

DOES WORKING FROM HOME WORK?

EVIDENCE FROM A CHINESE EXPERIMENT

Nicholas Bloom^a, James Liang^b, John Roberts^c and Zhichun Jenny Ying^d

December 18th, 2012

Abstract:

Over 10% of US employees now regularly work from home (WFH), but there is widespread skepticism over its impact and worries about “shirking from home”. We report the results of a WFH experiment at CTrip, a 16,000 employee NASDAQ-listed Chinese multinational. Call center employees who volunteered to WFH were randomly assigned to work from home or in the office for 9 months. Working from home led to a 13% performance increase, of which about 8.6% is from working more minutes per shift (fewer breaks and sick-days) and 4% from more calls per minute (attributed to a quieter working environment). Home workers also reported improved work satisfaction and their job attrition rate fell by 50%. After the experiment, the firm rolled the program out to all employees, letting them choose home or office working. Interestingly, only half of the treatment group decided to work at home, with the other half reallocating in favor of office working. After employees were allowed to choose where to work, the performance impact of WFH almost doubled, highlighting the benefits of choice when adopting modern management practices like home working.

Keywords: working from home, organization, productivity, field experiment, and China

Acknowledgements: We wish to thank Jennifer Cao, Mimi Qi and Maria Sun from Ctrip, and Itay Saporta, Stephen Terry and Edison Yu from Stanford for their help in this research project. We thank Chris Palauni for organizing our trip to Jet Blue, and David Butler, Jared Fletcher and Michelle Rowan for their time discussing the call-center and home-working industry. We thank in particular our discussants Mushfiq Mobarak, Rachael Heath, Sabrina Pabilonia, Shing-Yi Wang and seminar audiences at the AEA, Brown, CEPR, Columbia, the London School of Economics, Harvard, Stanford GSB, Texas A&M, the World Bank and the NBER for comments. We wish to thank Stanford Economics, Stanford GSB and the Toulouse Network for Information Technology (which is supported by Microsoft) for funding for this project.

Conflict of interest statement: No funding was received from CTrip. James Liang is the co-founder, ex-CEO and current Chairman of CTrip. No other co-author has any financial relationship (or received any funding) from CTrip. The results or paper were not pre-screened by anyone.

^a Stanford Economics, SIEPR, CEPR and NBER; ^b Beijing University Guanghua Management School and CTrip, ^c Stanford GSB and SIEPR, ^d Stanford Economics

I. INTRODUCTION

Working from home (also called telecommuting or telework, but hereafter referred to as “WFH”) is becoming an increasingly common practice. In the United States, over 10% of the workforce reports working from home at least one day a week (Census 2010), while the proportion that primarily work-from home has almost doubled over the past 30 years, from 2.3% in 1980 to 4.3% in 2010.¹ At the same time, the wage discount (after controlling for observables) from working exclusively at home has fallen, from 30% in 1980 to zero in 2000 as WFH moved from being predominantly found in low-skilled jobs to encompassing a wider set of occupations (Oettinger, 2010). Home-based workers now span a wide spectrum of jobs, ranging from sales assistants to managers and software engineers, with a correspondingly wide range of incomes (Figure 1).

Having employees work from home raises three issues. First, is it a useful management practice? This is an important question with no systematic evidence or consensus.² Thus, even within a single industry, practices often vary dramatically. For example, Jet Blue Airlines’ call-center employees work from home, American Airlines does not allow any home work, and United Airlines has a mix of practices. More generally, Bloom et al. (2009) reported that 30% of U.S. and 33% of European manufacturing firms offer opportunities for at least some managers to work from home, with wide variation in adoption rates within every 3-digit SIC code surveyed.

The second issue relates to the concerns over deteriorating work-life balance in the US, and the potential of working from home to help address this. The number of households in the US with both parents working has increased from 25% in 1968 to 48% in 2008 (Council of Economic Advisors, 2010). The increasing pressure for both parents to work is leading governments in the US and Europe to investigate ways to promote work-life balance, with again a shortage of evidence:

“A factor hindering a deeper understanding of the benefits and costs of flexibility is a lack of data on the prevalence of workplace flexibility and arrangements, and more research is needed on the mechanisms through which flexibility influences workers’ job satisfaction and firms’ profits to help policy makers and managers alike.”

Council of Economic Advisers, 2010

The third issue concerns the impact of WFH on changing urban land use and commuting. If people need not travel daily to and from work, the relative attractiveness of different locations changes, both as places to live and as locations for offices. Freeing large numbers of people from commuting to work could have substantial effects on traffic levels and pollution.

The efficacy of WFH as a management practice was what concerned CTrip International Corporation (“CTrip”) – China’s largest travel agency, with 16,000 employees and a \$3bn valuation on NASDAQ. Its leaders were interested in allowing its call-center employees to work from home because they perceived potential benefits from reducing office rental costs, which were becoming increasingly onerous due to rising rental rates at the Shanghai headquarters. But the executives

¹ In the E.U. in 2005, an average of 7.0% of employees worked from home at least a quarter of the time, and 1.7% did so almost all the time (EIROnline, 2010).

² See, for example, Bailey and Kurland (2002).

worried that allowing employees to work at home, away from the supervision of their shift managers, would lead to a large increase in shirking. The call center workforce was mainly younger employees, many of whom might well have struggled to remain focused working from home.

Given the uncertainty surrounding the effects of working from home in the research literature as well as in practice, the firm's leaders decided to run a randomized control trial. The authors assisted in designing the experiment and had complete access to the resulting data, as well as to data from surveys conducted by the firm. We have also conducted various surveys ourselves and numerous interviews with employees and managers.

In summary, CTrip decided to run a nine-month experiment on working from home. They asked the 996 employees in the airfare and hotel departments of the Shanghai call center whether they would be interested in working from home four days a week. Approximately half of the employees (508) were interested. Of these, 252 were qualified to take part in the experiment by virtue of having at least six months tenure, broadband access and a private room at home (in which they could work). After a lottery draw, those with even-numbered birthdays were selected to work at home while those with odd-numbered birthdates stayed in the office to act as the control group.

Throughout, the workers were organized in "teams,"³ each under a team leader, and all members of the team worked the same shifts. Assignments to teams remained unchanged through the experiment, so some members of a given team would be in the treatment group and others in the control. The home and office employees in each team had to work the same shift because they worked under a common team manager. The two groups also used the same IT equipment, faced the same work order flow from a common central server, and were compensated under the same pay system, which involved an element of individual performance pay. Hence, the only difference between the two groups was the location of work.⁴ This allows us to isolate the impact of working-from-home (flexi-place) versus other practices that are commonly bundled alongside this practice, such as flexible work hours (flexi-time). CTrip also had almost identical call center in Nan Tong which served as an additional control for any possible spillovers on to the control group.

