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

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 of the ESRC and IGC. We would like to thank 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
JEL No. J01,J10,J22,J24,J46

ABSTRACT

A growing literature explores the impact of home-based versus office-based work. Differences in productivity may arise due to a treatment effect of the office or from workers with different abilities sorting into office or home work. We conduct an RCT in the data entry sector in India that exogenously allocates workers to the home or office. We find that the productivity of workers randomly assigned to working from home 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 by office workers over time. We find negative selection effects for office-based work: workers who prefer home-based work are 12% faster and more accurate at baseline. We also find a negative selection on treatment: workers who prefer home work are substantially less productive at home than at the office (27% less compared to 13% less for workers who prefer the office). These negative selection effects are partially explained by subgroups that likely face bigger constraints on selecting into office work, such as those with children or other home care responsibilities as well as poorer households.

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 have documented 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 the high growth rates both in Europe's Industrial Revolution and in the East Asian "miracle" economies. In the popular debate, these findings are often seen as causal evidence that reallocating workers from household enterprises into office- or factory-based employment 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 working 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 working from home, see for example, Barrero et al. (2021), Bloom et al. (2022), Brynjolfsson et al. (2020), and Emanuel and Harrington (2023).

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 how these vary with workers' cost of effort or preferences. 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. (2014), Bloom et al. (2022), Kaur et al. (2015), and Schoar (2009)).

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 a lower cost to do so), they might select into working in the office more readily than low-ability or low-effort workers (a selection on ability effect).[1] 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, workers might choose the work location where they expect to receive the highest treatment effect (selection on treatment effects). For example, high-ability workers might possess skills that complement office work, such as the ability to learn from their peers. In this scenario, treatment effects would be larger and thus they may be more likely to choose office work. Alternatively, lower ability workers might have the larger treatment effects if the discipline of being in the office is especially important for this group. 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 or even negatively corre-

[1] This self-selection could be particularly accelerated if promotions or other benefits are more likely for workers who are in the office and thus might be easier to monitor or assess. The importance of personal interactions in the office has been a major issue for many firms and employees following the move to remote work as a result of the Covid19 pandemic.

lated with it. For example, some productive workers might have commitments at home such as child or elderly care, or use home-based work as a way to bridge unemployment spells while looking for another job. In these cases the selection of workers might be negatively related to productivity, depending on the magnitude of the constraints.[2]

While the organizations literature has explored the impacts of the productivity-enhancing practices used in offices and factories, and more recently the benefits of work-from-home, we are not aware of a literature examining this second explanation for such productivity differences; that offices and factories may act as sorting devices.

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 depend directly on the work of 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. We were able to establish our own data entry operation with several hundred workers in order to control both work conditions and allocation to home and office work. 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. Potential data entry workers are recruited through ads in leading local newspapers. Qualifying applicants are invited to an entry interview where they complete 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 are 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 is greater than one half). All applicants are 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 are recorded through the data entry software.[3]

Our analysis follows three steps: First, treatment effects of home versus office work are measured by

[2] 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.

[3] These software features are relatively standard for data entry work and so workers with data entry experience would have seen similar software and metrics at other firms.

comparing the performance of people randomized into home work to those randomized into the office-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 be 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 treatment effect of being allocated to the office. The productivity of workers randomly assigned to work from home is 18% lower than that of the workers assigned to work from the 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 that is adjusted for task difficulty. The treatment effect of working in the office is especially large 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. But treatment effects are indistinguishable from zero for certain groups, e.g. poorer workers, workers preferring part-time work, 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. What lies behind these negative selection effects?

One possible explanation for this negative selection 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 (selection on treatment effects). 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 state a preference for working from home have a particularly large negative treatment effect when 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.[4] 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 chose 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 investigating the exact form of these constraints and the characteristics associated with more productive workers selecting into home work. For example, norms may prevent educated women or those with home-care responsibilities from working outside the house. 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 that is 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.

A caveat in interpreting the magnitudes of the selection effects is that, in order 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.[5] 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 chose to work from home—particularly those who have high ability and those 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

[4] Most directly, we find no significant interaction between treatment and initial ability, although higher initial ability workers do have steeper learning curves in the office.

[5] That said, our sample still included many with strong preferences: in the earliest waves of hiring for the experiment, 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 not only applicants who are close to indifferent to their work location but also those who have strong preferences but can be incentivized with a sufficiently large bonus at the end of the first week to work in an environment not of their choosing.

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 burgeoning literature that highlights large productivity differences between 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), Bartelsman et al. (2013) and Hsieh and Klenow (2014)). Most relevant, McCaig and Pavcnik (2018) show substantial labor productivity differences between household enterprises and non-household firms in Vietnam. 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 second relevant literature is the strand of work in organizational economics that 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 are willing to choose dominated contracts that help them solve these issues. These papers highlights how better managerial practices can improve worker productivity in the office. However, this literature remains silent about WFH or the role of worker sorting in generating productivity gains.

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 at the expense of agriculture, and within manufacturing the movement from the putting out system where manufacturing was done 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 being attributed to factory work itself rather than worker selection, a hypothesis we test directly.

Finally, there is a small but growing literature on the productivity effects of working from home. Bloom et al. (2014) find substantial productivity improvements from workers in a large Chinese travel agency who were allowed to WFH compared to those who remained in the office. More recently, Bloom et al. (2022) find that hybrid WFH at the same firm reduced worker attrition and had small positive impacts on output. These two experimental studies differ from ours in two ways. First, the workers in Bloom et al. (2014) and Bloom et al. (2022) were selected from the subset of workers who were already at the firm (and in Bloom et al. (2014), the subset who additionally volunteered to work from home), thereby shutting down the selection channel at the center of our analysis. Our work is complementary to these studies as we set out to analyze the role that sorting plays in driving productivity differences between home-based employment and office-based work settings. Second, and closely related, the employees in Bloom et al. (2014) and Bloom et al. (2022) were existing office workers who had previously been working in an office environment in the same firm and thus might have already absorbed office-work norms. In contrast, our study population is not only poorer and less educated, but also many of our job applicants have not previously worked in formal office environments, let alone an office-based data entry job, and so will not have already learned many of the productivity-enhancing work habits that office work may foster.

There is also a recent non-experimental literature showing that home-based work arrangements lead to lower productivity. For example, Gibbs et al. (forthcoming) 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 workers report better work-life balance (e.g. Choudhury et al. (2022), Angelici and Profeta (2023)). Our experimental design allows us to deal with potential confounds inherent in comparisons of productivity across different workers assigned to or choosing different work environments, or the same worker whose assignment or preference changes over time.

We also must caveat that we purposefully held the management technology fixed across both work locations and at low levels. 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 telephone hotline. We chose this level of management to be similar to existing data entry operations. However, thanks in part to the recent uptake in WFH, new management practices have been developed to better cater for the needs of home-based workers where face-to-face communication is limited and monitoring is more invasive. As these new practices become more developed, an alternative experiment would be to compare office and home work under different sets of management practices optimized for each setting.

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 inter-

[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), Angelici and Profeta (2023))

act 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.

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. The first theoretical reason behind such sorting is that the office is more demanding given fixed schedules, has stricter norms and more peer pressure—demands less costly or unpleasant for more productive workers. There may be additional, long-term benefits (e.g. promotions) from office work due to greater interactions with supervisors—an attractive feature to ambitious types. In either case, we expect the office to attract relatively high-ability workers. The productivity impacts of such sorting are magnified if high-ability or high-effort workers are complements with each other in production, either through peer learning dynamics or an O-ring production function.

Sorting along ability may also occur due to preference-relevant characteristics that correlate to 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 levels and education, potentially generating selection within genders. Relatedly, office work might be a status good, particularly so for workers with low social status.

Selection on treatment effects Workers may experience heterogeneous impacts of home versus office work with those with relatively high returns to office work more likely to select into such environments. If the productivity-enhancing features of the office complement ability (e.g. higher ability workers learn more from those around them), this selection on treatment effects may drive selection on ability. If low ability workers expect to gain more (e.g. from being close to supervisors) the reverse will be true. 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

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

on ability is unclear—high-ability workers may be more conscious of their self-control problems, but such problems may plague low-ability workers relatively more (potentially because of greater distraction and less privacy at home).

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. We selected this sector since it provides five 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, due to the discrete nature of the tasks, the work can easily be done from home without support from colleagues using the same technology as in an office setting. This last feature is crucial to ensure that productivity differences across the two work environments 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 to set up the office environment and the support structure for the work: upfront training, technical help with equipment problems, and compensation schemes were all modeled after 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. Home-based workers had access to a telephone helpline to call in with problems. 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]

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

[9] A picture of the user interface of 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 begin-

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.

The data entry tasks that data entry workers had to complete were constructed by us. Each data entry task consisted of four sections with each section focusing on a different type of data entry such as typing type-set text, typing strings of random alpha-numeric characters, etc. We had two types of task, easy and difficult. The difficult tasks had the same number of sections as 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 random alpha-numeric and special characters which made typing difficult.[11] Workers were assigned easy tasks from weeks 1 to 3, followed by harder tasks in weeks 4 to 6, and finally, in last the two weeks, a random mix of both the difficult and easy tasks was assigned. Table 1 presents the timeline from recruitment through week 8.

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-base data entry jobs.[12]. 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 appendix A.3.

The interview process was designed to both elicit baseline characteristics and initial typing speed and accuracy. Applicants who responded to the ads had to answer a number of interview questions as well as typing speed tests. Furthermore, we asked applicants to choose between home and office-based work to elicit their preferences regarding the work environment. This question was incentivized as applicants were told that they would be more likely to get their first choice than their second but it was not guaranteed.[13]

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. All workers passing this screening 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 in total from an applicant sample of 892

ning of the work. The experiment received IRB approval for capturing pictures of workers.

