, Panel A presents our baseline results where we assumed no heterogeneity. Panel B reports the coefficients on interaction of Alloc_home$_i$, Pref_home$_i$, or Alloc_home$_i$ $*$ Pref_home$_i$ with a worker characteristic dummy Sub_group$_i$ obtained by bisecting the sample by the median value of that characteristic. Section 5.1 describes the characteristic variables. For columns (1)–(2) and (5)–(6), the dependent variable is the log of net speed. For columns (3)–(4) it is the log of initial net speed during speed tests conducted as part of the interview and training process. The regressions reported in columns (1)–(2) and columns (5)–(6) include section, week, and wave fixed effects. The unit of observation is the worker-task pair and observations are re-weighted to give equal weight to each worker. The regressions reported in columns (3)–(4) include speed test type and wave fixed effects and the unit of observation is the worker-speed test pair. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistically significant at the 10%, 5%, and 1% level, respectively.

The groups within which selection on treatment is largest are certainly suggestive of societal constraints lying behind the unexpected negative selection. For example, widely varying norms and socioeconomic conditions across women with family care commitments mean that expectations regarding childcare or the acceptability of work outside the home may vary greatly even within this group. This within-group heterogeneity makes it possible that those who choose WFH are those who have the greatest non-data-entry demands on their time while at home and those who choose office work have fewer demands on them while at home. For example, the latter group may live with their mother-in-law who helps with housework while the former group do not. In contrast, we find no negative selection within groups that are free of these types of constraint, and thus face no heterogeneity in the severity of the constraint. For example, women without family obligations. The presence and importance of such heterogeneity within constrained subgroups deserves further investigation in future work.

6 Heterogeneity in Treatment and Selection Effects

Finally, we turn to studying heterogeneity in our treatment and selection effects. This serves two purposes. First, it is of independent interest. For example, whether women have higher treatment effects from working in an office than men, or poorer households compared to richer ones, is of value to policymakers interested in the functioning of labor markets. Second, just as was the case in the analysis of selection on treatment effects above, the heterogeneity we find may shed further light on the origins of the selection effects by highlighting the groups for which these effects are particularly substantial.

As with the analysis above, we bisect our sample into two subgroups, above and below the median, for each characteristic discussed at the start of Section 5.1. Columns (1) and (2) of Table 8 first considers heterogeneity in treatment effects and interacts Sub_group_i dummies with Alloc_home_i:

$$\text{Worker Performance}_{i,t} = \alpha \text{Alloc\_home}_i + \alpha' \text{Alloc\_home}_i * \text{Sub\_group}_i + \theta \text{Sub\_group}_i + \gamma X_{i,t} + \epsilon_{i,t} \tag{7}$$

Again, each row is a separate regression. The coefficient α' on the interaction between Alloc_home and Sub_group provides an estimate of treatment heterogeneity by subgroup.

We find limited evidence for heterogeneity in the size of the treatment effect.[29] Only three characteristics reveal heterogeneity that is significant at the 10% level. WFH treatment effects are 18 p.p. less negative for low income households, 31 p.p. less negative for female workers with family care responsibilities, and 23 p.p. more negative for older female workers. Which subgroups have the largest and smallest treatment effects from being allocated to WFH? The subgroups with the most negative WFH treatment effects are older female workers, workers with high family income, and married workers, with treatment effects of -38%, -28%, and -26%, respectively (all significantly different from zero at the 5% level, even if the difference

[29] Recall from the discussion in Section 4.4, Appendix Table A.12 also explores heterogeneity in treatment effects but by initial ability and finds no significant heterogeneity.

between the subgroup and its complement is not). Female workers with family care responsibilities are the only group that is more productive at home than at the office (by 11%) but this treatment effect is not significantly different from zero. From the earlier discussion of heterogeneity in selection on treatment effects in columns (5) and (6), the women in this group who experience these positive treatment effects of WFH are disproportionately those who choose to work in the office, and those choosing to WFH instead experience large negative treatment effects of WFH.

Finally, we repeat the specification exploring worker selection on initial ability but interact the indicator variable for whether a worker prefers WFH with subgroup indicator variables:

$$\text{Initial Worker Performance}_{i,n} = \alpha \text{Pref\_home}_i + \alpha' \text{Pref\_home}_i * \text{Sub\_group}_i + \theta \text{Sub\_group}_i + \gamma X_{i,t} + \epsilon_{i,t} \quad (8)$$

Columns (3) and (4) of Table 8 report these results. Recall that workers who prefer WFH are 12% faster than ones who prefer the office (repeated in row (1)). There is relatively little heterogeneity in the size of this coefficient with only 2 of the 16 α' coefficients significantly different from zero (although with 234 workers, estimates are noisy). For example, the first row of the Female Constraints section of Panel B considers whether selection on ability varies by gender. Men who prefer WFH are 11% more productive prior to starting the job than those who prefer office, while women who prefer WFH are 13% more productive, with the 2 p.p. differential not significant.

The two characteristics with significantly higher selection on initial ability are applicants with a high discount parameter and older workers. In both cases, those who prefer WFH are almost 20% more productive than those who do not. There is also much greater selection for those who prefer part time work, have family care responsibilities, have children, or are married (particularly for women in the last three cases). But these sizable differences are not statistically different from zero. These patterns do, however, complement the finding above that the selection on treatment effects were particularly large within groups who are often constrained in the labor market choices they can make.