We found five main results. First, the performance of the home workers went up dramatically, increasing by 13% over the nine months of the experiment. This improvement came mainly from a 9% increase in the number of minutes they worked during their shifts (i.e., the time they were logged in to take calls). This was due to a reduction in breaks and sick-days taken by the home workers. The remaining 4% improvement came from home workers increasing the number of calls per minute worked. In interviews, the workers attributed this gain to the quieter working conditions at home. Second, there were no spillovers to the rest of the group. Comparing the control group to similar workers in CTrip's other call center in Nan Tong we see no performance drop despite their losing the treatment lottery. Third, attrition fell sharply among the home workers, dropping by 50% versus the control group. Home workers also reported substantially higher work satisfaction and had more positive attitudinal survey outcomes. Fourth, the one downside of WFH appears to be that

³ There is no sense in which the call-center jobs involved "teamwork" and there was no group-based pay, but we stick with the term "team" because that was what CTrip called the work groups operating under a common "team" manager.

⁴ This of course had implications that were potentially relevant to the experiment. In particular, employees at home did not have on-going, immediate contact with the team leaders and they worked in a different environment than those in the office. We discuss these points more below.

conditional on performance it reduced rates of promotion by about 50%. Finally, at the end of the experiment, the firm estimated it saved about \$2,000 per year per employee working at home, leading it to offer the option to work from home to the entire firm. This allowed the treatment and control groups to re-select their working arrangements. Almost half of the treatment group changed their minds and returned to the office, while two thirds of the control group (who initially had all volunteered to work from home) decided to stay in the office. This selection led to a doubling of the impact from working at home, as workers with better performance at home remained at home while those performing poorly returned to the office.⁵

This experiment thus highlights the extensive learning by both the firm and employees around the adoption of a modern management practice like working from home. *Ex-ante*, both groups were unsure about its impact, and the 9-month experiment and subsequent roll-out process were essential for their ability to evaluate it. These gradual learning effects are one factor behind the slow adoption of modern management practices, and we see the results as similar to the adoption process for product innovations, like hybrid seed-corn as emphasized in Griliches' (1957) classic article.

This experiment is unique as the first randomized experiment on working from home. As such, it provides causal evidence to supplement the prior case-study and survey research. A second way in which this study is unusual is the fact that it is a randomized controlled experiment within a large firm. Third, we were granted exceptional access not only to data but also to CTrip management's thinking about the experiment and its results. This was because one of the co-authors of this paper - James Liang - the co-founder, first CEO and current Chairman of CTrip, was also a doctoral student at Stanford at the time. As a result, this paper benefits from unusually detailed insight into the rollout and adoption of a new management practice in a large, multinational firm.

As suggested, the paper connects to three strands of literature. First, there is a long literature that links the puzzling dispersion of productivity between firms to differences in management practices (see the literature from Walker (1887) to Leibenstein (1966) to Syversson (2011), Gibbons and Henderson (2012) and Bloom et al (2013)). This paper suggests that uncertainty about the efficacy of new practices may play a role in the slow diffusion of these practices, including those addressing issues of work-life balance such as the WFH under study. These practices have potentially large effects on measured productivity. For example, based on the methodology that is used to measure productivity using Census data (e.g. Foster, Haltiwanger and Krizan, 2000) CTrip would have increased productivity by 30% after introducing working from home, even before accounting for selection effects, because it increased output while cutting capital and labor inputs.

The second strand of literature is on the adoption of work-life balance practices. It is based primarily on case-studies and surveys across firms. These tend to show large positive associations of adoption with lower employee turnover and absenteeism and with higher productivity and profitability.⁶ However, these studies are hard to evaluate because of the non-randomized nature of the programs. This is both true in terms of the selection of firms into adopting working-from-home programs and also the selection of employees to work at home.

⁵ Strikingly, this is the same ratio of benefits from selection to total benefits from the intervention found in Lazear's well-known study of Safelite Auto Glass (Lazear 2000).

⁶ See, for example, see the surveys CEA (2010) and Lazear and Oyer (2011).

Finally, there is the connection to the urban economics literature. Reducing the frequency of commuting will reduce vehicle miles travelled, lowering emissions, but also reducing population centrality as people move out to the suburbs (Bento et al. 2005).⁷ WFH is also part of the wider impact of IT on firm-fragmentation arising from the increasing ease of long-distance communicating (e.g. Rossi-Hansberg et al., 2009). CTrip is now setting up regional offices to employ workers in lower-wage inland Chinese cities using the same working-from home technology they used in this experiment.

Section II describes the experiment in more detail, while section III presents the firm results, and section IV the impact on employees. Section V discusses the roll-out and finally section VI provides a set of concluding comments.

II. THE EXPERIMENT

II.A. The Company

Our experiment took place at CTrip International Corporation, a leading travel agency in China with operations also in Hong Kong and Taiwan. CTrip aggregates information on hotels and flights, makes reservations and obtains tickets for clients, and generates revenue through commissions from hotels, airlines and tour operators. The services provided by CTrip are comparable to Expedia, Orbitz or Travelocity, although, because of lower Internet penetration in China, CTrip did much more of its business on the telephone. CTrip was established in 1999, was quoted on NASDAQ in 2003, and was worth about \$5bn at the time of the experiment. It is the largest travel agent in China in terms of hotel nights and airline tickets booked, with over 50% market share in 2010. Exhibit A displays photos of the CTrip Shanghai office, a modern multi-story building that housed the call center in which the experiment took place as well as several other CTrip divisions and its top management team. The firm also operated a second call center in Nan Tong, a city about 1 hour away from Shanghai, which employed about 5,000 call center staff. Both locations operate in the same fashion, with the same equipment under the same procedures.

Call center representatives are organized into small teams of around 10 to 15 people (mean of 11.7 and median of 11), grouped by department and the type of work. There are four jobs in each of the two departments – hotel and airline – involved in the experiment. These are order takers, who answer customer calls, take orders, and enter them into the CTrip information system; order placers, who deal with the airlines and hotels and then notify the clients; order correctors, who resolve problems such as a flight being canceled, plus a night shift that both places and corrects orders. The members of a given team sit together in one area of the floor, typically occupying an entire aisle. Each call center representative works in a cubical with equipment including a computer, a telephone and a headset. When team members are ready to start work, they log on to CTrip's IT system and, in the case of order takers, client calls are automatically routed into their headsets. Order placers and order correctors also are allocated tasks automatically. The allocations between the two Shanghai and Nan Tong call centers are handled centrally, using a standard call queuing system. When employees want to take a break, they log out of the system. The team leaders patrol the aisles

⁷ For example, Jet Blue allows home-based call center employees to be up to 3 hours drive from the office as they only need to come into the office one day per month. Due to this lax requirement, many Jet Blue employees now live in rural areas outside Salt Lake City where the firm is headquartered.

to monitor employees' performance as well as help resolve issues with reservations, provide ongoing training and give emotional support to employees dealing with difficult clients.