[11] Examples of data entry tasks for both difficulty levels can be found in the Appendix Figure A.4.

[12] Sample newspaper ads for home-based and office-based jobs can be found in the Appendix Figure A.3.

[13] Due to an implementation error by the field team, rather than workers being given their preference with a probability of 0.55 they were given it with a probability of 0.5.

Table 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

over a period of 15 months beginning in January 2017.[14] Workers were hired in batches since the number of office-based workers that could be employed at any given time was 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.[15] Four groups were formed through this process: Preferred home allocated home (labeled as HH); Preferred home, allocated office (HO); Preferred office, allocated home (OH); referred office, allocated office (OO).

The randomization allows us to estimate the treatment effect of being allocated to home or office independent of workers' preferences. 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.

[14] Approximately half of the 892 applicants met the two screening criteria and were invited to the training but only 280 showed for the training. Of the 280 applicants who received training 45 dropped out prior to the beginning of the work. Hence the working sample consists of 235 workers.

[15] As discussed in footnote 13 above, although workers were told that the probability was greater than 0.5, due to an implementation error the probability was equal to 0.5.

3.4 Compensation Structure

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

As we discuss in more detail in Section 3.6, we made several incremental modifications in the earliest pilot waves before arriving at our final compensation structure—primarily to address issues with selective attrition. This compensation structure consisted of a fixed component and a performance-based variable component. The fixed component was equal to INR 8500 ($ 128.80)[16] per month which the workers were eligible to receive on completing 35 hours per week and a certain target number of data entry tasks. These data entry task targets increased with each week to accommodate learning. If workers failed to meet either of these targets for three weeks their contracts were terminated. For every data entry task completed beyond the weekly target, workers were compensated an additional INR 65 ($ 1). This constituted the performance-based variable component of the compensation. A retention bonus of INR 2000 ($ 30.30) was paid after the completion of week 1.[17]

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). Any additional tasks they completed counted towards the variable component of the compensation. Because of our sorting, these were necessarily the tasks they did most poorly on and they faced a penalty schedule based on the difficulty of the task. For error rates between 0-7.5% for easy tasks and 0-15% for hard tasks, we reduced the per-task variable component of INR 65 proportionately with the share of errors. For error rates between 7.5-10% for easy tasks and 15-20% for hard tasks, 1.5 times of error rate was deducted. Finally, for errors greater than 10% for easy tasks and 20% for hard tasks, 2 times of error rate was penalized.[18]

3.5 Outcome Measures

As part of the hiring process, we administered a short survey collecting 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 for candidates selected for the experiment, 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. Along with the

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

[17] A more detailed weekly compensation structure is presented in the Appendix Table A.1.

[18] The full penalty structure is presented in Appendix Table A.2.

baseline survey, selected 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, all applicants were required to complete a typing speed test at the time of the job interview. During the training, workers were required to complete both a cash-incentivized and non-incentivized typing speed test.

A variety of data entry job outcomes were collected over the 8-week work period. The data entry job outcomes that were collected over time were the same as the ones the data entry companies collect for administrative purposes and so they did not pose additional burdens on the workers. We hired developers 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 in 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 we 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 (particularly those who preferred home-based work and were allocated to the office). Of the 50 workers in these early waves who preferred home but were allocated office, 40 quit immediately upon learning their assignment (furthermore, 30% of all workers who preffered office attrited before the start of work).

To address attrition and incentivize workers to stay longer, in later waves we adjusted the compensation structure. Most importantly, a retention bonus of INR 2000 ($ 30.3) was introduced, which was paid upon the completion of the first week of work. This amount was approximately equivalent to the average weekly earnings of workers.[19] These changes substantially reduced attrition for all groups, and crucially there were no longer differences in attrition between those allocated their preffered choice and those not.[20] Appendix Figure A.5 and Table A.4 present this analysis. As differential attrition makes interpreting treatment effects difficult, our analysis presented in the paper focuses on these later waves after this retention bonus was put in place and this issue was addressed. This leaves us with 280 workers of

[19] The job duration was also reduced from 12 to 8 weeks to limit experimental costs by reducing the time between waves when computers were unused for attrited workers, and the initial filtering for workers unwilling to work in both environments was strengthened. In earlier waves our surveyors would ask whether the worker was willing. In later waves the office managers would ask this question and probe whether the worker was sure of their answer. Appendix Table A.3 presents a complete list of these modifications.

[20] 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.

whom 235 completed training and commenced work. Results for the earlier (high-attrition) waves are relegated to Appendix Table A.7.[21]

4 Results: Treatment and Selection Effects

4.1 Baseline Characteristics

We first check that our randomization led to a balance on baseline characteristics for the groups of workers assigned to the home and office work locations.

Columns (1)–(3) of Table 2 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 measured by the speed test conducted during their initial interview or the two speed tests administered as part of training—including one test that was incentivized through cash payments based on performance.[22] We also find no differences in the proportion of workers who preferred to work from home across the randomly assigned groups (37% of workers preferred to work from home for both groups). Significant differences (at the 10% level) appear for only 4 out of 22 characteristics. Of the workers who were assigned to work from home, 58% are women whereas only 43% are women in the assigned-office group (significant at the 2% level). The home group has 6% fewer workers with family care responsibilities and 7% more workers who have used a computer before (significant at the 7% and 3% level, respectively).

The last three columns of Table 2 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 clearly non-random and are 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.

[21] The sample from the earliest (high attrition) waves also appears to contain relatively more applicants who had been out of the labor market for extended periods and is less representative of the flow population that enters the job market. During the early waves, 8 advertisements over 3 months attracted 79 applicants per ad. During the later waves, 33 ads over 16 months attracted 27 applicants per ad.

[22] The first test, conducted during the interview process, was an hour-long typing test that a novice with no introduction to data entry could complete. The next two tests were shorter 25-minute tests that workers took as part of the first day of training. We incentivized the second of these by paying a reward based on the total number of correct characters. All three tests were conducted in an office as that is where the interview and training took place.

Table 2: 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
# of Kids	0.21	0.20	0.87	0.25	0.18	0.29
Has family care responsibility (==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
No. 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						
Least concerned with Last Minute Effort (Ranking 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 (Yes)	0.94	0.98	0.13	0.93	0.98	0.06
At Home Commitments -Aspiration (Yes)	0.32	0.35	0.64	0.32	0.34	0.72

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

4.2 Treatment Effects

To estimate the impact of the random assignment to working from home on worker performance, we run 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$_{i,t}$ are our outcome measures described below and measured at the worker i and task t level. 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 (either wave 4 or wave 5), week fixed effects capturing the week of employment the outcome is measured in (ranging from week 1 to week 8),[23] and task-type fixed effects capturing the difficulty of the data entry task being performed (easy or hard). Although our unit of observation is the performance on a particular data entry task, i.e. an individual survey that requires data entry and takes about 2 hours to enter, the regression is essentially at the individual level as we reweight the regressions so that each worker has a total weight of 1 over all its observations and standard errors are clustered at the level of the individual. Table 3 reports these regression results.

Our primary measure of worker performance is log(Net Speed) where Net Speed is defined as correct entries typed per minute. We find that the employees allocated to work from home experience a drop of 18% in net speed (Table 3 column (1)). This effect is statistically significant at the 1% level. Columns (2) and (3) report treatment effects for Gross speed and Accuracy which are defined as total entries typed per minute and the ratio of correct entries typed to total entries typed, respectively.[24] Employees working from home see a drop of 12% in the Gross speed and 2.48 % in accuracy. Thus, the drop in net speed is mostly attributable to the drop in gross speed.

The magnitude of the treatment effect is larger when we use alternative measures of worker performance. In column (4) of Table 3, 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 would require workers to concentrate harder and expend higher cognitive effort). We find that participants assigned to work from home display 30% lower net speed on hard tasks. To incentivize workers to make fewer errors, in keeping with industry practice, we imposed an exponentially increasing penalty for remuneration. Thus employees were paid on the basis of remunerated speed and not net speed. We find that the magnitude of the treatment effect is a larger 24% when measured by the remunerated speed that punishes errors more heavily than net speed (column (5)).

One key benefit of working from home is the flexibility that it affords workers regarding their time use. We consider four outcomes pertaining to time use. In the first two measures, we explore how total time worked differs across home and office. Irrespective of work locations, all employees were mandated to work 35 hours per week. 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). These 35 hours constituted two components—time spent while 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 in and pay received for prior weeks' work).

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

[24] Net speed, Gross speed and Accuracy are related as follows: Net speed $=$ Gross speed $*$ Accuracy.

Table 3: Treatment Effect—Main Table

	(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) Attrition & Warning Weeks Included	(7) Time on Data Entry (hrs/wk) Attrition & Warning Weeks Excluded	(8) Prop of Time 9–6, M–F	(9) Idle (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	3.67***	3.80***	86.6***	3.47***	3.45***	34.7***	33.7***	0.95***	14.6***
	(0.058)	(0.041)	(1.43)	(0.047)	(0.10)	(0.93)	(0.13)	(0.017)	(0.83)
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	0.194	0.205	0.266	0.108	0.233	0.121	0.012	0.800	0.072

Notes: This table contains estimated treatment effects of allocating workers to home-based work environments. The regression specification is given by Equation 1. Column (1) regresses the log of net speed on a dummy for being allocated to home-based work. Net speed is defined as the number of accurate characters typed per minute. In columns (2) and (3), the dependent variable is the log of gross speed and accuracy. Gross speed is the number of total characters typed per minute, and accuracy is the ratio of accurate characters typed to total characters typed in percentage terms. In column (4), the dependent variable is the same as column (1), the log of net speed, and in column (5), the dependent variable is the log number of remunerated characters typed per minute (the total characters typed minus an exponentially increasing penalty for incorrectly typed characters). The dependent variable in columns (6) and (7) is the time spent actively entering data (in hours per week). In columns (8) and (9), the dependent variables are the proportion of work during office hours (i.e., between 9 am and 6 pm, Monday to Friday) and idle time, the ratio of the total time spent not moving the mouse or keyboard to the total time spent entering data, respectively. $Alloc\_home_i$ 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. All regressions account for variation arising from the type of survey section being attempted, week of employment, and the cohort of workers using section, week, and wave fixed effects, respectively. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively. Except for columns (6), (7), and (8), the unit of observation is the individual survey task pair. In columns (6), (7), and (8), the unit of observation is the week. Column (6) considers all 8 weeks for 235 workers (including those when workers received a warning or they had quit). Column (8) considers weeks when workers had done some work (including the warning weeks but not when they quit). Finally, column(7) considers only those weeks where workers had completed 35 hours of work time target (thus excluding attrition and warning weeks). Despite observations being at survey tasks or week levels, all regressions are re-weighted to give a total weight of 1 to each worker across all observations.