7 Conclusions

We set up a randomized control trial that allocates workers to home- or office-based work while holding all other dimensions of the work constant. Our first finding is a large and negative WFH treatment effect of -18%. Two-thirds of the effect exists from the first day of work and the rest is due to quicker learning by office workers over the subsequent weeks. Second, we find negative selection on ability into office based work. Those who prefer home-based work are 12% more productive at baseline. Finally, we find negative selection on treatment effects: workers who prefer WFH have larger negative productivity effects when allocated to home than those who preferred office. This misallocation of workers away from their most productive work environment reduces aggregate labor productivity although welfare impacts depend on whether these choices are the result of cultural or personal preferences rather than external constraints.

We find somewhat limited evidence that these negative selection effects are due to workers with different characteristics, such as those with high discount rates or those living far from the office, facing different constraints. However, the negative selection on ability and negative selection on treatment effects are larger within subgroups that typically face bigger constraints on working outside the home, such as those with children or other home care responsibilities, particularly women, as well as for poorer households. For example, we would see such patterns if the women who struggle to find substitutes for their home labor both prefer to WFH and are most interrupted by demands for home labor when they do so. This heterogeneity that we find within constrained subgroups requires further investigation.

Our results add experimental evidence to the burgeoning literature showing the potential negative productivity impacts of WFH. The finding that those with the most negative treatment effects are more, not less, likely to select WFH shows that the self-selection of workers into different work locations is of first-order importance when evaluating the merits of policies that aim to alter the allocation of workers to different work environments. Our results can also help evaluate the productivity impacts of industrial policies that are not directly aimed at changing constraints to or selection into WFH, but change the availability of office- and factory-based jobs versus home-based ones (e.g., support for microenterprises).

References

Andreoni, James, Michael A Kuhn, and Charles Sprenger, "Measuring time preferences: A comparison of experimental methods," *Journal of Economic Behavior & Organization*, 2015, *116*, 451–464.

Angelici, Marta and Paola Profeta, "Smart working: work flexibility without constraints," *Management Science*, 2023.

Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis, "Why Working from Home Will Stick," Working Paper 28731, National Bureau of Economic Research April 2021.

Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta, "Cross-Country Differences in Productivity: The Role of Allocation and Selection," *American Economic Review*, February 2013, *103* (1), 305–34.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts, "Does management matter? Evidence from India," *The Quarterly Journal of Economics*, 2013, *128* (1), 1–51.

_ , James Liang, John Roberts, and Zhichun Jenny Ying, " Does Working from Home Work? Evidence from a Chinese Experiment *," *The Quarterly Journal of Economics*, 02 2015, *130* (1), 165–218.

_ , Ruobing Han, and James Liang, "How Hybrid Working From Home Works Out," Working Paper 30292, National Bureau of Economic Research July 2022.

Boltz, Marie, Bart Cockx, Ana Maria Diaz, and Luz Magdalena Salas, "How does working-time flexibility affect workers' productivity in a routine job? Evidence from a field experiment," *British Journal of Industrial Relations*, 2023, *61* (1), 159–187.

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," Working Paper 27344, National Bureau of Economic Research June 2020.

Choudhury, Prithwiraj, Tarun Khanna, Christos Makridis, and Kyle Schirmann, "Is hybrid work the best of both worlds? Evidence from a field experiment," *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, 2022, (22-063), 22–063.

Clark, Gregory, "Factory Discipline," *The Journal of Economic History*, 1994, *54* (1), 128–163.

Emanuel, Natalia and Emma Harrington, "Working remotely? Selection, treatment, and the market for remote work," *Selection, Treatment, and the Market for Remote Work (June 2023). FRB of New York Staff Report*, 2023, (1061).

Forquesato, Pedro, "Social norms of work ethic and incentives in organizations," *Journal of Economic Behavior Organization*, 2016, *128*, 231–250.

Gibbs, Michael, Friederike Mengel, and Christoph Siemroth, "Work from Home and Productivity: Evidence from Personnel and Analytics Data on Information Technology Professionals," *Journal of Political Economy Microeconomics*, 2023, *1* (1), 7–41.

Hiller, Victor, "Work organization, preferences dynamics and the industrialization process," *European Economic Review*, 2011, *55* (7), 1007–1025.

— , "Self-control and the rise and fall of factory discipline," *Journal of Development Economics*, 2018, *133*, 187–200.

Ho, Lisa, Suhani Jalota, and Anahita Karandikar, "Bringing Work Home: Flexible Arrangements as Gateway Jobs for Women in West Bengal," Technical Report, MIT Working Paper 2023.

Hsieh, Chang-Tai and Peter J Klenow, "Misallocation and manufacturing TFP in China and India," *The Quarterly journal of economics*, 2009, *124* (4), 1403–1448.

— **and** — , "The life cycle of plants in India and Mexico," *The Quarterly Journal of Economics*, 2014, *129* (3), 1035–1084.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan, "Self-control and the development of work arrangements," *American Economic Review*, 2010, *100* (2), 624–28.

— **,** — **, and** — , "Self-control at work," *Journal of Political Economy*, 2015, *123* (6), 1227–1277.

Kim, Hyuncheol Bryant, Hyunseob Kim, and John Y Zhu, "The Selection Effects of Part-Time Work: Experimental Evidence from a Large-Scale Recruitment Drive," 2022.

Künn, Steffen, Christian Seel, and Dainis Zegners, "Cognitive Performance in Remote Work: Evidence from Professional Chess," *The Economic Journal*, 2022, *132* (643), 1218–1232.

McCaig, Brian and Nina Pavcnik, "Export markets and labor allocation in a low-income country," *American Economic Review*, 2018, *108* (7), 1899–1941.

Monte, Ferdinando, Charly Porcher, and Esteban Rossi-Hansberg, "Remote Work and City Structure," Working Paper 31494, National Bureau of Economic Research July 2023.