The employees typically work 5 shifts a week, scheduled by the firm in advance. All members of a team work on the same schedule, so individuals do not choose their shifts. A team shares the same team leader, the same work schedule, as well as the same working area on the call center floor. Despite the team organization, the job of a CTrip call center representative is not based on teamwork, nor does individual performance of one team member affect the salaries of other team members.

Monthly salary is composed of a flat wage and a bonus. The flat wage depends on seniority, education and prior experience, averaging around ¥1300 per month. The bonus portion depends on monthly performance, and also averages about ¥1300 per month. The bonus is primarily a linear function of call and order volumes, but with small adjustments for call quality (penalties are applied for call quality scores below certain thresholds) and shift type (night shifts, for example, are paid a higher flat rate). Promotion to team-leader is also based on performance, so both salary and career concerns provide incentives for employees to perform well.

Since no other Chinese firm had moved to allowing home-working among its call center employees, there was no local precedent. In the US, the decision to allow employees in call centers to work from home varied across firms, even those within the same industry, suggesting a lack of any consensus on its impact. Meanwhile, the prior academic literature on call centers also offered limited guidance, being based on case studies of individual, firm-level interventions. Due to this uncertainty, management decided to run an experiment.

II.B. The Experimental Design

The experiment took place in the airfare and hotel booking departments in the Shanghai call center. The treatment in our experiment was to work 4 shifts a week at home and to work the 5th shift in the office on a fixed day of the week determined by the firm. Treatment employees still worked on the same schedule as their teammates because they had to work under the supervision of the team leader (who is always office-based), but they operated from home for 4 of their 5 shifts. For example, in a team the treatment employees might work from home from 9am to 5pm on Monday, Tuesday, Wednesday and Friday and in the office from 9am to 5pm on Thursday. The control employees from that team would work in the office from 9am to 5pm on all five days. Hence, the experiment only changed the location of work, not the type or the hours of work. Because all incoming phone-calls and work orders are distributed by central servers, the work flow was also identical between work and home locations. Home workers also used the same CTrip-provided computer terminals, communications equipment and software, faced the same pay structure and undertook the same training as the control group (although for the treatment employees this occurred only the day they were in the office).

Importantly, individual employees are not allowed to work overtime outside their team shift as it would require their team leader to supervise their work. Hence, entire teams could have their hours changed – for example all teams had their shifts increased during the week before Chinese New Year – but individuals were not able to work overtime on their own. Thus, eliminating commuting

time, which was 80 minutes a day for the average employee, did not permit the treatment group to work overtime and so this is not a factor directly driving the results.

Three factors other than location did differ between treatment and control. First, the treatment group's spending less time commuting meant that they would sometimes be able to take care of personal and family responsibilities without taking breaks or leaving early from work. As we will see, this appears to have had a significant effect. Second, the treatment workers did not have as much support from their team leaders, because technological limitations meant that they could not get face-to-face help from their supervisors. If anything, this presumably reduced the effectiveness of the treatment workers. Finally, the work environment differed between treatment and control. The former were working alone, typically reported to be in a quieter environment. This had some negative effects on willingness to work from home, but positive effects on productivity.

In early November 2010, employees in the airfare and hotel booking departments were informed of the working from home program. They all took an extensive survey on demographics, working conditions and their willingness to join the program. When inquired of their willingness to join the program, employees were not told the set of criterion that they would have to qualify in order to participate in the program. Employees who were both willing and qualified to join the program were recruited for the experiment. Of the 996 employees in the airfare and hotel booking departments, 508 (51%) volunteered for the experiment, with those with a more expensive and longer commute, with less tenure in the firm, with less education and with their own bedroom significantly more likely to want to work from home (see Table 1). Importantly, prior-performance (measured by the gross-wage given that almost 50% of salary is performance related pay) was not predictive for the take-up of working from home. This helped to assuage one concern of the firm that lower performing employees would be more tempted to work from home to avoid the direct supervision of their team-leaders.

Interestingly 49% of employees did not opt to work from home despite the considerable saving in commuting time and cost. The major reason given for this in later interviews was the loneliness of working from home and the lack of opportunities to socialize in the office and after work. Another reason was the negative impact for WFH on promotion, which as we discuss in section IV below, appears to have been a legitimate concern.

To qualify, an employee also needed to have tenure of at least 6 months, have broadband Internet at home to connect to the network, and an independent workspace at home during their shift (such as their own bedroom). Among the volunteers, 252 (50%) of the employees met the eligibility requirements and were recruited into the experiment.

The treatment and control groups were then determined from this group of 252 employees through a public lottery. Employees with an even birth date (a day ending 2, 4, 6, 8, etc.) were selected into the treatment and those with an odd birth date were in the control group. This selection of even birthdates into the treatment group was randomly determined by the Chairman, James Liang, by drawing a ping pong ball from an urn in a public ceremony one week prior to the experiment's start

date (see Exhibit B).⁸ Treatment group was notified and equipment is installed at each treatment participant's home the following week. Comparison between treatment and control groups shows there is no significant difference between the two groups across 18 characteristics other than the control group is more likely to be married and have children (see Appendix A2).

The experiment commenced on December 6, 2010 and lasted for 9 months, which employees knew in advance. All treatment employees had to remain at home for this period, even if they changed their mind and wanted to return to the office. On August 15, 2011, employees were notified that the experiment had been a success and CTrip would roll out the experiment to those who were qualified and wanted to work at home in the airfare and hotels departments on September 1st, 2011.

Throughout the experiment, employees were told the experiment would be evaluated to guide future company policies, but they did not learn the actual policy roll-out decision until August 15th. Because of the large scale of the experiment and the lack of dissemination of experimental results beyond the management team, prior to the roll-out decision, employees were uncertain about what that decision would be. Employees in the treatment group who wished to come back to work in the office full-time were only allowed to do so after August 15th, while control group had to stay in the office for the full duration of the experiment. Hence, the treatment and control assignments were fixed for the full 9 months, except for a few cases where equipment broke down or employees moved apartments and lost access to their own room.⁹

Figure 2 shows compliance with the experiment throughout the experimental period, and for one year after the general roll-out. During the experiment, the percentage of treatment group working at home hovered between 80% and 90%. Since compliance was not perfect, our estimators take even birth date status as the treatment status, so we estimate an intention to treat result. The impact on performance if compliance were perfect would be about 10% larger.