For the first measure of time spent on data entry in column (6), we consider all worker-week pairs of observations including the weeks when the worker failed to complete the work hours targets and received either a warning, had been fired, or had quit. We find workers randomly assigned to work from home on average worked 2.71 fewer hours (significant at the 10% level). Next, we exclude those worker-week pairs of observations where the worker either received a warning for not completing mandated work hours or the worker had stopped working. This allows us to focus on time spent working absent any attrition effects. We find that employees across both locations spent 33.7 hours actively entering data while 1.3 hours were spent on ancillary tasks pertaining to data entry with no significant difference in this behavior across both work locations (column (7)).[25]

Although the total time spent entering data does not differ across work locations, when in the week the work is done differs substantially. Individuals working in the office completed 97% of their work during office hours (i.e., between 9 am to 6 pm from Monday to Friday) (column (8)). The remaining 3% of the work that occurred during non-office hours was due to office employees being allowed to stay later 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. Finally, along with choosing when to work, working from home provides workers greater autonomy regarding breaks during working hours and potentially helps workers deal with moderate distractions. The software was programmed to measure intervals of time when no action was performed by the worker using either the mouse or the keyboard while logged in to the data entry system. The ratio of the total time spent in such intervals to the total time spent entering data is defined as idle time. We consider idle time to be a measure of small breaks and distractions while typing. Employees working from the office spend 14.6% of their time idle and this rises by 2.46% for those working from home, indicating that home workers had more of such small breaks and distractions (column (9))—although such a difference only explains a small fraction of the total productivity difference.

In Table 4, we run several robustness checks to explore the sensitivity of the treatment effect to differences in workers' baseline speed, different tasks and more granular worker characteristics. Column 1 reports our baseline estimate (column (1) of Table 3 above). Column 2 controls for workers' baseline speed during the cash-incentivized speed test (Initial Log(Net Speed)). This control should increase precision and control for bias if, despite randomization, initial performance differences are driving the productivity drop. The 18% drop in 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 performances of workers from these pre-bonus waves as well. The treatment effect remains unchanged at -18% and standard errors fall.

Our baseline re-weights each observation of the data entry task such that each employee has equal

[25] 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 infer that the rest of the time was spent on ancillary tasks.

Table 4: Treatment Effect—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***	-0.19***
	(0.050)	(0.043)	(0.043)	(0.054)	(0.049)
Initial Log(Net Speed)		0.74***			
		(0.14)			
Constant	3.67***	1.01**	3.65***	3.88***	3.64***
	(0.058)	(0.50)	(0.082)	(0.040)	(0.095)
Characteristics Controls					Yes
Section, Week and Wave FE	Yes	Yes	Yes	Yes	Yes
Observations	138,646	138,646	213,859	138,646	138,646
R-squared	0.194	0.266	0.189	0.197	0.196

Notes: This table reports various robustness checks. In all the columns, the dependent variable is the log of net speed and the independent variable is Alloc_home$_i$, a binary variable representing that 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. Column (1) repeats the baseline from column (1) of Table 3. Column (2) controls for ln_n_speed_cash, the log of net speed during the incentivized speed test conducted during training and prior to assignment. All regressions account for variation arising from the type of survey section being attempted, duration of employment, and cohort of workers using section, week, and wave fixed effects, respectively. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively. For all specifications, the unit of observation is the worker survey task pair. All regressions are re-weighted to equally weight each worker except column (4) where each survey task is weighted equally. In column (3), observations covering both pre and post-retention bonus waves are included. In column (4), workers attempting a greater number of survey tasks get proportionately greater weight assigned to their observations. Column (5) controls for three characteristics for which we observed baseline imbalance. These three characteristics are gender, family care responsibilities, and prior computer usage.

weight in the estimation of Equation 1. Thus, an individual data entry task receives lower weight for workers performing more tasks, either because their typing speed was faster or they attrited later. Instead, column (4) uses task-based weights, with each survey task carrying equal weight. The table shows that the 18% drop in productivity from working from home is almost unchanged, falling only slightly to 16%. Finally, as discussed earlier, the two groups formed by randomly assigning work locations are reasonably but not perfectly balanced (columns (1)-(3), Table 2). In column (5), we control for three characteristics across which the two groups are not balanced: gender, family care responsibilities and prior computer usage. Controlling for these characteristics, we find that the treatment effect marginally increases to 19%. Taken together, there is strong evidence that workers are more productive when completing the same number of work hours in an office environment compared to a home environment.

4.2.1 Cumulative Learning

We now analyze how much of the treatment effect in the office is due to differential learning versus a static treatment effect that is apparent from day one. Workers in both home and office locations experience an increase in productivity over the period of employment. The top-left panel of Figure 1 plots the average net speed of workers in both locations over the 8-week employment spell, with the drop in week 4 due to the fact that workers were assigned harder tasks in weeks 4 to 6 (and a mix of hard and easy in weeks 7 and 8). The top-right panel of Figure 1 separately plots the average net speed for each difficulty level for each work location across the 8 weeks. Finally, the bottom-left panel plots the cumulative learning—the percentage increase in net speed relative to the initial speed for tasks of that difficulty level, as measured by the speed for the first four tasks of that type—against the number of weeks experience the worker has with that type of task. 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 initial gain from experience is particularly substantial for office workers performing hard tasks (i.e. the change in speed observed already in week 1, shown in the lower-left panel). However, in subsequent weeks the gap between office and home workers narrows slightly for hard tasks while, if anything, widening, for easy tasks.

To test how much of the total productivity advantage of the office is due to differential learning we now look at the change in productivity over time. Columns (1)–(3) of Table 5 again uses the specification in Equation 1 but with the log of day 1 net speed, the log change in net speed relative to day 1, and the log of net speed excluding day 1 as the dependent variables (separately by task type). The sum of the day 1 and learning coefficients in columns (1) and (2) equal the post day 1 effect in column (3).[26] Office workers

[26] To be more precise, we use the first four surveys as our day-1 measure as each survey takes approximately 2 hours. The negative 20% treatment effect reported in column (3) is slightly different than the treatment effect reported in column (1) of Table 3 both because we exclude the first 4 surveys and because the regressions in the Table 5 are weighted slightly differently to ensure that the sum of columns (1) and (2) equal column (3). While Table 3 re-weights each observation such that each employee has an equal weight across all observations, here we re-weight each observation such that each employee task difficulty type

Figure 1: Learning Over Time

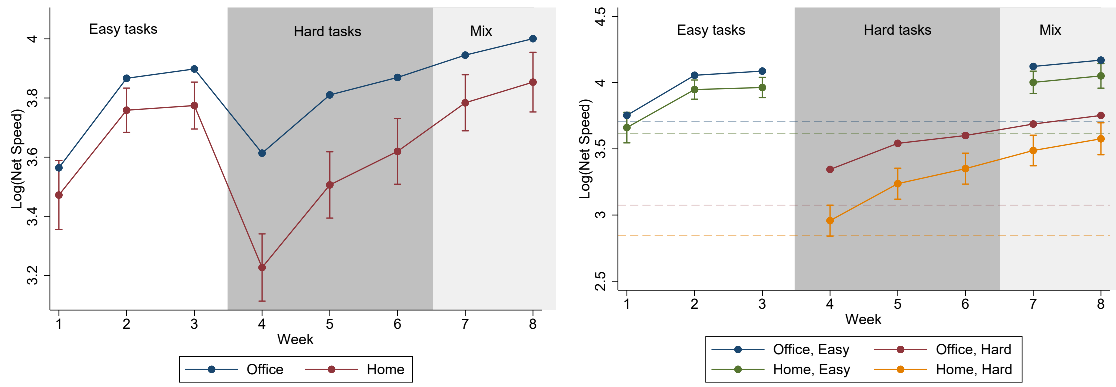

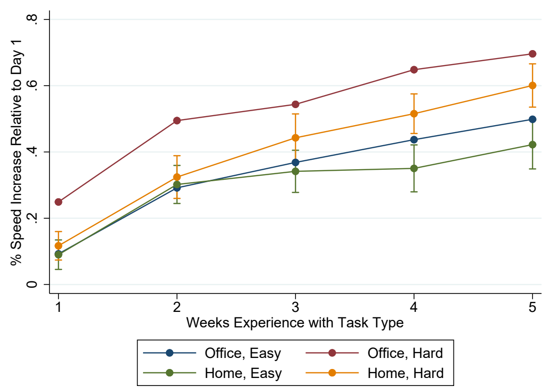

Notes: These figures plot net speed and learning over time. The top-left panel plots the average net speed for home- and office-based workers over 8 weeks. The top-right panel plots the average net speed separated by task difficulty. The dashed lines represent the average net speed on the first day for 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 type.

are approximately 13% more productive on day 1 of a task type and this difference rises another 7% over time (primarily in week 1 as seen in Figure 1) to obtain a difference of 20% on average. Columns (4)–(6) repeat the exercise only for easy tasks, and (7)–(9) only for hard tasks, with all the learning occurring on the 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 become steeper with the complexity and difficulty of the task.