Monteiro, Natalia P, Odd Rune Straume, and Marieta Valente, "When does remote electronic access (not) boost productivity? Longitudinal evidence from Portugal," *Information Economics and Policy*, 2021, *56*, 100923.

Morikawa, Masayuki, "Productivity dynamics of remote work during the COVID-19 pandemic," *Industrial Relations: A Journal of Economy and Society*, 2023.

Porta, Rafael La and Andrei Shleifer, "Informality and development," *Journal of economic perspectives*, 2014, *28* (3), 109–26.

Shen, Lucas, "Does working from home work? A natural experiment from lockdowns," *European Economic Review*, 2023, *151*, 104323.

Yang, Longqi, David Holtz, Sonia Jaffe, Siddharth Suri, Shilpi Sinha, Jeffrey Weston, Connor Joyce, Neha Shah, Kevin Sherman, Brent Hecht et al., "The effects of remote work on collaboration among information workers," *Nature human behaviour*, 2022, *6* (1), 43–54.

Appendix to: Working from Home, Worker Sorting and Development

David Atkin, Antoinette Schoar, and Sumit Shinde

Figure A.1: Pictures of the Office and Home Work Settings

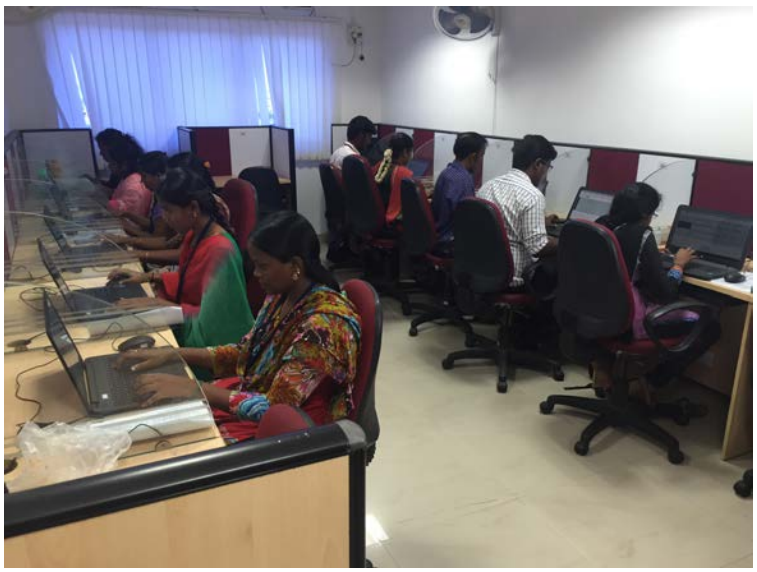

(a) The Office

(b) Home work setups

Figure A.2: User Interface of Sample Data Entry Tasks in the Proprietary Software

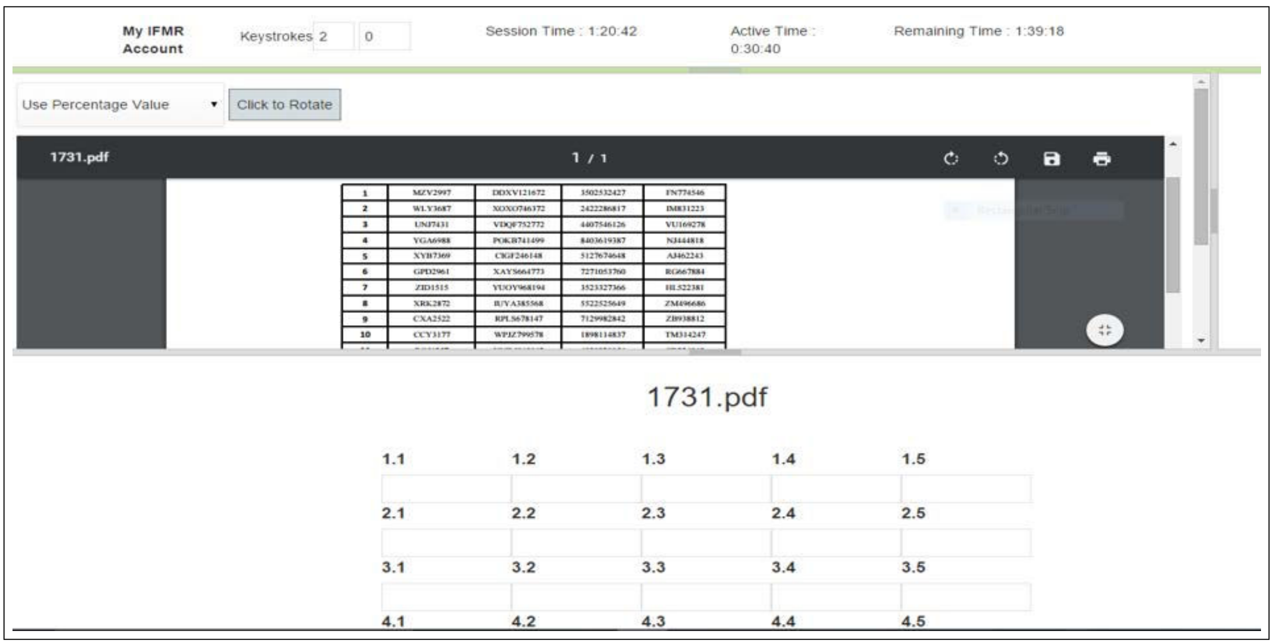

Figure A.3: Examples of Data Entry Tasks by Difficulty

(a) strings of random alpha-numeric characters vs alpha-numeric and special characters

Easy task:

S.no வ. எண்	Information 1 தகவல் 1	Information 2 தகவல் 2	Information 3 தகவல் 3	Information 4 தகவல் 4	Information 5 தகவல் 5
1	EQY6267	UGKI733669	8981753224	WM578562	OZ441532
2	DHQ1499	TSUA974617	4773422856	QD647325	NV663391
3	YAN8395	YJVW199368	6553632731	CW344523	UC189451
4	SQN6386	ZNCQ587129	3070840773	KW478175	XG635848
5	HQT4833	LYHS997811	2157713174	CN687268	LY694874

Hard task:

S.no வ. எண்	Information 1 தகவல் 1	Information 2 தகவல் 2	Information 3 தகவல் 3	Information 4 தகவல் 4	Information 4 தகவல் 4
1	?z j?G~L_	oaFeDc-,lg	4:bcw AsBoe	-r*/o~n8	)ekO
2	X tw +DfoB	b!`x #*Rk,s	VLW eNoCArM	|s.,J@=u	8X Z -o
3	5,~;honQ	k4TF|"?y#[	%~4w q?5@:t	1.$}d"^M	x 2tY9
4	]nX ^x %1m	\#*vz Z .#no	YbBm+44P35	i|{")=cY	4G*l
5	tNP$k#|C	2-/)6aCn)o	mw 9gOe[IC`	h$K~DUQr	ps;M:

(b) Type-set vs Handwritten text

Easy task:

their mother tongue and unsure in the official language. To remedy the situation we need a radically new approach to the teaching of languages. It is essential that children are taught only in their mother tongue and simultaneously learn Hindi up to grade six. This will give them the necessary grounding in their own milieu, their own folklore, mythology and literature, and help them develop a love and respect

Hard task:

mashindano. ambayo wewe kushinda au
kupoteza. kama unafikiri maishani mchezo.
kwa hiyo ni pia ni muhimu kuuliza ni aina gani
ya mchezo. Baadhi ya michezo ni alicheza kwa
ajili ya kujifurahisha peke yake. Baadhi ya michezo
ni distinctively kijamii (daraja): baadhi ni makusudi

Figure A.4: Newspaper Ads Examples

(a) Office-based work ad

(b) Home-based work ad

Figure A.5: Attrition Before and After Retention bonus

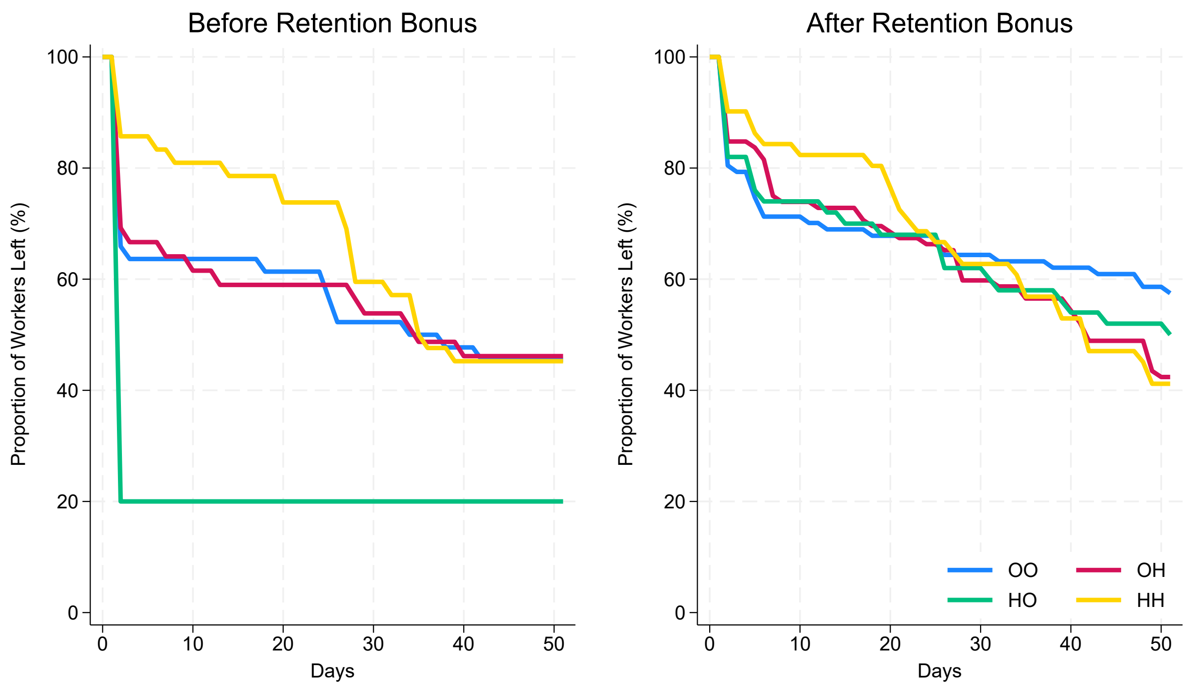

Notes: This figure plots the proportion of workers continuing to carry out data entry work against the number of days since the start of training. Left panel shows attrition prior to the introduction of a retention bonus paid at the end of the first week, right panel show attrition after. Each plot shows attrition separately for four worker groups. OO represents the workers who preferred office and were assigned office. HO represents the workers who preferred home but were assigned office. OH represents the workers who preferred office but were assigned home. Finally, HH represents workers who preferred home and were assigned home.