After the experiment, we see about 50% of the treatment group immediately decided to return to the office despite having to incur the financial and time costs of commuting. Strikingly, only about 35% of the control employees – who also all initially volunteered to work from home – actually moved home when they were allowed to do so. The main reason both gave for changing their mind in interviews were concerns over being isolated at home. Finally, we also see that about 10% of the workers who did not initially volunteer changed their minds after the experiment and decided to work from home.

It is worth noting that the firm's management was surprised by two of the findings. First, they were struck by how many employees changed their minds about working from home. More than 50% of the volunteer group and 10% of the non-volunteer group switched preferences after the 9-month experiment. Employees reported that after working from home for a few months they started to get lonely and wanted to return to the office.

⁸ It was important to have this draw in an open ceremony so that managers and employees could not complain of "favoritism" in the randomization process. The choice of odd/even birthdate was made to ensure the randomization was straightforward and transparent.

⁹ In all estimations, we use the even birthdate as the indicator for working-at-home so these individuals are treated as home workers. In a probit for actually working from home during the experiment none of the observables are significant, suggesting that returning to the office was effectively random. One reason is that the IT group policed this heavily to prevent employees fabricating stories to enable them to return to the office.

Second, despite the time and financial savings from not having to commute, more than half of the workers eligible to work at home decided to work in the office, suggesting they place a high value on social interactions at work (Hamermesh, 1990). This is particularly striking because, as we note below, we find no negative impact of home working on performance or quality of service.

II.C. Data Collection

CTrip had an extremely comprehensive central data collection system, as its founding team came from Oracle with extensive database software experience. The majority of data we used in our paper were directly extracted from the firms' central database, providing extremely high data accuracy. The data we collected can be categorized in 7 fields: performance, labor supply, attrition, promotions, reported employee work satisfaction, detailed demographic information, and survey information on attitudes towards the program.

Performance measures vary by the broad type of workers – the 134 order takers and the 118 order placers, order correctors, and night shift workers (details in Appendix 1). Order takers' key performance measures are the number of phone calls answered and number of orders taken. The key measures for the other three groups are the numbers of different types of calls made. For order takers, we can also accurately measure time spent working in terms of minutes on the phone because phone calls and call lengths are logged in the central CTrip database. The firm uses these measures to monitor the work of their employees. We also calculate phone calls answered per minute on the phone as a measure of labor productivity for these workers.

We have daily performance measures of all employees in the airfare and hotel booking departments from January 1st, 2010 onwards. We also have daily minutes on the phone for order takers and daily records of hours of leave for the airfare department. We also collect data on different types of promotion by September 2012, almost 2 years after the experiment commenced. The firm ran internal surveys of the employees during the experiment on work exhaustion and, positive and negative attitudes (see details in Appendix A3). Finally, we conducted two rounds of surveys, in November 2010 and August 2011, to collect detailed information on all the employees in the two departments including basic demographics, income, and attitudes toward the program, and ran extensive interviews and focus groups with the employees.

III. IMPACT ON THE FIRM

We analyzed the effect of home-working both in terms of its impact on the firm, which we cover in this section, and the impact on the employees, which we cover in the next section.

III.A. Individual Employee Performance

We first estimated the intention to treat effect on weekly employee performance on pre- and during experimental period data via equation (1)

$$Employee\ Performance_{i,t} = \alpha Treat_i \times Experiment_t + \beta_t + \gamma_i + \epsilon_{i,t} \quad (1)$$

where $Treat_i$ is a dummy variable that equals 1 if an individual belongs to the treatment group defined by having an even-numbered birthday; $Experiment_t$ is a dummy variable that equals 1 for the experimental period December 6th to August 15th; and $Employee\ Performance_{i,t}$ is one of the key measures of work performance, including an overall performance z-score measure, log of weekly phone calls answered, log of phone calls answered per minute on the phone, and log of weekly sum of minutes on the phone. Finally, β_t reflects a full set of weekly time dummies to account for seasonal variation in travel demand, such as the World Expo in 2010 and the Chinese New Year, and γ_i reflect a full set of individual fixed effects.

To make performance of different types of workers comparable, we use performance z-scores. For each individual we subtract the pre-experiment mean for their worker type, and divide by the pre-experiment standard deviation for their worker type. Hence, this normalized z-score measure has a mean 0 and standard deviation 1 across all employees within each type of worker during the pre-experiment period.

In column (1) of Table 2, overall performance of the treatment group is found to be 0.232 standard deviations higher than the control group after the experiment started, significant at the 1% level. If we limit the sample to the 137 order-takers we can use phone calls answered as the key performance measure. The z-scores of phone calls account for different volume and average length of phone calls in two departments. In column (2), we look at only the phone calls performance measure and find this is 0.247 standard-deviations higher for the treatment group. In column (3), we look at the log of phone calls and find a coefficient of 0.120, meaning that treatment employees were making 13% (noting that $13\% = \exp(0.120)$) more phone calls per week.

We can also see these results in Figure 3 where we plot the raw number of phone calls per week for the treatment and control groups from Jan 1st 2010 until the end of the experiment in August 15th 2011. Before the experiment started, the treatment group trends closely together with the control group, both of which bounce around due to seasonal fluctuations in demand. But once the experiment began, the treatment group started to outperform the control group, answering about 40 more phonecalls per person per week.

Interestingly, the difference in performance was greatest during the middle of the experiment, from about 2 to 6 months. The reason for the 2 month initial performance improvement was installation and learning effects. It took several weeks for all the IT and logistical bugs to be addressed. The gradual decline in the performance gap from 6 months onwards reflects two trends. First, poorly performing employees in the control group were more likely to quit than those in the treatment group (see section IV.B and Table 8), boosting the control group. Second, from focus groups and interviews we learned that some employees in treatment groups felt lonely working at home after a few months and wanted to return to the office (but because of the experiment could not return until 9 months), potentially impacting their motivation.

Figure 4 plots the cross-sectional distribution of performance for treatment and control groups at 3 months into the experiment, displaying a broad distributional improvement from working-from-home (rather than the results being driven by a few outliers).

We further decomposed the difference in performance observed in column (3) into phone calls answered per minute on the phone (a measure of productivity), and minutes on the phone (a measure of high-frequency labor supply). In column (4), we found treatment employees were making 3.5% (note that $3.5\% = \exp(0.034)$) more phone calls per minute, which the employees working from home identified as resulting from home being quieter than the office. They told us this meant it was easier to hear the customers, so they did not have to ask them to repeat themselves as often and could process the information more quickly. This suggestion matches the psychology literature which has shown that background office noise can reduce cognitive performance (see, for example, Banbury and Berry, 1998).