4.2.2 Daily and Weekly Work Patterns

It is also interesting to see how workers assigned to home work take advantage of the flexibility and allocate their time over the week and over the 24 hours in the day. In particular one might wonder whether these workers are particularly able to work outside regular hours and show high productivity during these times. In Figure 2 top left panel, we plot the hours worked per average weekday. For people working from home, the smallest share of work was done on Mondays as home-based workers were required to visit the office to upload the data entry tasks completed in the prior week and to receive a new set of tasks to com-

combination has an equal weight across all observations.

Table 5: 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 differences and learning across home and office environments. In columns (1), (4) & (7) the dependent variable is the log of initial net speed which is defined as the average net speed for the initial four surveys completed for a particular difficulty level by each worker. In columns (2), (5) & (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 type. In columns (3), (6) & (9), we have the log of net speed excluding the first four surveys. 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. Alloc_home$_i$ is a binary variable that takes a value equal to one if the worker was randomly assigned to work from home and zero if assigned to work from the office. All regressions include section, week, and 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. For all specifications, the unit of observation is a worker survey task pair. All regressions are re-weighted to give a total weight of 1 to each worker across all observations.

plete over the following week. In a typical week, the proportion of work done steadily rises as Monday approaches, with the highest proportion of work being done on Sunday and Saturday. In contrast, office workers have a very stable work allocation Monday to Friday and do no work on the weekend. When looking at the allocation of work across a typical workday in Figure 2 top-right panel, we see that home workers start their workday a bit later than office workers, around 10 am, and then spread out the work over the day between 11 am to 10 pm. In contrast, office-based workers do almost all their work between 9 am and 4 pm with a dip around lunchtime.

Figure 2: Daily and Weekly Distribution of Work and Typing Speed

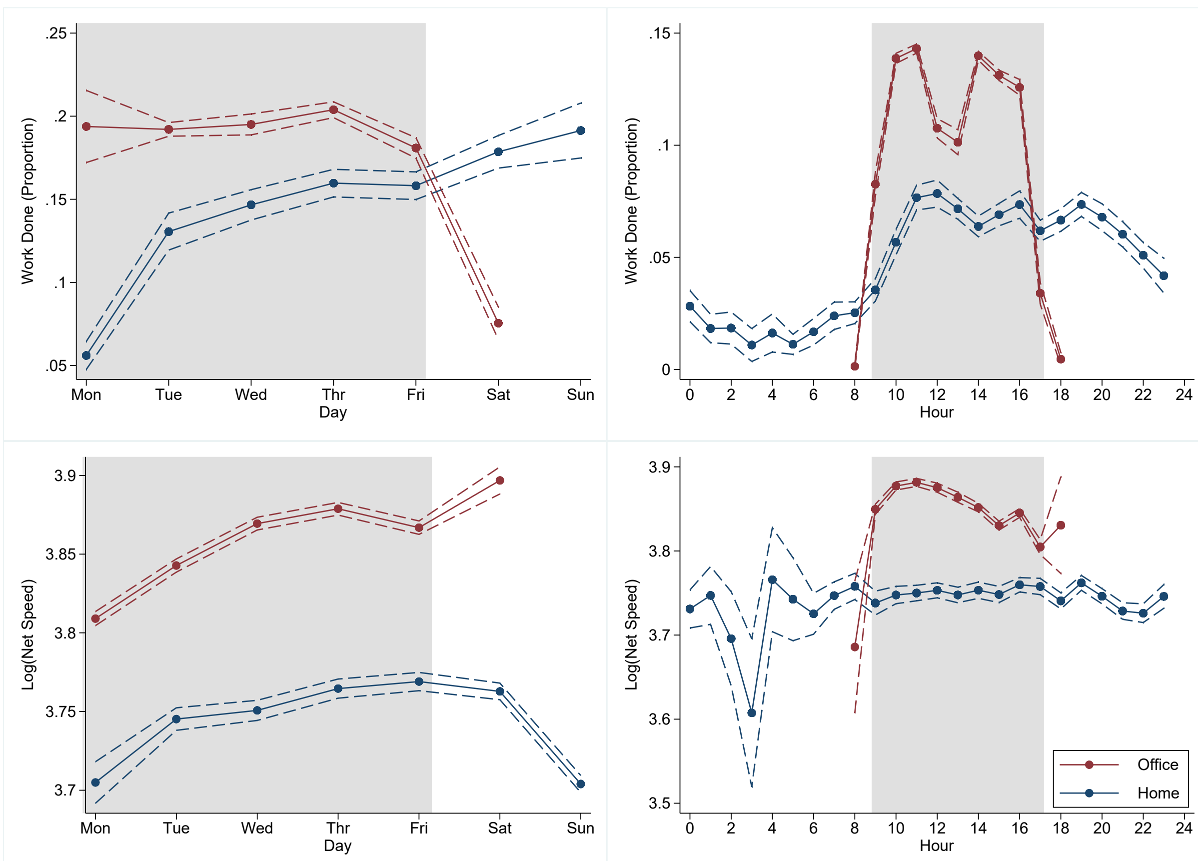

Notes: This figure plots the distribution of work done and typing speed by work location over days of the week and hours of the day. The top-left panel plots the proportion of work completed on each day of the week, across both work locations. The top-right panel plots the proportion of work completed on each hour of the day, across both work locations. The bottom-left and bottom-right panel plots the average log(net speed) for both work locations over days of the week and hours of the day, respectively.

How much does worker productivity vary by when the work is done? The bottom panels of Figure 2 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 considerable drop only observed in the middle of the night (2 am–4 am).

In addition, we analyze how workers utilize the flexibility of being assigned to WFH. Do those who request WFH and are assigned to it have different work patterns or are particularly productive when working outside regular office hours? We do not find support either of these hypotheses. There is no significant difference in the time allocation of workers who are assigned to WFH and chose WFH versus those who were assigned to WFH but chose working from the office, as measured either by the fraction of work done outside regular office hours, including the weekends, or just the fraction of work done on weekends.[27] Just focusing on those working at home, those who requested WFH are no faster in after hours or weekend hours compared to those who requested to work from the office. These results are reported in Appendix Table A.8.

In sum, we find that people working from home take advantage of the flexibility of working outside regular business hours. But we do not see that the people who choose WFH use the flexibility in different ways or to greater effect than workers who prefer office environments.

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. We run two specifications to investigate if higher ability workers select into office work. The first specification regresses 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 a binary variable representing the (incentivized) work location choice of the employee prior to being allocated a work location and takes value equal to one if the worker preferred to work from home and is equal to zero otherwise; X_n are controls, specifically fixed effects for each of the three different speeds tests that were administered (which we include only when all three test scores are combined in the same regression).

Table 6 presents the results of estimating Equation 2. Column (1) considers the sample of all 884 applicants who showed up for walk-in interviews and did a speed test. As the coefficient on Pref_home$_i$ in column (1) indicates, contrary to the hypothesis that there may be positive selection on ability into office work, applicants preferring home-based work typed 15% faster in terms of net speed during the hour-long

[27] We drop the people who were assigned to the office since they did not have any flexibility to adjust their work hours.

Table 6: Selection on Ability—Main Table

	(1)	(2)	(3)	(4)	(5)
			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*	0.083*
	(0.025)	(0.049)	(0.033)	(0.050)	(0.048)
Alloc_home					-0.18***
					(0.049)
Constant	3.08***	3.13***	3.22***	3.55***	3.64***
	(0.023)	(0.037)	(0.032)	(0.058)	(0.063)
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	0.089	0.040	0.148	0.181	0.197

Notes: This table contains estimates of the degree of selection based on initial ability. 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$_i$, 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$_i$ 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. 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 who were selected to start working for us and turned up on day 1 of the job. Column (3) adds observations from two additional tests performed by hired workers. All specification control for wave fixed effects. 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. 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 3. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

data entry test conducted as part of the interview. This difference relative to those who preferred home is significant at a 1 % significance level. In column (2), we restrict our sample to only include the 234 workers who moved forward to training.[28] As discussed previously, applicants were offered the work only if they were between the ages of 18-40 and, more importantly, were able to confirm that they could work in a work environment that is not their first choice. Thus applicants with the most extreme preferences may have been filtered out. When the restricted sample is considered, the selection effect persists although it is diminished to a 10% difference (significant at the 5 % level). Our preferred selection specification is presented in column (3) where we include speeds from the three different speed tests conducted prior to the start of work. By including three tests we increase precision and can also include the cash-incentivized

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

test. As two of these tests were part of the training, we focus only on the sample of workers who progressed from the interview stage and started training. We find workers preferring home are 12% faster than workers preferring office, significant at the 1% significance level. In sum, whether we look at the full sample of job applicants or those ultimately selected for work, we see that more productive workers at baseline are more likely to prefer working from home. Additionally, the results indicate that the filtering of candidates to only include those who are willing to work in either location does not generate the observed selection effect in the worker sample but, if anything, diminished the size of the selection effect.

We next investigate whether the same selection effect is present in the performance of employees over the subsequent two months of employment. To do so, we run the specification given by Equation 3:

$$\text{Worker Performance}_{i,t} = \alpha \text{Pref\_home}_i + \gamma X_{i,t} + \epsilon_{i,t} \tag{3}$$

where $X_{i,t}$ captures the wave, week, and section fixed effects. The measure of worker performance is again the log of net speed. Results for this specification are reported in columns (4) and (5) of Table 6. Though slightly smaller in magnitude than the initial difference, we again find those who prefer home work perform better in whatever location they were assigned to, with 8.4% higher speed (statistically significant at 10% level), see Table 6 column (4). Finally, in column (5) of Table 6 we explore what happens to this selection effect once we control for the work location they were allocated to. Since the allocation of work location is randomized and so should be uncorrelated with preferences, it is reassuring that controlling for work location barely changes the magnitude and significance of the selection effect.