Table A.1: Compensation Structure

(1)	(2)	(3)	(4)	(5)
Week	Fixed component		Performance-based variable component	Retention Bonus
	Tasks Target	Amount Paid INR ($)	INR/task ($ / task)	INR ($)
1	18	2125 (32.2)	65 (1)	2000 (30.3)
2	20	2125	65	0
3	24	2125	65	0
4	24	2125	65	0
5-8	26	2125	65	0

Notes: This table explains the compensation structure for workers in both work locations. Each row indicates the compensation structure for a particular week. The weeks are displayed in column (1). Columns (2) and (3) display the fixed component of the compensation structure. Upon completing the number of tasks listed in column (2), workers were paid the amount listed in column (3). Column (4) lists the performance-based pay which paid a piece rate per task completed beyond the weekly task target. Finally, column (5) displays the retention bonus that was paid at the end of week 1. Figures in parenthesis are amounts in dollars at the exchange rate of INR 66 $\approx$ \$ 1.

Table A.2: Compensation Penalty for Errors

Penalty	Easy Task	Hard Task
	Error rate between (%)	
1X	0 - 7.5	0 - 15
1.5X	7.5-10	15-20
2X	10+	20+

Notes: This table explains the penalty schedule imposed for various levels of error rates.

Table A.3: Changes Made Across Study Waves

Detail	**Wave 1** Mar'16 - Jun'16	**Wave 2** Aug'16- Nov'16	**Wave 3** Dec '16- Feb '17	**Wave 4** Jan '17- Nov '17	**Wave 5** Oct '17- Apr '18
# Applicants **# Workers**	221 30	263 50	154 25	339 121	553 104
Hiring Batch	One per wave			Multiple overlapping per wave	
Work Duration	3 months		2 months		
Salary- Fixed component	None	• Month 1 (M1): INR 8500 ($ 128.8) • M2: INR 4400 ($ 66.7) • M3: None	• M1 : INR 8500 • M2 : INR 8500		
Variable component	Paid 6 paisa per 4 correct characters ($1 per 4000 correct characters)	INR 65 ($ 1)/ DE task completed above the given target of surveys in each week			
Completion bonus	None			INR 2000 on completion of 1st week of work	• High Incentive: INR 2000 on completion of 1st week of work • Low Incentive: INR 2000 on completion of 8th (last) week of work
Targets to retain the job	• Time: 40 hours			• Time: 35 hours (weeks 1-4) • Tasks: 26 surveys (weeks 5-8)	
Selection Criterion	• Age: 18 to 40 years • Education: 9th grade to graduates • Speed: 10-30 words per minute • Residence: Within Chennai • DE work ex: 0-6 months	• Age: 18 to 40 years • Willingness to work in both locations (surveyor assessment) • Time commitment: Full time		• Age: 18 to 40 years • Willingness to work in both locations (Manager assessment)	

Cells left blank imply that no changes to the previous setting were made. We use the average exchange rate between Indian Rupee and United States Dollar during the period of the experiment which is INR 66 $\approx$ $ 1.

Table A.4: Attrition—Pre and Post Retention Bonus

	(1) Pre	(2) Post	(3) Pre	(4) Post
Dependent Variable	Days Worked		Worked (yes)	
1{Preferred Home, Allocated Office}	-27.0***	-2.98	-0.68***	-0.082
	(1.96)	(6.61)	(0.045)	(0.17)
1{Preferred Office, Allocated Home}	-7.11	-1.88	-0.20	-0.054
	(5.62)	(0.88)	(0.068)	(0.047)
1{Preferred Office, Allocated Office}	-6.21	-1.33	-0.22	-0.097
	(7.43)	(4.40)	(0.16)	(0.12)
Constant	34.2***	38.3**	0.82***	0.90*
	(1.12)	(2.82)	(0.049)	(0.083)
Wave FE	Yes	Yes	Yes	Yes
Observations	175	280	175	280

Notes: This table regresses two measures of attrition, the number of days worked (columns (1) and (2)) and a binary variable taking value one if the worker started work after being offered the job (columns (3) and (4)), on membership of the four intervention groups (with the preferred home allocated home group being the omitted baseline group). Regressions are run separately for the sample of workers who were provided the retention bonus and those who were not. Columns (1) and (3) present results for pre-bonus waves where the issue of high and differential attrition existed. Columns (2) and (4) present results for post-bonus waves where these issues were resolved by providing workers with first-week completion bonuses. Standard errors (in parentheses) are clustered at the wave level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively. For all specifications, the unit of observation is a worker.

Table A.5: Attrition—Dependence of Days Worked on Ad Type, Location Preference, and Location Allocation

	(1) daysworked	(2) daysworked	(3) daysworked	(4) daysworked
Ad_home	-3.29	-3.21		
	(2.81)	(2.78)		
Pref_home	0.58		0.14	
	(2.84)		(2.82)	
Alloc_home	0.71			0.72
	(2.71)			(2.70)
Constant	37.6***	38.2***	36.7***	36.3***
	(2.76)	(2.21)	(2.07)	(2.28)
Wave FE	Yes	Yes	Yes	Yes
Observations	280	280	280	280
R-squared	0.011	0.010	0.006	0.006

Notes: This table presents the result of the number of days worked regressed on the type of ad workers responded to, their preference of work location, and their assigned work location. The dependent variable across all regressions is the number of days worked. Variable ad_home is a dummy variable taking a value equal to one when the worker responded to a home-based work ad or takes a value equal to zero. Variable pref_home is a dummy variable taking a value equal to one when the worker requested to work from home and zero otherwise. Variable alloc_home takes a value equal to one when the work is randomly assigned to work from home and zero when randomly assigned to work from the office. Standard errors (in parentheses) are clustered at the wave level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively. For all specifications, the unit of observation is a worker.