The largest factor increasing the home-workers performance is that, as shown in column (5), they worked 9.0% ($9.0\% = \exp(0.086)$) more minutes per day. This was despite the fact that home and office workers both worked the same nominal shift – for example, 9am to 5pm on Monday to Friday – as members of both groups worked in the same team under the same team manager. The reason home-workers could increase minutes on the phone was that, within their shift, they were logged on and available to take calls for more time, meaning they were taking less time off during their shift.

III.B. Individual Employee Labor Supply

In Table 3, we investigate the factors driving this increase in minutes worked within each shift. Because we have accurate records of hours of leave from the airfare booking department only, we limit the sample further to the 89 order takers in the airfare department. Column (1) repeats the results from the final column of Table 2, while Column (2) of Table 3 shows that these order takers show a very similar increase to the full group.

Columns (3)-(5) break this difference in minutes on the phone down into three pieces. In column (3), we look at whether treatment workers spent more minutes on the phone per hour they spent logged into their computer¹⁰, column (4) looks at whether they spent more logged in per day worked, and column (5) looks at whether they worked for more days.

What we found is that in column (3), there is no difference between the number of minutes on the phone while logged in for the treatment and control employees. This is not surprising because both groups operated using the same call routing server and so are on the same queuing system.¹¹

Column (4) shows that about two-thirds of the difference in the time on the phone was accounted for by taking less time off per day worked. This is because: (a) they started work more punctually, a phenomenon which they attributed to avoiding the effects of events that disrupted commuting like the heavy snow in Shanghai in February 2011¹²; and (b) they could schedule personal matters, like doctor's appointments, in the time they saved by not commuting rather than having to leave early. Finally, in column (5) we see that the other third of the difference in time worked between

¹⁰ When workers arrive at their desk to start work they need to log into their computer. Once they have done this they can start receiving and making calls. The firm both records the times they are logged in, and the times they are actively on the phone (rather than, for example, waiting for calls to arrive or performing follow-up work between calls).

¹¹ Moreover, it shows that home workers are not picking busier times to log in to the system (i.e. they are not timing their breaks to coincide with quiet periods when demand is lower). We thank Wouter Dessein for pointing this out.

¹² CTrip is strict with punctuality of its workers. If a worker comes to work late, it is recorded as taking time off.

treatment and control was explained by treatment employees working more days because they took fewer sick-days. Employees explained this was because they continued to work at home when they felt somewhat ill but would not have felt up to commuting into work.

III.C. Quality, Spillovers

One question is whether quality of the service was compromised for the increase in output in the treatment group. We constructed two quality measures: conversion rates and weekly recording scores. Conversion rates were calculated as the percentage of phone calls answered that resulted in orders, while the weekly recording scores came from the 1% of phone-calls that are randomly evaluated by an external monitoring team. In summary (with the full details in table A4 in the appendix), we find no impact of working from home on call quality using either measure.

Another related question is whether the improvement associated with working from home came from an improvement in the treatment group or from a deterioration in the control group. Perhaps the gap between treatment and control was caused not by the treatment group performing better but by the control group performing worse after they “lost” the randomization lottery. The group winning the treatment lottery saved themselves 9 months of commuting time and costs, a substantial gain worth about 17% of their salary, evaluated at their CTrip wage rate.¹³

We collected data on two other “quasi” control groups to answer this question. The first group is the employees who would be eligible in the Nan Tong call center. This was CTrip’s other large call center, located in Nan Tong, a city one hours drive from Shanghai. This call center also had airfare and hotel departments, and calls were allocated across the Shanghai and Nan-Tong call centers randomly from the same central server. The second group was the 253 eligible employees who did not volunteer to participate in the WFH experiment in the Shanghai call center. These were the individuals that were eligible for the experiment (own room, 6+ months of tenure and broadband), but did not want to work from home. We think these two groups are comparable to the treatment and control groups for two reasons. First, all four groups face the same demand for their service. Second, they all meet the requirements for eligibility to participate in the experiment.

Comparisons of these alternative control groups are also reported in Table 4. Results in the top panel of Table 4 compare the treatment and control groups to Nan Tong, revealing no significant differences in overall performance between Nan Tong and the control group. The bottom panel compares treatment and control group to the eligible non-experimental group in Shanghai. Again, we find no significant difference between the control group and the eligible non-experimental group. These results suggest that the gap between the treatment and control group reflects an improvement in the performance of the treatment group rather than any deterioration of the control group. That is, although the control group and the treatment group work in the same team, we find – perhaps surprisingly – no evidence of the control group’s being discouraged by losing the working-from-home lottery.¹⁴

¹³ The average employee makes about \$100 per week for a 40 hour week. The commuting time is 40 minutes each way and the out of pocket cost \$0.5 on average. Hence, the saving in time is about \$13 a week in time costs and about \$4 per week in out of pocket costs.

¹⁴ Alternatively, we could compare the performance of treatment and control group during the days treatment group work at home with the days they work in the office. Unfortunately, they were not comparable because team leader often scheduled weekly team meetings during the day treatment group worked in the office.

III.D. Potential Hawthorne Effects

Another explanation for the superior performance of the treatment group are Hawthorne effects, which suggests that the employees were motivated by the experiment, possibly deliberately so that the firm would roll out WFH permanently. We should note three things, however, that make this appear unlikely. First, there were 121 employees working from home, so each individual employee has little impact on, and thus little incentive to try to manipulate, the overall evaluation of the experiment. Second, the gap between home-based employees and office-based employees widened after the experiment ended as we shall see in the next section. Finally, the firm was itself so convinced that the success of the experiment was not due to Hawthorne type effects that it rolled out WFH to the entire airfare and hotel divisions.

III.D. Post-Experiment Selection

In August 2011, the management estimated that each working from home employee saved CTrip about \$2,000, so they decided to immediately roll out the option to work from home to the entire hotel and airfare departments. Employees in these departments were notified that the experiment had ended and they were entitled to choose their location of work conditioning on being eligible, so control employees who still wished to could move home, and treatment employees that wanted to return to the office could do so.¹⁵

As shown in Figure 5 – which plots the difference in normalized phone-calls between home and office workers – post-experiment selection substantially increased the performance impact from working from home. The differential increase in performance z-scores (versus the pre-experiment baseline) from home-working was about 50 calls per week, rising to about 100 calls per week within 9 months after the experiment. This is also evaluated in Table 5 which estimates the performance impact of working-from home during and after the experiment. We see after the experiment the average z-score rises by 26% in column (2) and the average number of phone calls by 78% in column (5)¹⁶. Once you control for quits by using a balanced panel we find even larger increases in performance, of 34% for z-scores in column (3) and 93% in column (6). The fact that the balanced panel results show an even larger increase in performance shows that it is sorting of employees between home and the office rather than differential attrition that is driving the improvement from WFH home during the post-experiment roll-out period.