We also explore the robustness of these selection effects. Appendix Table A.9 analyzes how the sorting on work location preference manifests itself in other productivity measures. Overall, just like net speed, both applicants' and workers' samples reflect a positive selection on home-based work in gross speed, accuracy, and proportion of 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 other weeks, 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.10, we run an identical set of specifications as the preceding Table 6 except we control for the selection effect introduced by the type of newspaper ad workers responded to. On adding the type of newspaper ad control, the selection effect based on self-reported work location preference in the 884 applicants sample declines to 12% from 15% (still significant at the 1% significance level). In contrast, the selection effects rise relative to Table 6 when restricting the sample to only workers who attended training (column (2)) and when considering performance across the full employment period (with the selection effect now significant at the 5% level in the latter case).

The variation generated by the different newspaper ads provides another dimension of selection to explore, rather than just serve as a control as in our previous selection analysis. In the applicant sample, we find that applicants responding to home-based work newspaper ads are faster than those responding to office-based work ads. The direction of selection driven by the type of ad responded to is the same as the selection on self-reported work location preference. Column (1) of Appendix Table A.10 illustrates the performance among those responding to home ads is 6.7% higher and is significant at the 5% level.[29].

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 productivity differences across office and home-based production in observational data. In Section 5.1 we will try to understand the origins of this sorting, i.e. whether there are differences in the preferences or constraints that different groups face when selecting a work environment by exploring how the selection effect attenuates with the addition of different sets of observable worker characteristics.

4.4 Selection on Treatment Effects

Before looking at the heterogeneity by worker characteristics in the next section, we first explore whether there are selection on treatment effects. The negative selection we documented above might be the result of low-ability workers expecting to benefit more than high ability workers from being in the office, and thus they might be more likely to choose to work there. For example, they might face more distractions at home or need more help from peers to learn. This selection on treatment effects would mean that the workers who prefer office-based work should see greater improvements in their productivity if they are allocated to the office, compared to workers who prefer home-based work. To test for this possibility, we explore how the treatment effect interacts with the preference for office- or home-based work.

To test for selection on treatment effects, we run the following regression specification:

$$\text{Worker Performance}_{i,t} = \alpha \text{Alloc\_home}_i + \delta \text{Pref\_home}_i + \lambda \text{Pref\_home}_i * \text{Alloc\_home}_i + \gamma X_{i,t} + \epsilon_{i,t} \quad (4)$$

Worker Performance$_{i,t}$ for worker i and survey t is measured along a number of dimensions as before, by log net speed, log gross speed, accuracy, and share of time idle; Alloc_home$_i$ and Pref_home$_i$ are binary variables taking a value equal to 1 when the worker was randomly assigned to work at home and when workers self-reported preference is to work from home, respectively; and $X_{i,t}$ capture week, section and wave fixed effects. The coefficient λ on the interaction between the preference for home and allocation to home-work, Pref_home$_i$ $*$ Alloc_home$_i$, captures the selection on treatment effect. In other words, do those whose productivity falls relatively less from home work disproportionately prefer it?

[29] We do not further consider the selection driven by advertisement type in part because our filtering process to select workers from the applicant sample tampers the selection by ad type. This is evident in columns (2) and (3) of the workers' sample where the selection driven by ad type operates in the opposite direction. The applicants from home-based work adverts are 6.7% faster than applicants from office-based adverts but the selected workers from home-based work adverts are 3.6% to 8.6% slower compared to the workers selected from office-based work advert. This flipped sorting effect grows even stronger when considering worker performance across 8 weeks of employment. However, we provide separate results of the main specifications for samples recruited by home- and office-based work ads in Appendix Table A.6 and observe that limited heterogeneity exists along with dimension

Table 7: Selection on Treatment Effects—Main Table

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log(Net speed)		Log(Gross speed)		Accuracy		Idle time	
Pref_home	0.14**	0.067	0.062	0.028	3.83***	3.91***	-0.32	-0.34
	(0.064)	(0.058)	(0.049)	(0.032)	(1.37)	(1.37)	(1.04)	(1.04)
Alloc_home	-0.14**	-0.13**	-0.086*	-0.055*	-1.67	-1.60	1.76*	1.74*
	(0.066)	(0.055)	(0.044)	(0.030)	(1.53)	(1.55)	(1.05)	(1.04)
Pref_home*Alloc_home	-0.12	-0.14*	-0.081	-0.13***	-2.17	-2.26	1.88	1.99
	(0.094)	(0.082)	(0.067)	(0.049)	(2.11)	(2.12)	(1.68)	(1.68)
Baseline Worker		0.75***		0.57***		-0.033		0.042
Performance		(0.14)		(0.066)		(0.043)		(0.11)
Constant	3.62***	0.96*	3.78***	1.52***	85.1***	87.3***	14.7***	14.6***
	(0.067)	(0.49)	(0.047)	(0.26)	(1.60)	(3.14)	(0.90)	(1.05)
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	0.198	0.268	0.208	0.357	0.270	0.270	0.074	0.075

Notes: This table contains estimates of selection on treatment effects. The regression specification for all columns is given by Equation 4. In columns (1) & (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 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. pref_home is a binary variable representing workers' preference for work location taking a value of one if they chose home-based work and zero if they chose 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. Baseline worker performance controls are the same performance measures as the dependent variable but calculated from the cash-incentivized speed test conducted during the training. Regressions include section, week, and 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. Standard errors (in parentheses) are clustered at the individual level. *, **, and *** denote statistically significant at the 10%, 5%, and 1% level, respectively.

Table 7 presents these results. In column (1), we see that in terms of net speed, workers who prefer to work from home are 14% faster than those who choose office, while workers randomly allocated to home are 14% slower. However, the interaction between preferring to work from home and being allocated to work from home is equal to negative 12%. I.e. those who prefer to work from home are 12 percent slower when working from home than those who prefer to work from the office are when working from home. This negative selection on treatment effect is sizeable but insignificant. However, when we control for baseline speed to increase precision, as we do in column (2), we find that the negative selection on treatment effects increases to -14% and are statistically significant at the 10% level. 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). In addition, treatment effects on accuracy are more negative and idle time more positive among those preferring office, although these are smaller and not significant.

Appendix Table A.11 looks directly at whether high-ability individuals have larger treatment effects of office work. To do so, we replace the Pref_home_i dummy in Equation 4 with the (logged) baseline speed measure that captures initial ability. 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 not significant. 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 prefer home-based work experience a large increase in productivity when they work from the office compared to home, even though they preferred to work at home. People who prefer office-based work also experience higher productivity in the office, but the treatment effect is only about half as large. Thus, there is negative selection on treatment effects. The hypothesis that low ability or low self control workers have the most to gain from office work in terms of productivity and so sort into such an environment is rejected, so it cannot explain the negative selection on ability found in the preceding section. Instead, the results again point to an explanation where 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 exploiting heterogeneity in effects by worker characteristics.

5 Exploring the Origins of the Negative Selection Effects

The prior analysis shows that despite a strong positive treatment effect of the office, high ability workers are less likely to sort into the office. Furthermore, high ability workers are the ones who benefit most from being in office. In this section, we explore what might be the origin of the negative selection on ability and the negative selection on treatment effects.

5.1 Constraints to Selection on Ability

We now compile potential hypotheses for why high ability workers might be constrained from selecting into the most productive workplace or have preferences for a different allocation. Based on these hypotheses, 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. If we directly include proxies for these omitted factors in our selection regressions, initial ability should no longer be (as) correlated with work environment preferences if the mechanism is important. Thus, in the final step, we control for these characteristics with our proxies for various potential omitted variables and explore the degree to which these proxies attenuate the negative selection on ability.

We explore the following six hypotheses that can generate negative sorting into office work based on initial ability (with the corresponding controls in parentheses):

1. High-ability workers tend to live further away from the office and so incur higher time and effort costs commuting every day. [Distance between home and the office location.[30]]
2. The office serves as a commitment device for low self-control/low-productivity workers. [Agreement with statement "I never leave things to the last minute", the time discount parameter.[31]]
3. Office work is a status good for lower-ability workers. [Number of previous office jobs, total monthly household income, the interaction of the two.]
4. High-ability workers combine the job with the search for better job opportunities or study which is easier to do with a more flexible schedule. [Prefer full-time work, have additional time commitments.]
5. High-ability workers have more responsibilities at home (or low-ability workers anticipate more distractions at home and thus choose the 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 thus choose office). [Female; interactions of female with family care responsibilities, married, has kids, age.]

Specifically, we run the following specification:

$$\text{Initial Worker Performance}_{i,n} = \beta \text{Pref\_home}_i + \sum_h \gamma_{1,h} \text{Characteristic}_{i,h} + \gamma_2 X_n + \epsilon_{i,n} \tag{5}$$

where Initial Worker Performance$_{i,n}$ is the log of speed of worker i in the three initial speed tests indexed by n; Pref_home$_i$ takes a value equal to one if worker i prefers to work from home; $\{\text{Characteristic}_{i,h}\}_{h=1}^{H}$

[30] 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.

[31] 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.

denotes the set of characteristics that proxy for the hypothesis; and X_n are three fixed effects that control for aggregate speed difference on the three different speed tests.