Table A.6: Treatment Effects By Ad-Type

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sample		All			Home Ads			Office Ads	
Effect	TE	SAB	SOT	TE	SAB	SOT	TE	SAB	SOT
Dependent Variable		Log(Net Speed)			Log(Net Speed)			Log(Net Speed)	
Alloc_home	-0.18***		-0.13**	-0.19**		-0.17**	-0.17**		-0.11
	(0.050)		(0.055)	(0.072)		(0.078)	(0.065)		(0.071)
Pref_home		0.12***	0.067		0.16***	0.088		0.089**	0.069
		(0.033)	(0.058)		(0.059)	(0.078)		(0.038)	(0.078)
Pref_home*Alloc_home			-0.14*			-0.17			-0.099
			(0.082)			(0.11)			(0.11)
Constant	3.45***	3.22***	0.73	3.32***	3.12***	0.62*	3.55***	3.28***	0.88
	(0.057)	(0.032)	(0.49)	(0.083)	(0.051)	(0.33)	(0.074)	(0.041)	(0.83)
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Test FE		Yes			Yes			Yes	
Section and Week FE	Yes		Yes	Yes		Yes	Yes		Yes
Observations	138,646	704	138,646	47,253	269	47,253	91,393	435	91,393
R-squared	0.260	0.148	0.335	0.272	0.165	0.358	0.264	0.163	0.328

Notes: This table presents the paper's main results for sub-samples split by ad types. Columns (1)-(3) present the main results for the entire workers sample, whereas columns (4)-(6) and (7)-(9) present the same results for home- and office-based work ads, respectively. Columns (1), (4), and (7) present the treatment effect regressions for the three samples and include section, week, and wave fixed effects. Columns (2), (5), and (8) present the regression estimating the sorting at baseline effect and include speed test type and wave fixed effects. Columns (3), (6) and (9) present the selection on treatment effect regressions and include section, week, and wave fixed effects. Variable pref_home is a dummy variable taking value equal to one when the worker requested to work from home and is zero otherwise. Variable alloc_home takes value equal to one when the work is randomly assigned to work from home and is equal to zero when the worker is randomly assigned to work from office. All regressions are based on eight weeks of work data except the ones in columns (2), (5), and (8), which are based on the 3-speed test conducted for each worker. In columns (2), (5), and (8), the unit of observation is the test attempted. In columns (1),(3), (4), (6), (7) and (9), the unit of observation is the survey task attempted, and observations are re-weighted to give equal weight to each worker. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A.7: Treatment and Selection Effects for All Waves

	(1)	(2)	(3)	(4)	(5)	(6)
Wave	Post-Retention Bonus			All Waves		
Effect	TE	SAB	SOT	TE	SAB	SOT
Dependent Variable	Log(Net Speed)			Log(Net Speed)		
Alloc_home	-0.18***		-0.13**	-0.18***		-0.12***
	(0.050)		(0.055)	(0.042)		(0.045)
Pref_home		0.12***	0.068		0.038	0.059
		(0.033)	(0.058)		(0.034)	(0.050)
Pref_home*Alloc_home			-0.14*			-0.11
			(0.082)			(0.071)
Constant	3.45***	3.22***	0.73	3.40***	3.34***	0.36
	(0.057)	(0.032)	(0.49)	(0.081)	(0.064)	(0.39)
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes
Test FE		Yes			Yes	
Section and Week FE	Yes		Yes	Yes		Yes
Observations	138,646	704	138,646	212,823	986	212,823
R-squared	0.260	0.148	0.335	0.268	0.225	0.382

Notes: This table presents the main results of the paper replicated for all waves. Columns (1)-(3) present the main results for the post-retention bonus sample whereas columns (4)-(6) present the same results for all the waves. Columns (1) and (4) present the treatment effect regressions for the two samples and include section, week, and wave-fixed effects. Columns (2) and (5) present the regression estimating selection at baseline effect and include speed test type and wave fixed effects. Columns (3) and (6) present the selection on treatment effect regressions and include section, week, and wave fixed effects. Variable pref_home is a dummy variable taking a value equal to one when the worker requested to work from home and zero otherwise. Variable alloc_home takes a value equal to one when the work is randomly assigned to work from home and zero when randomly assigned to work from the office. In columns (2) and (5), the unit of observation is the test attempted. In columns (1),(3), (4), and (6), the unit of observation is the survey task attempted and observations are re-weighted to give equal weight to each worker. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A.8: Utilization of the Flexibility of WFH

	(1)	(2)	(3)	(4)	(5)
	Proportion of work completed				
	outside regular office hours	during weekend	Log(Net Speed)		
Pref_home	-0.0029	0.013	-0.17***	-0.17***	-0.17***
	(0.038)	(0.034)	(0.025)	(0.023)	(0.025)
Pref_home*1{Outside Office Hours}			0.0024		0.024
			(0.020)		(0.019)
Pref_home*1{Weekend}				-0.019	-0.035
				(0.023)	(0.024)
1{Outside Office Hours}			0.020		-0.017
			(0.014)		(0.012)
1{Weekend}				0.049***	0.060***
				(0.018)	(0.018)
Constant	0.57***	0.31***	4.25***	4.26***	4.26***
	(0.034)	(0.034)	(0.031)	(0.030)	(0.031)
Week and Wave FE	Yes	Yes	Yes	Yes	Yes
Worker and Section FE			Yes	Yes	Yes
Observations	738	738	64,214	64,214	64,214
R-squared	0.037	0.051	0.532	0.532	0.532