This sorting is driven by treatment workers who had performed relatively badly at home returning to the office. This is shown in Table 6, top panel, columns (1) to (4), which runs probits on whether a treatment worker returns to the office. The results show that treatment workers who performed relatively worse at home versus the office returned to the office. This was despite the fact that all treatment workers had initially volunteered to work from home, suggesting that many of them subsequently discovered home working was not as attractive as they initially believed it would be. In the bottom panel we find no significant evidence of differential sorting in the office-based control group, suggesting potentially less learning occurred in this group who did not get to experiment with home based working.

¹⁵ Treatment group employees who were working at home and wanted to come back to work in the office full-time were allowed to come back on September 1st, 2011, 2 weeks after the announcement of the roll-out. Control group employees who wanted to work at home started to do so gradually starting at the beginning of November 2011.

¹⁶ $78\% = (\exp(0.208) - 1) / (\exp(0.122) - 1) - 1$

IV. IMPACT ON THE EMPLOYEES

III.A. Employees' self-reported outcomes

CTrip management was also interested in how employee self-reported wellbeing was affected by the program. They thus ran two sets of surveys: a satisfaction survey and a work attitude survey. Details of survey questions and methodology are listed in Appendix A3. In summary, these were standard employee satisfaction tests developed by Berkeley psychologists Christina Maslach and Susan Jackson in the 1970s. The satisfaction survey was conducted five times throughout the experimental period: once in early November before the randomization took place and four times after the experiment had started. The first three columns of Table 7 show three different satisfaction measures. The treatment group reported no difference in satisfaction levels from the control group at the first survey, but the treatment group reported statistically significantly higher satisfaction once the experiment began.

The work attitude survey was conducted every week. The first weekly survey was conducted in late November 2010, before the experiment began but after the randomization had occurred. So individuals had been informed of their sorting into the treatment or control groups. Interestingly, the treatment group already reported higher positive attitude (significant at the 10% level), less negative attitude and less exhaustion from work. This group had yet to move home, so this difference was presumably due to the control group's learning they lost the working from home lottery while the treatment group learned they had won, highlighting the importance of comparing our treatment groups with other control groups like Nan Tong and the eligible non-experiment group. After starting the experiment, the gap between the treatment and control group rose further, so that the treatment group reported statistically significantly higher positive attitude and less work exhaustion.

IV.B. Attrition

One of the key initial reasons CTrip was interested in running the experiment was to lower their rate of employee attrition because it takes them about 6 weeks salary to train each new employee. The turnover rate among CTrip call center representatives had historically hovered around 50% per year, which was typical of the call center industry in China¹⁷. Figure 6 plots the cumulative attrition rate of treatment and control group separately over the experimental period. Shortly after the commencement of the experiment, cumulative attrition rates diverged between the two groups and the difference is statistically significant. By the end of the experiment, the total attrition rate in the treatment group (17%) was less than half of that in the control group (35%) during the 9 months period.¹⁸

We further tested whether selective attrition existed by running probit regressions in Table 8. The dependent variable is whether an employee quit the job during the experimental period between December 6th 2010 and August 15th 2011. Column (1) confirms the finding in Figure 6, that treatment employees rate of attrition was about half that of the control group. In column (2), we

¹⁷ 2010 Report on Chinese Call Center Operation and Management. Note that CTrip could in principle fire employees, but this was rare and no employees in these two divisions were fired over this period as far as we are aware.

¹⁸ Of course these figures are dependent on the market circumstance: the fact that no other call-centers offered home-working in Shanghai is likely to make this particularly effective in reducing attrition.

tested whether employees with worse performance were more likely to leave the firm from the treatment group compared to the control group. To measure performance we used the average individual weekly performance z-scores during the experimental period (from December 6th 2010 until their quit date or August 15th 2011) or for employees that quit their performance until their last full week. We found that low performers were significantly more likely to quit, particularly those in the control group.

In columns (3) and (4), we estimated the impact of experimental period performance on quitting in the treatment and control groups separately and we found a large and strongly significant impact in the control group and smaller weakly significant impact in the treatment group. From interviewing the employees, we heard that control group employees who underperformed tended to quit for other similar call-center or office jobs, which were easy to find. Treatment employees, however, were much less likely to quit because no other home-working jobs existed, substantially reducing selection from the treatment group.

This differential attrition, of course, also raises the question of whether our estimated impact of WFH is biased. Specifically, the concern is that the estimated impact is mainly driven by differential attrition. We note that this is unlikely because in Table 7 we see that employees with worse performance in the control group are more likely to quit. This suggests that the observed control group performance level is larger than it would have been, hence generating a smaller performance gap between treatment and control than if no attrition had occurred. In other words, our estimated treatment effect is actually biased downwards.¹⁹

To address this issue, we use the Lee (2008) bounds estimator. This provides upper and lower bounds on the differential selection on performance across groups, assuming that attrition is monotonically driven by the performance variable. This allows us to generate two bounds – the upper bound that assumes that the extra attrition in the control group is based on a negative correlation between performance (as we saw in Table 7) while the lower bound assumes a positive correlation (the reverse of what we see in Table 7, but included for completeness). These bounds are shown in Appendix Figure A1. We see that the upper bound lies above the actual treatment-control estimated impact, suggesting that the actual treatment effect on attrition is, if anything, larger than we estimated, because the attrition of the worst performers from the control group biases our results down. Hence, even our large estimated positive impact of WFH is likely to be underestimated by attrition bias.

IV.C. Promotions

One possible negative effect from working at home is that long-run career performance could be damaged by less on-the-job training from team leaders and less “face-time” in the office. To investigate this, we collected promotions data on the 252-employee experimental sample. We consider promotion either being promoted to team leader or being moved to a more advanced function (i.e., switch to quality-control team) which often occurs shortly before promotion. Both would require a formal evaluation by team leaders and departmental managers. During the period from the start of the experiment in December 2010 until the end of September 2012, a total of 17

¹⁹ Of course, to the extent that lower performing employees are leaving the firm is less concerned with their loss, although, the firm still would have rather kept them because of the substantial costs of recruitment and training.

employees from the treatment group received promotions or more advanced job functions and 23 from the control group.

In Table 9 we show the results from running probit regressions on performance before the experiment, performance during the experiment, as well as some demographic controls. We see in column (1) that working from home has no overall impact on promotion. However, we see in column (2) that once you control for their performance during the experimental period WFH has a negative and weakly significant impact on performance. It appears that WFH improves performance – increasing the probability of promotion – but this is offset by a roughly equal negative effect from being home based. In column (3) we add demographics and the story is very similar. Finally, in column (4) we interact WFH with the performance measure and find a negative coefficient in addition to a negative WFH levels effect.