We summarize our results in Panel A of Table 8 (for completeness Appendix Table A.12 contains the full set of coefficients). Columns (1)–(2) of the first row of Table 8 report the coefficient and standard errors for our baseline selection estimation that includes no characteristic controls (Table 6 column (3)). Subsequent rows report regressions that include the additional sets of controls discussed above but only the selection effects are reported, the coefficient β on Pref_home$_i$. The final row includes all sets of controls concurrently. As the number of characteristics representing a particular hypothesis varies, it is challenging to compare across explanations. Thus, we also conduct a principal component analysis using each hypothesis' complete set of characteristics and use the first component as a control for that hypothesis. The β coefficient from these regressions are reported in columns (3)–(4).

Although the selection effects attenuate with controls, the amount of 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 workers choose the office as a commitment device.

The final row of Panel A simultaneously controls for all the hypotheses. Unsurprisingly, the attenuation is greater than with any single set of controls with the coefficient falling to 0.051 (significant at the 10% level). However, even when including all these 17 controls, there is still a substantial negative selection on ability effect 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 see if accounting for the same sets of characteristics as above explains the selection on treatment,

$$\begin{aligned}\text{Worker Performance}_{i,t} &= \tau \text{Alloc\_home}_i + \delta \text{Pref\_home}_i + \lambda \text{Alloc\_home}_i * \text{Pref\_home}_i + \\ &\sum_h \gamma_{1,h} \text{Alloc\_home}_i * \text{Characteristic}_{i,h} + \sum_h \gamma_{2,h} \text{Characteristic}_{i,h} + \gamma X_{i,t} + \epsilon_{i,t}\end{aligned} \tag{6}$$

where variables are defined as in previous specifications.

Table 8: Selection Effects—Controlling for Characteristics

Panel A: Selection on Ability

		(1)	(2)	(3)	(4)
		Controls for			
		All Characteristics		**1st Principal Components**	
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 Option	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 Characteristic		**1st Principal Components**	
Regression Specification		**Pref_home* Alloc_home**	**(SE)**	**Pref_home* Alloc_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 Option	-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: Panel A and B report estimates of selection on initial ability and selection on treatment effects after conditioning on worker characteristics. The dependent variable in Panel A 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 (1) and (3) present the coefficient estimates on Pref_home$_i$ from running regression equation (5). The dependent variable in Panel B is the log of net speed (the number of accurate characters per minute) during data entry task performed as part of employment and columns (1) and (3) of Panel B report the coefficient estimates on Pref_home$_i$ $*$ Alloc_home$_i$ from running the regression equation (6). For both Panels, columns (2) and (4) report the standard error of that estimate. 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, all regressions control for the speed test type fixed effects and wave fixed effect, which accounts for the hiring batch of the worker. In Panel A the unit of observation is the test attempted. In Panel B, the unit of observation is the survey task attempted and all regressions include section, week, and wave fixed effects. The 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 statistically significant at the 10%, 5%, and 1% level, respectively.

To understand this specification, suppose that our negative selection on treatment effects is coming from women being both more likely to prefer to work at home due to cultural constraints, and less productive at home compared to the office because of home responsibilities. In this scenario, by interacting a female gender dummy in $\text{Characteristic}_{i,h}$ with Alloc_home_i, we would no longer observe selection on treatment effects (i.e. $\lambda = 0$).

Panel B of Table 8 presents the λ coefficients after the inclusion of the controls. The first row repeats the -14% coefficient from our baseline regression. Rather than the inclusion of controls attenuating this coefficient, for all six sets of hypotheses above, λ becomes more negative, growing to -18% when all controls are included. Thus, heterogeneity in treatment effects by worker characteristics coupled with correlations between characteristics and home preferences are not behind our negative selection on treatment effects.

Note that this result does not fully rule out the possibility that cultural constraints for 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 $\text{Gender}_i * \text{Alloc\_home}_i$ is included since the treatment-effect heterogeneity is not across genders per se but across 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 bisect 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. We then interact an indicator for one of the two groups resulting from the bisection, sub_group_i, with alloc_home_i, pref_home_i, and the interaction of the two:

$$\begin{aligned}\text{Worker Performance}_{i,t} = {} & \tau \text{Alloc\_home}_i + \delta \text{Pref\_home}_i + \lambda \text{Pref\_home}_i * \text{Alloc\_home}_i + \\ & \tau' \text{Alloc\_home}_i * \text{sub\_group}_i + \delta' \text{Pref\_home}_i * \text{sub\_group}_i + \\ & \lambda' \text{Pref\_home}_i * \text{Alloc\_home}_i * \text{sub\_group}_i + \theta \text{sub\_group}_i + \gamma X_{i,t} + \epsilon_{i,t} \quad (7)\end{aligned}$$

where the coefficient λ tells us about selection on treatment effects for the subgroup for which $\text{sub\_group}_i = 0$, and $\lambda + \lambda'$ tells us the selection on treatment effects for the subgroup for which $\text{sub\_group}_i = 1$.

Columns (5) and (6) of Table 9 present these two coefficients, λ and λ', for each of the characteristics discussed at the start of this section. (The earlier columns explore heterogeneity in Selection and in Treatment Effects that we discuss in the next section.) For a number of characteristics, selection on treatment effects occurs only within one of the two subgroups and is particularly pronounced for households with low family income, those with family care obligations (particularly women), those with children, and older workers. In all these cases, 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 selection on the treatment effect of -14%, these subgroups exhibit additional negative selection on treatment effects of between 37% and 81%.

The groups within which selection on treatment is largest are certainly suggestive of societal constraints lying behind the unexpected negative sign of the 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 heterogeneity makes it possible that those who choose home work 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 (e.g. their mother-in-law also helps with housework). In contrast, we find no negative selection within groups that are free of these constraints, and thus face no heterogeneity in the severity of the constraint, such as women without family obligations. 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 designing labor market policies.

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 discussion in the previous section, we bisect our sample into two subgroups, above and below the median, for each characteristic discussed at the start of Section 5.1. We then repeat the specification exploring worker selection on ability but interact the indicator variable representing worker's preference to work from home with an indicator variable for membership of one of these subgroups:

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

where our measure of Initial Worker Performance is the log of net typing speed averaged over the three different speed tests that were conducted prior to the beginning of the job.

Columns (1) and (2) of Table 9 report these results. Recall that workers who prefer WFH are 12% faster than the ones who prefer to work in the office (repeated in row (1)). There is relatively little heterogeneity in the size of this coefficient with only two of the 16 α' interaction coefficients significantly different from zero (although with 704 tests across 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 home are 11% more productive prior to starting the job than those who prefer office, while women who prefer home are 13% more productive, with the 2% difference not significant.

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

Columns (3) and (4) of Table 9 present a similar exercise but where we interact sub_group_i dummies with Alloc_home_i to explore heterogeneity in treatment effects:

$$\text{Worker Performance}_{i,t} = \alpha \text{Alloc\_home}_i + \alpha' \text{Alloc\_home}_i * \text{sub\_group}_i + \theta \text{sub\_group}_i + \gamma X_{i,t} + \epsilon_{i,t} \tag{9}$$

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 treatment effect.[32] Only three characteristics show significant interactions at a 10% significance level: monthly family income of the worker's household, female workers with family care responsibilities, and older female workers. Interestingly, female workers with family care responsibilities are 11% faster at home and are the only subgroup showing positive (although not significantly different from zero) WFH treatment effects. (Although recall, from the discussion of selection on treatment effects in columns (5) and (6), that those who experience these large positive treatment effects of WFH are disproportionately the workers who choose to work in the office, and that those choosing to WFH experience large negative treatment effects of home work.)

Similarly, subgroups such as workers with family care responsibilities, workers preferring to work part-time, and workers with low family income experience the smallest treatment effect at -5%, -8%, and -10%, respectively. These effects are statistically indistinguishable from zero. On the contrary, subgroups such as older female workers, workers with high family income, and married workers exhibit the strongest treatment effect at -38%, -28%, and -25%, respectively. All these effects are significantly different from zero at a 5% significance level, even if the difference in treatment effects is not.

[32] Recall from the brief discussion in Section 4.4, Appendix Table A.11 finds no significant evidence of heterogeneity in treatment effects by initial ability.

Table 9: Heterogeneity in Selection on Ability, Treatment Effects, and Selection on Treatment