Notes: This table analyzes work patterns among those assigned to WFH, comparing those who prefer WFH those who prefer office work. The dependent variables in columns (1) and (2) are the proportion of work done outside regular office hours (i.e. outside 9am–6pm, Monday to Friday) and during the weekend, respectively. In columns (3)–(5), the dependent variable is the log of net speed, the number of accurate characters per minute. The independent variable, Pref_home, is a binary variable representing workers' preference for work location taking the value one if the choice is home-based work and zero if the choice is office-based work. Independent variables, 1{Outside Office Hours} and 1{Weekend} take value equal to one if the task was completed outside regular office hours or during weekends, respectively. All regressions account for variation arising from week of employment and the cohort of workers using week and wave fixed effects, respectively. Additionally, regressions in columns (3)–(5) account for variation arising from the type and difficulty of the attempted survey section and the worker attempting the data entry task using section and worker fixed effects. In columns (1) and (2) the unit of observation is the worker-week. In columns (3)-(5), the unit of observation is the worker-survey task pair. Despite observations being at the worker-task or -week levels, all regressions are re-weighted to give a total weight of 1 to each worker across all observations. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A.9: Work Performance after Completion of Weekly Targets

	(1)	(2)
	log(Net Speed)	
1{Above Target}	-0.0031	0.0011
	(0.014)	(0.016)
Alloc_home*1{Above Target}		-0.010
		(0.014)
Alloc_home		0.56***
		(0.017)
Constant	3.90***	3.34***
	(0.019)	(0.025)
Section and Week FE	Yes	Yes
Attempt Sequence FE	Yes	Yes
Worker FE	Yes	Yes
Observations	138,646	138,646
R-squared	0.545	0.545

Notes: This table analyzes whether workers shirk once they meet the target of the weekly task that qualifies them to receive the fixed component of the salary. The dependent variable for all specifications is the log of net speed, the number of accurate characters per minute. The independent variable, 1{Above Target} takes a value equal to one if the task was completed after the weekly tasks target was met. The independent variable, Alloc_home, is a binary variable representing workers' work location taking the value one if the worker was randomly allocated to home-based work and zero if assigned to office-based work. All regressions account for variation arising from the week of employment, the type and difficulty of survey section being attempted, learning within each week, and the worker attempting the data entry task using week, section, attempt sequence, and worker fixed effects, respectively. Both columns (1) and (2), consider the sample of all workers for all weeks and the unit of observation is the worker-survey task pair. Despite observations being at the worker-task level, all regressions are re-weighted to give a total weight of 1 to each worker across all observations. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A.10: Selection on Initial Ability—Other Outcome Measures

Sample	(1) Applicants net speed	(2) Workers net speed	(3) Applicants gross speed	(4) Workers gross speed	(5) Applicants accuracy	(6) Workers accuracy	(7) Applicants idle time	(8) Workers idle time
Pref_home	0.15***	0.10**	0.14***	0.095*	0.71	0.49	-1.96***	-1.57*
	(0.025)	(0.049)	(0.028)	(0.052)	(0.87)	(1.67)	(0.42)	(0.82)
Constant	3.08***	3.13***	3.61***	3.62***	60.6***	62.3***	14.5***	13.9***
	(0.023)	(0.037)	(0.026)	(0.039)	(0.83)	(1.24)	(0.40)	(0.61)
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	884	234	884	234	884	234	884	234
R-squared	0.089	0.040	0.037	0.018	0.045	0.025	0.045	0.044

Notes: This table contains estimates of the degree of sorting based on initial ability using additional outcome measures. The regression specification for all columns is given by Equation 2. In columns (1) & (2), the dependent variable is the log of net speed, the number of accurate characters typed per minute (these results were earlier presented in Table 5 in columns (1) and (2), respectively). In columns (3) and (4), the dependent variable is the log of gross speed, the number of total characters typed per minute. In columns (5) and (6), the dependent variable is accuracy, the ratio of accurate characters typed to total characters typed in percentage terms. In columns (7) and (8), the dependent variable is idle time, the ratio of the total time spent not moving the mouse or keyboard to the total time spent entering data. Columns (1), (3), (5) and (7) use data from speed tests attempted by all applicants. Columns (2), (4), (6) and (8) restrict the sample to include only workers who started working for us. pref_home is a binary variable representing workers' choice of work location taking the value one if the choice is home-based work and zero if the choice is office-based work. Each observation in these regressions is a worker survey pair. All specification control for wave fixed effects. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A.11: Selection on Ability—Controlling for Ad Type

	(1) Pre-Filter 1 test	(2) Post-filter 1 test	(3) Post-filter 3 test	(4) Post-filter Work data	(5) Post-filter Work data
Pref_home	0.12***	0.11**	0.12***	0.10**	0.10**
	(0.028)	(0.049)	(0.033)	(0.049)	(0.047)
Alloc_home					-0.18***
					(0.049)
Ad_home	0.067**	-0.086*	-0.036	-0.16***	-0.15***
	(0.028)	(0.049)	(0.035)	(0.050)	(0.049)
Constant	3.06***	3.17***	3.23***	3.40***	3.49***
	(0.025)	(0.042)	(0.033)	(0.060)	(0.065)
Speed Test FE			Yes		
Wave FE	Yes	Yes	Yes	Yes	Yes
Section and Week FE				Yes	Yes
Observations	884	234	704	138,646	138,646
R-squared	0.095	0.052	0.151	0.258	0.274

This table contains estimates of the degree of sorting based on initial ability, controlling for the type of advertisement workers responded to. In all columns, the dependent variable is the log of net speed, the number of accurate characters per minute. The main dependent variable, pref_home, is a binary variable representing workers' preference for work location taking the value one if the choice is home-based work and zero if the choice is office-based work. alloc_home is a binary variable representing the treatment and takes a value equal to one if the worker was randomly assigned to work from home and zero if assigned to work in the office. ad_home is a binary variable taking value one if the worker responded to employment advertising home-based jobs and zero if responded to office-based jobs. Column (3) includes Speed Test fixed effects to account for each of the three specific typing speed test performed by workers prior to beginning work. Column (1) uses data from the speed tests attempted by all applicants who showed up for walk-in interviews. Column (2) filters the sample of applicants to include only workers selected to start working for us and turned up on first day of the job. Column (3) adds observations from two additional tests performed by hired workers. The regression specification for columns (1) to (3) is given by Equation 2. Regressions (4) and (5) consider log net speed over two months of employment and further include section and week fixed effects. All specification control for wave fixed effects. Each observation in these regressions is a worker survey pair and observations are re-weighted to give a weight of one to each worker. The regression specification for columns (4) and (5) is given by Equation **??**. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A.12: Heterogeneity in Treatment Effects with Initial Ability