One story that is consistent with this is that home-based employees are “out of sight, out of mind”. As a result managers do not notice their performance as much and are less likely to promote them. We heard anecdotal evidence for this from employees and managers during focus groups and interviews, and it was one factor that led some employees to return to the office to avoid what they perceived as a WFH promotion ‘discrimination’ penalty. Another possibility is that WFH employees lack opportunities to develop the interpersonal skills to succeed in managerial jobs and therefore are less likely to be promoted. Either way, it is clear that WFH comes with a promotion penalty.

V. PROFIT, PRODUCTIVITY AND FIRM LEARNING

One of the most interesting aspects of the experiment was the learning process for both the firm and the individual employees on the costs and benefits of working from home. Both groups were initially unsure about its impact, because no other Chinese call centers had ever offered this option before. However, we were able to monitor both management’s and employees’ learning over the course of the experiment because of our extensive access to the CTrip’s management team and frequent employee surveys and interviews. Before discussing this we first present the estimated impacts on firm profits and productivity from allowing employees working-from home.

V.A. Profit and productivity impact

The firm saw working from home as a way to save on office costs, but was worried that employees would shirk at home or that call quality would decline due to multi-tasking on other activities which are prohibited in the office like playing computer games or watching TV. While managers had previously been allowed to work from home on an ad hoc basis, no non-managerial level employees had been allowed to home work. The research literature provided very little guidance on what might happen, with what little evidence there was suggesting that typically routine jobs are if anything less effectively carried at home.²⁰

²⁰ For example, Dutcher (2012) runs lab experiments on routine and non-routine tasks with and without remote monitoring, and finds the more routine ones are negatively impacted by mimicking a home-based environment. He conjectures that the lack of peer and manager effects, which have been shown to be important in low level tasks in

Running the experiment revealed, however, that working from home actually generated an improvement in employee performance, worth about \$375 per employee per year (evaluated at the 13% performance improvement from the Table 3). In addition, they estimated office cost savings of about \$1250 per employee and reduced turnover savings of about \$400 per employee per year. Hence, given the saving of about \$2000 per employee, the firm rolled the program out in August 2011, accompanied by an aggressive poster campaign to persuade employees to take up the home working option.

A related question is what was the impact on total factor productivity (TFP)? We estimate TFP would increase by about 30% from moving every employee home, using the methodology adopted on US Census data by, for example, papers like Foster, Haltiwanger and Syverson (2008) and Syverson (2011).

This 30% rise in measured TFP comes from three sources. First, output (as measured by the number phonecalls) increased 13% from working from home. While 9.5% of this increase comes from employees working more hours, this increase in attendance would not be measured in US Census survey data, since this collects information on shift-hours (i.e. 40 hours per week), not actual hours worked. Second, the reduction in attrition from 50% to 25% would reduce steady-state labor hours lost to training by 3%, since new employees need 6 weeks of training.²¹ Finally, the capital per employee is comprised of about \$5k of desktop IT equipment, \$10k of central IT equipment (servers and the network) and \$24k of office space (total office capital divided by the total number of employees). Moving employees home for 4 days a week reduces the office space required by 80%²², although it increases the desktop IT requirement by 20% (equipment lies idle at home for 1 day a week). On aggregate this reduces capital by 48%. Assuming a coefficient of 1/3 on capital and 2/3 on labor this yields an estimated TFP increase of 30%. Given that the cross-sectional standard-deviation of TFP reported in Foster, Haltiwanger and Syverson (2008) is 26%, this highlights how differences in the adoption of WFH across firms could potentially account for quite large differences in measured productivity.

V.B. Firm learning

The firm learned four important results from running the formal experiment versus the non-randomized pilot that they had initially been considering. First, they learned that WFH improves performance. Without running a formal experiment, their view was that they could have interpreted the drop in treatment performance shown in Figure 3 as a negative treatment effect. The period of the experiment (December 2010 to August 2011) coincided with a business slow-down for CTrip due to a combination of the (predicted) end of Shanghai Expo 2010 and an (unpredicted) increase in competition from other travel agencies. As a result, the *difference* in performance for the treatment group relative to their pre-experiment baseline was negative, and is only positive when evaluated as

field environments by Falk and Ichino (2006), Bandiera, Barankay and Rasul (2007) and Mas and Moretti (2009), could explain this.

²¹ Training costs 6 weeks of employee time, which given a 50% rate of employee turnover, has to be amortized over 2 years, meaning in steady state about 6% of employee output is lost from training. Hence, reducing attrition by 50% reduces training time by 3%. This ignores the labor savings from a smaller central recruitment team, as it only focuses on individuals lost output from training.

²² Note that since employees hot-desk (share desks), and because teams come in on different days of the week, the firm can reduce the numbers of desks in line with the reduction in the number of days spent working in the office.

a *difference of differences* against the control group. This highlights the importance of having a well matched (ideally randomized) control group to strip out these kinds of seasonal and demand effects.

Second, *ex ante* there was very little discussion of selection effects on employee performance, but by running the experiment and then rolling this out it became clear that allowing employee choice generated a far greater effect than requiring work from home. The impact of working from home is positive, on average, but appears to have a large variance, so that employee choice leads to a much higher effect, as shown in Figure 5.

Third, having the large sample of treatment and control employees allowed the firm to evaluate the impact on different types of employees. Somewhat surprisingly, they found a very homogeneous impact across all types of employees. For example, in Figure 7, we plot the impact on the top half of the treatment versus control distribution and the bottom half of the treatment vs control distribution. To calculate this, both groups were split in half by the pre-experiment median performance and then compared. What we see is a similar improvement in performance for both groups. CTrip's *ex ante* expectation was that the bottom half of employees were the less motivated ones, and they would perform far worse at home. Table A5 shows a similar result that the impact of WFH was homogeneous across a range of other characteristics, including gender, commute time, age, prior experience and living arrangements. In other words, in the current context observed characteristics are not good predictors of who will benefit from work from home more *ex ante*. These results have led the firm to offer WFH to all employee groups going forwards rather than any selected sub-samples (such as high-performers), which they were initially intending to target.

Finally, they were surprised by the dramatic drop in attrition that highlighted how many of their employees valued working-from home. They anticipated a reduction, but nothing like the 50% cut they observed.

V.C. Employees' learning

One direct measure of the extent of employee learning is the number of employees who changed their minds about working from home. Figure 2 shows that after the experiment about 50% of the initial treatment and control volunteers changed their minds and decided to work in the office after the end of the experiment, while 10% of the initial non-volunteer group opted to work from home.