		(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Baseline Regression							
		Pref_home		Alloc_home		Alloc_home* Pref_home	
Baseline		0.12***		-0.18***		-0.14*	
		(0.033)		(0.050)		(0.082)	
Panel B: Regressions with Heterogeneity							
Hypothesis	Characteristic	Pref_home	Pref_home* sub_group	Alloc_home	Alloc_home* sub_group	Alloc_home* Pref_home	Alloc_home* Pref_home* sub_group
Costs	abv_avg_dist	0.15***	-0.07	-0.20***	0.05	-0.03	-0.21
		(0.05)	(0.07)	(0.07)	(0.10)	(0.11)	(0.17)
Low Discipline	last_min_effrt_yes	0.11**	0.01	-0.18**	-0.01	-0.16	0.03
		(0.05)	(0.07)	(0.07)	(0.10)	(0.11)	(0.17)
	high_discount	0.05	0.14**	-0.12**	-0.11	-0.14	0.00
		(0.04)	(0.06)	(0.06)	(0.10)	(0.10)	(0.16)
Status	low_fam_income	0.09*	0.05	-0.28***	0.18*	0.09	-0.42**
		(0.05)	(0.07)	(0.08)	(0.10)	(0.12)	(0.17)
	no_prior_off_job	0.10***	0.03	-0.16***	-0.05	-0.18*	0.11
		(0.04)	(0.08)	(0.06)	(0.10)	(0.10)	(0.18)
Outside Option	fulltime_pref	0.24**	-0.14	-0.08	-0.12	-0.19	0.04
		(0.10)	(0.11)	(0.13)	(0.14)	(0.17)	(0.18)
	commit_prof_asp	0.09**	0.08	-0.16***	-0.06	-0.17*	0.07
		(0.04)	(0.07)	(0.06)	(0.11)	(0.10)	(0.19)
Home Pressure	fam_care	0.10***	0.13	-0.19***	0.14	-0.10	-0.74***
		(0.03)	(0.12)	(0.05)	(0.16)	(0.09)	(0.14)
	married	0.08**	0.14	-0.17***	-0.09	-0.11	-0.13
		(0.04)	(0.09)	(0.06)	(0.12)	(0.09)	(0.20)
	child_yes	0.10***	0.12	-0.18***	-0.07	-0.07	-0.48**
		(0.03)	(0.10)	(0.05)	(0.14)	(0.09)	(0.21)
	abv_avg_age	0.06	0.12*	-0.16***	-0.07	-0.01	-0.37**
		(0.04)	(0.07)	(0.06)	(0.10)	(0.10)	(0.18)
Female Constraints	female	0.11**	0.02	-0.16**	-0.07	-0.21	0.10
		(0.04)	(0.06)	(0.07)	(0.10)	(0.13)	(0.17)
	female_fam_care	0.11***	0.13	-0.20***	0.31*	-0.11	-0.81***
		(0.03)	(0.16)	(0.05)	(0.17)	(0.09)	(0.16)
	female_married	0.10***	0.16	-0.18***	-0.02	-0.13	-0.00
		(0.03)	(0.11)	(0.05)	(0.14)	(0.09)	(0.22)
	female_child	0.10***	0.16	-0.18***	-0.07	-0.10	-0.31
		(0.03)	(0.14)	(0.05)	(0.15)	(0.09)	(0.22)
	female_old	0.10***	0.12	-0.15***	-0.23*	-0.13	-0.02
		(0.03)	(0.10)	(0.05)	(0.13)	(0.10)	(0.22)

Notes: This table explores the heterogeneity in selection on ability (in columns (1) & (2)), the treatment effect (in columns (3) & (4)), and selection on treatment effects (in columns (5) & (6)). For each pair of columns, Panel A presents our baseline result when we assumed no heterogeneity. Panel B reports the main effect and interaction when we interact either Pref_home_i, Alloc_home_i, or $\text{Pref\_home}_i * \text{Alloc\_home}_i$ with worker characteristic dummies sub_group_i obtained by bisecting the sample by the median value of the characteristic. Section 5.1 describes the characteristic variables. Across all regressions, the dependent variable is the log of net speed (the number of accurate characters typed per minute). Across both panels, the regressions reported in columns (1) and (2) include speed test type and wave fixed effects. The regressions reported in columns (3) and (4), and in columns (5) and (6), include section, week, and wave fixed effects. In columns (1) and (2) the unit of observation is the test attempted. In columns (3) and (4), 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 statistically significant at the 10%, 5%, and 1% level, respectively.

7 Conclusions

We set up a randomized control trial to exogenously allocate workers to home-based or office-based work while holding all other dimensions of the work constant. We first find a large positive and significant treatment effect of 18% from working in the office. Two-thirds of the effect exists from the first day of working and the rest is due to quicker learning by office workers over the subsequent weeks. Second, we find negative selection effects for office based workers. Those who prefer home-based work are 12% faster and more accurate at baseline. Finally, we also find negative selection on treatment: workers who prefer home have larger negative productivity effects when allocated to home. The negative selection effects are stronger within subgroups that typically face bigger constraints in selecting office work, such as workers with children and with other home care responsibilities, as well as poorer households.

Our results show that understanding 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 working from home, but change the availability of of office- and factory-based jobs versus home-based ones in the economy (e.g., policies that support microenterprises).

Finally, our suggestive evidence that some parts of the population are constrained in allocating their labor to their most productive work environment has implications for inequality between, for example, women and men or lower and higher class groups. This misallocation also leads to distortions in the productivity of the labor force overall. To the extent that these constraints are due to social or family pressures, policies that explicitly reduce such constraints, could improve the allocation of workers to jobs. Of course, some of these choices might be the result of cultural or personal preferences rather than external constraints. For example, in some societies, women themselves might feel that they should not work outside the home independent of their work productivity. Under these circumstances, even policies that increase childcare access might not have a large effect on female labor force participation. If preferences of some workers for WFH are lexicographic in this way, home based production might allow workers who would otherwise not be able to take up a job to participate in the economy.

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*, 11 2014, *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 & productivity: Evidence from personnel & analytics data on IT professionals," *Journal of Political Economy Microeconomics*, forthcoming.

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.

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.

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.

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.

Schoar, Antoinette, "Supervisory Skills Training in the Cambodian Garment Industry: A Randomized Impact Evaluation," Technical Report, IFC Advisory Services in East Asia and the Pacific, World Bank Group. 2009.

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.

A Appendix

Figure A.1: Pictures of the Office and a few of Home work settings

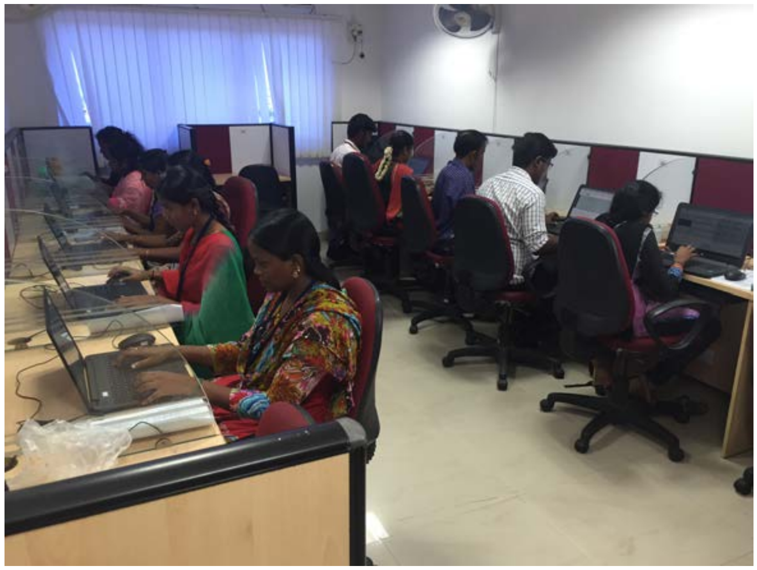

(a) The Office

(b) Home work setups

Figure A.2: User interface of a sample data entry tasks in the proprietary software

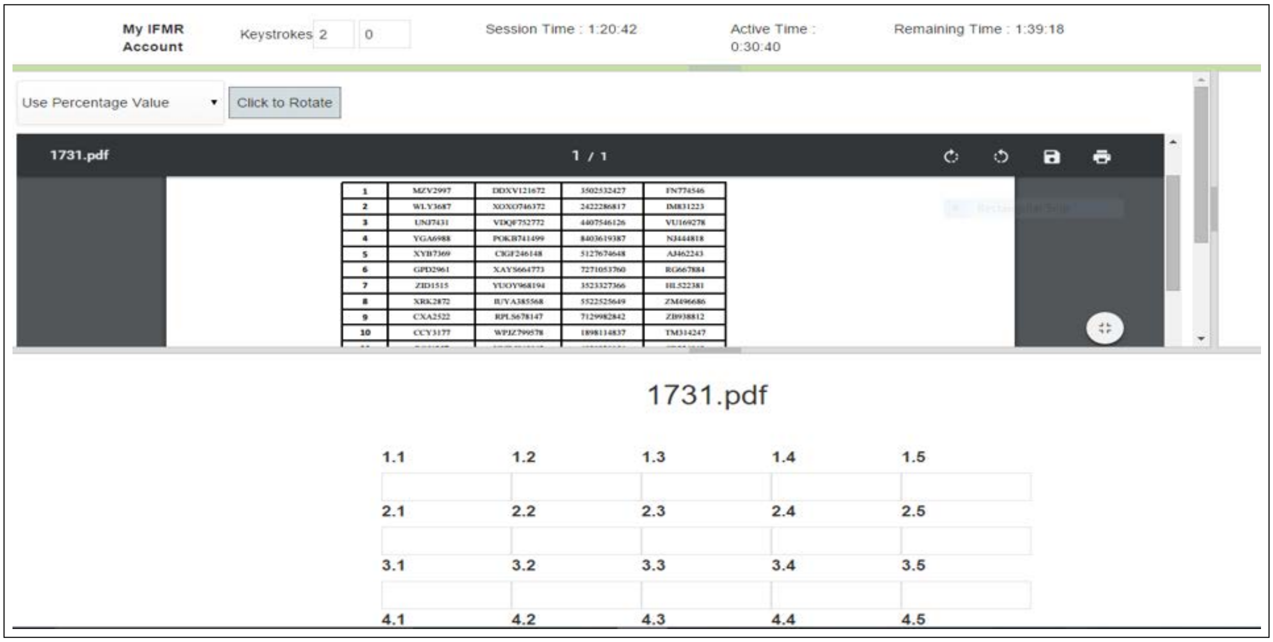

Figure A.3: Newspaper Ads sample

(a) Office-based work ad

(b) Home-based work ad

Figure A.4: 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:

<table>
<tr><th>S.no
வ. எண்</th><th>Information 1
தகவல் 1</th><th>Information 2 தகவல் 2</th><th>Information 3 தகவல் 3</th><th>Information 4 தகவல் 4</th><th>Information 4
தகவல் 4</th></tr>
<tr><td>1</td><td>?Z j?G~L_</td><td>oaFeDC-,lg</td><td>4:bcw AsBoe</td><td>-r*/o~n8</td><td>)ekO</td></tr>
<tr><td>2</td><td>X tw +DfoB</td><td>b!`x #*RK,s</td><td>VLW eNoCArM</td><td>|s.,J@=u</td><td>8X Z -o</td></tr>
<tr><td>3</td><td>5,~;honQ</td><td>K4TF|"?y#[</td><td>%~4w q?5@:t</td><td>1.$}d"^M</td><td>x 2tY9</td></tr>
<tr><td>4</td><td>]nX ^x %1m</td><td>\#*vz Z .#no</td><td>YbBm+44P35</td><td>i|{")=cY</td><td>4G*l</td></tr>
<tr><td>5</td><td>tNP$K#|C</td><td>2-/)6aCn)o</td><td>mw 9gOe[lC`</td><td>h$K~DUQr</td><td>ps;M:</td></tr>
</table>