	(1)	(2)	(3)	(4)
	Baseline	Speed Test		
		Cash Incentive	No Incentive	Walk-in
Alloc_home	-0.18***	0.46	-0.51	0.39
	(0.050)	(0.79)	(0.52)	(0.45)
Alloc_home*Initial Log(Net Speed)		-0.19	0.091	-0.18
		(0.22)	(0.15)	(0.14)
Initial Log(Net Speed)		0.84***	0.84***	0.71***
		(0.090)	(0.081)	(0.076)
Constant	3.45***	0.44	0.57**	1.20***
	(0.057)	(0.33)	(0.29)	(0.26)
Section, Week and Wave FE	Yes	Yes	Yes	Yes
Observations	138,646	138,646	138,646	137,429
R-squared	0.260	0.334	0.359	0.354

Notes: This table presents the heterogeneity in the treatment effect of allocating workers to home-based work environments by initial ability. Across all specifications, the dependent variable is the log of net speed. Variable Alloc_home takes a value equal to one when the worker is randomly assigned to work from home and zero when randomly assigned to work from the office. Column (1) presents the baseline regression of the treatment effect. Subsequent columns present the heterogeneity in treatment effect based on initial ability measured by three different initial speed tests: namely—column (2) considers speed that was incentivized through cash payments based on performance, column (3) considers speed from an identical test with no such incentive, and column (4) considers speed from a test conducted during the initial walk-in interview. All regressions account for variation arising from the type and difficulty of the attempted survey section, the week of employment, and the cohort of workers using section, week, and wave fixed effects, respectively. All regressions are based on eight weeks of work data with the unit of observation being the individual survey task pair. Despite observations being at the survey tasks level, all regressions are re-weighted to give a total weight of 1 to each worker across all observations. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A.13: Selection on Initial Ability—Controlling for Characteristics

Panel A: Regressions with ATE

	Regression Specification		(1) Pref_Home	(2) (SE)	(3)	(4)	(5)
	Baseline		0.116***	(0.033)			

Panel B: Controlling for individual hypothesis

	Hypothesis	PCA	Pref_Home	(SE)	Characteristics	Pref_Home	(SE)
(1)	Costs	1st PC	0.116***	(0.033)	Distance to Office	0.116***	(0.033)
(2)					Never Leave Last Min	0.105***	(0.03)
(3)	Low Self-Control				Time Discount Rate	0.110***	(0.03)
(4)		1st PC	0.101***	(0.032)	All characteristics	0.102***	(0.032)
(5)					Num. Prior Office Jobs	0.091***	(0.03)
(6)	Status				Family Income	0.115***	(0.03)
(7)		1st PC	0.109***	(0.033)	All characteristics	0.089***	(0.032)
(8)					Prefer Fulltime	0.107***	(0.03)
(9)	Outside Options				Additional Commit	0.118***	(0.03)
(10)		1st PC	0.116***	(0.032)	All characteristics	0.108***	(0.031)
(11)					Fam_care	0.110***	(0.03)
(12)					Married	0.113***	(0.03)
(13)	Home Responsibilities				Has Kids	0.117***	(0.03)
(14)					Age	0.103***	(0.03)
(15)		1st PC	0.109***	(0.034)	All characteristics	0.084**	(0.034)
(16)					Female	0.115***	(0.03)
(17)					Female*Fam Care	0.112***	(0.03)
(18)	Female Constraints				Female*Married	0.120***	(0.03)
(19)					Female*Has Kids	0.117***	(0.03)
(20)					Female*Age	0.117***	(0.03)
(21)		1st PC	0.118***	(0.033)	All characteristics	0.096***	(0.032)

Panel C: Controlling for all hypothesis

	Regression Specification	PCA	Pref_Home	(SE)	Characteristics	Pref_Home	(SE)
	Control for all Hypotheses	All 1st PCs	0.088***	(0.032)	All characteristics	0.051*	(0.031)

Notes: This table reports estimates of the effect of workers selecting home work based on initial ability after conditioning on worker characteristics. The regression specification is given by Equation 4. The dependent variable in Panels A and B is the log of net speed (the number of accurate characters per minute) during three speed tests conducted during the job interview and training process prior to beginning work. Columns (2) and (5) present the coefficient estimates on Pref_home$_i$ from running regression Equation 4. The corresponding standard errors are presented in columns (3) and (6). Rows describe the controls included in the Characteristic$_{i,h}$ controls. The first row presents the baseline effect when no characteristics are controlled with each section separated by dashed lines representing one hypothesis. Section 5.1 describes the characteristic variables. The final line of each section denoted by "All characteristics" in column (4) represents the selection effect when controlled for all characteristics listed in the particular hypothesis section. Columns (1)-(3) represents the selection effect when controlled for the first principal component of the set of all characteristics representing a particular hypothesis. Finally, Panel C represents the results of the selection effect when we control for all hypotheses. Columns (1)-(3) use all the 1st principal components as control whereas columns (4)-(6) use all the characteristics as controls. All regressions control for the Speed Test fixed effect, which accounts for variation that occurs in productivity due to speed tests and wave fixed effect. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively. For all other specifications, each observation in these regressions is a worker survey pair