(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 alizcheza kwa
ajili ya kujifurahisha peke yake. Baadhi ya michezo
ni distinctively kijamii (daraja): baadhi ni makusudi

Figure A.5: Attrition before and after retention bonus

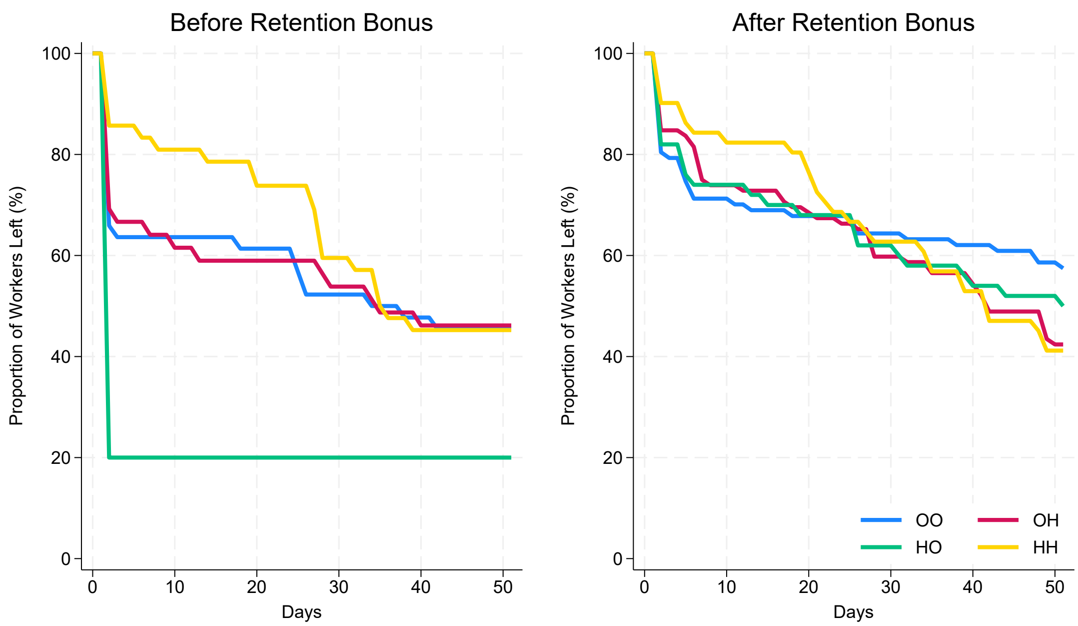

Notes: This figure plots the proportion of workers continuing to carry out data entry work against the number of days since the start of training. Left panel shows attrition prior to the introduction of a retention bonus paid at the end of the first week, right panel show attrition after. Each plot shows 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) Week	(2) Fixed component: Tasks Target	(3) Fixed component: Amount Paid INR ($)	(4) Performance-based variable component INR/task ($ / task)	(5) Retention Bonus 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 the waves

Detail	Wave 1	Wave 2	Wave 3	Wave 4	Wave 5
# Applicants **# Workers**	221 30	263 50	154 25	339 121	553 104
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)	(2)	(3)	(4)
	Pre	Post	Pre	Post
Dependent Variable	Days Worked		Worked (yes)	
$1\{\text{Preferred Home, Allocated Office}\}$	-27.0***	-2.98	-0.68***	-0.082
	(1.96)	(6.61)	(0.045)	(0.17)
$1\{\text{Preferred Office, Allocated Home}\}$	-7.11	-1.88	-0.20	-0.054
	(5.62)	(0.88)	(0.068)	(0.047)
$1\{\text{Preferred Office, Allocated Office}\}$	-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 statistically significant 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 is zero otherwise. Variable alloc_home takes a value equal to one when the work is randomly assigned to work from home and is equal to zero is the worker is randomly assigned to work from the office. Standard errors (in parentheses) are clustered at the wave level. *, **, and *** denote statistically significant 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)	
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)
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*Alloc_home			-0.14*			-0.17			-0.099
			(0.082)			(0.11)			(0.11)
Constant	3.67***	3.22***	0.96*	3.55***	3.12***	0.85**	3.78***	3.28***	1.11
	(0.058)	(0.032)	(0.49)	(0.088)	(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	0.194	0.148	0.268	0.204	0.165	0.290	0.195	0.163	0.260

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 is 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 statistically significant 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)		
Pref_home		0.12***	0.067		0.038	0.055
		(0.033)	(0.058)		(0.034)	(0.050)
Alloc_home	-0.18***		-0.13**	-0.18***		-0.12***
	(0.050)		(0.055)	(0.043)		(0.045)
Pref_home*Alloc_home			-0.14*			-0.11
			(0.082)			(0.072)
Constant	3.67***	3.22***	0.96*	3.65***	3.34***	0.58
	(0.058)	(0.032)	(0.49)	(0.083)	(0.064)	(0.40)
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	0.194	0.148	0.268	0.190	0.225	0.305

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 is zero otherwise. Variable alloc_home takes a value equal to one when the work is randomly assigned to work from home and is equal to zero is the worker is 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 statistically significant 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		Log(Net Speed)		
	outside regular office hours	during weekend			
Pref_home	-0.0029	0.013	-0.16***	-0.17***	-0.17***
	(0.038)	(0.034)	(0.025)	(0.023)	(0.025)
Pref_home*1{Outside Office Hours}			-0.0031		0.016
			(0.020)		(0.020)
Pref_home*1{Weekend}				-0.020	-0.030
				(0.024)	(0.025)
1{Outside Office Hours}			0.021		-0.016
			(0.015)		(0.012)
1{Weekend}				0.050**	0.060***
				(0.019)	(0.019)
Constant	0.57***	0.31***	4.49***	4.49***	4.50***
	(0.034)	(0.034)	(0.036)	(0.035)	(0.036)
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	0.463	0.463	0.463

Notes: This table analyzes whether those who request WFH and are assigned to it have different work patterns or are particularly productive outside regular office hours. 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 of survey section being attempted 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: 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. 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 statistically significant at the 10%, 5%, and 1% level, respectively.

Table A.10: 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***	3.63***	3.72***
	(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	0.191	0.207

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 who were selected to start working for us and turned up on the day 1 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 3. 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: Heterogeneity in Treatment Effects with Initial Ability

	(1)	(2)	(3)	(4)
	Baseline	Speed Test		
		Cash Incentive	No Incentive	Walk-in
Alloc_home	-0.18***	0.47	-0.50	0.39
	(0.050)	(0.79)	(0.52)	(0.45)
Alloc_home*Initial Log(Net Speed)		-0.19	0.089	-0.18
		(0.22)	(0.15)	(0.14)
Initial Log(Net Speed)		0.85***	0.84***	0.71***
		(0.090)	(0.081)	(0.076)
Constant	3.67***	0.67**	0.80***	1.43***
	(0.058)	(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	0.194	0.267	0.292	0.286

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 work is randomly assigned to work from home and is equal to zero is the worker is 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 initial walk-in interview. All regressions account for variation arising from the type of survey section being attempted, 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 statistically significant at the 10%, 5%, and 1% level, respectively.

Table A.12: Selection on Initial Ability- Controlling for Characteristics

Panel A: Regressions with ATE

	(1)		(2)	(3)	(4)	(5)	(6)
	Regression Specification		**Pref_Home**	**(SE)**			
	Baseline		0.116***	(0.033)			

Panel B: Controlling for individual hypothesis

	Hypothesis controlled for	**PCA**	**Pref_Home**	**(SE)**	**Characteristics**	**Pref_Home**	**(SE)**
(1)	Costs	1st principal comp	0.116***	(0.033)	distance	0.116***	(0.033)
(2)					last_min_effrt	0.105***	(0.03)
(3)	Low Dscpln				time_discount	0.110***	(0.03)
(4)		1st principal comp	0.101***	(0.032)	All characteristics	0.102***	(0.032)
(5)					off_jobs_num	0.091***	(0.03)
(6)	Status				fam_inc_scal	0.115***	(0.03)
(7)					fulltime_pref	0.107***	(0.03)
					commit_prof_asp	0.118***	(0.03)
(8)		1st principal comp	0.109***	(0.033)	All characteristics	0.089***	(0.032)
(9)					fulltime_pref	0.107***	(0.03)
(10)	Outside Opt				commit_prof_asp	0.118***	(0.03)
(11)		1st principal comp	0.116***	(0.032)	All characteristics	0.108***	(0.031)
(12)					fam_care	0.110***	(0.03)
(13)					married	0.113***	(0.03)
(14)	Home Press				child_yes	0.117***	(0.03)
(15)					age_scale	0.103***	(0.03)
(16)		1st principal comp	0.109***	(0.034)	All characteristics	0.084**	(0.034)
(17)					female	0.115***	(0.03)
(18)					female_fam_care	0.112***	(0.03)
(19)	Female Const				female_married	0.120***	(0.03)
(20)					female_child	0.117***	(0.03)
(21)					female_age_scale	0.117***	(0.03)
(22)		1st principal comp	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)**
Controlled for all Hypothesis	All 1st principal components	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 5. 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 5. 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 statistically significant at the 10%, 5%, and 1% level, respectively. For all other specifications, each observation in these regressions is a worker survey pair