We also designed a survey to inquire into employees' evolving views toward the Program from across all 996 airfare and hotel department employees. We administered the same survey with the help of the CTrip management in November 2010 and August 2011. Employees were asked specifically whether they were interested in participating in the Work-at-Home Program if they were eligible. They could choose from three answers: "yes", "no" or "undecided". We find of the 568 employees that took part in both surveys, that only 303 (53%) maintained their views, while the remaining 47% changed their minds. Of those, 24% went from "yes" or "undecided" to "no", while 12% went from "no" or "undecided" to "yes", with the remainder switching from "yes" or "no" into "undecided".

In follow-up interviews, most of the interviewed employees who had decided they no longer wanted to work from home cited social reasons. Another group who had thought working from home would be attractive found that it was troublesome for the people with whom they lived (often

parents), especially if they were called to work outside normal business hours. In reverse, a number of employees saw the success of their peers that worked from home and switched in favor of this.

V.D. Why did the firm not introduce working from home before?

Finally, one question that arises is why CTrip (or any other similar firm) did not introduce working from home earlier given it was highly profitable? From extensive discussion with the senior management team there seemed to be two reasons.

First, there was the classic free-rider problem that arises with all forms of process innovation where the absence of intellectual property makes it hard to prevent imitation. CTrip believed that the private benefits of home-based working would be short-lived (if it was successful) as rivals would copy the scheme and use it to drive down commission margins in the travel agent market, while the costs of experimentation would be borne entirely by the firm. Hence, they viewed themselves as paying the full cost of experimentation but only capturing part of the benefits because imitation due to knowledge spillovers.

Second, within CTrip the senior management had incentives that provided limited upside and extensive downside from the experimental outcomes. Senior managers were primarily motivated by career concerns with limited bonus or equity compensation. As a result their incentives to experiment were muted – they gained little from a successful experiment, but risked major career damage if the experiment was unsuccessful. James Liang, the Chairman and co-founder, had more balanced incentives to promote the experiment since he owned extensive CTrip equity and had no firm-level career concerns.

Both factors – the threat of imitation and risk-aversion from the career concerns of senior managers - are likely to represent forces deterring process innovations in most large firms, and highlights how these could lead to an under experimentation in management and operational practices.

VI. CONCLUSIONS

The frequency of working from home has been rising rapidly in the U.S. and Europe, but there is uncertainty and skepticism over the effectiveness of this, highlighted by phrases like “shirking from home”. We report the results of the first randomized experiment on working from home, run in a 16,000 employee, NASDAQ-listed Chinese firm, CTrip. Employees who volunteered to work from home were randomized by even/odd birth-date into a treatment group who worked from home four days a week for nine months and a control group who were in the office all five days of the work week. We found a highly significant 13% increase in performance from home-working, of which 9% was from working more minutes of their shift period (fewer breaks and sick days) and 4% from higher performance per minute. We found no negative spillovers onto workers who stayed in the office. Home workers also reported substantially higher work satisfaction and psychological attitude scores, and their job attrition rates fell by over 50%. Further, when the experiment ended and workers were allowed to choose whether to work at home or in the office, selection effects almost doubled the gains in performance.

This experiment highlights how complex the process of learning about new management practices is. For CTrip, having no precedent in terms of similar Chinese firms that had adopted working from home for their employees led them to run this extensive field experiment. Given their success, other firms are now likely to copy this practice, generating the type of gradual adoption of a new management practices that Griliches (1957) highlighted. More generally, given the large impact of this practice on employee performance – a \$2000 per employee reduction in costs and a 30% increase in TFP – this also provides a management-practice based explanation for heterogeneous firm performance.

While our results suggest a promising future for working from home, we should note that several distinctive factors at CTrip that contributed to the success of the experiment and the implementation of the practice. First, the job of a call center employee is particularly suitable for telecommuting. It requires neither teamwork nor in-person face time. Quantity and quality of performance can be easily quantified and evaluated. The link between effort and performance is direct. These conditions apply to a range of service jobs such as sales, IT support, and secretarial assistance, but they are far from universal. Second, the firm can closely monitor the performance and labor supply of the employees thanks to its extensive centralized database. Team leaders may generate a report from the database of the performance of the team members daily and efficiently detect problems in individual employees' performance. Third, the extent of WFH was moderate, so that it did not require a significant reorganization at the workplace. Because the participation rate ranged between 15% and 25%, team leaders continued to supervise their teams with a mix of home and office workers without major reshuffling of team formation.

While these features arguably favored successful implementation of working from home at CTrip, we believe the practice is worth further exploration. After all, much of the research for this paper and its writing were done by the authors working from home.

V. BIBLIOGRAPHY

- Bailey, Diane and Nancy Kurland. 2002. A review of the telework research: findings, new directions and lessons for the study of modern work. *Journal of Organizational Behavior* 23:283 – 400.
- Banbury, Simon and Berry, Dianne. (1998). “Disruption of office related tasks by speech and office noise”. *British Journal of Psychology*, 89, 494-517.
- Bandiera, Oriana, Barankay, Iwan and Imran Rasul (2007), “Incentives for Managers and Inequality Among Workers: Evidence from a Firm Level Experiment”, *Quarterly Journal of Economics*, 122(2), 729-773.
- Bento, Antonio, Cropper, Maureen, Ahmed Mobarak and "The Impact of Urban Spatial Structure on Travel Demand in the United States" (with A. Bento, M. Cropper, and K. Vinha), *Review of Economics and Statistics*, Vol. 87, No. 3, August, 2005
- Bloom, Nick, Benn Eifert, Aprajit Mahajan, David McKenzie and John Roberts, (2013), “Does Management Matter: Evidence from India,” *Quarterly Journal of Economics*, 1-51.
- Bloom, Nick and Van Reenen, John, (2011), “Human resources and management practices”, *Handbook of Labor Economics*, Volume 4, edited by Orley Ashenfelter and David Card.
- Bloom, Nick, Tobias Kretschmer and John Van Reenen, (2009), “Work-life Balance, Management Practices and Productivity”, in *International Differences in the Business Practice and Productivity of Firms*, Richard Freeman and Kathy Shaw (eds.). Chicago: University of Chicago Press.
- Clark and Tellegen (1988), “Development and validation of brief measures of positive and negative affect: the PANAS scales.”, *Journal of Personal Social Psychology*. 1988 Jun;54(6):1063-70.
- Council of Economic Advisors (2010), “Work-life balance and the economics of workplace flexibility”, <http://www.whitehouse.gov/files/documents/100331-cea-economics-workplace-flexibility.pdf>
- Dutcher, Glenn (2012), “The effects of telecommuting on productivity: An experimental examination. The role of dull and creative tasks”, *Journal of Economic Behavior and Organization*, April 2012.
- EIROnline (European Industrial Relations Observatory on-line) (2010), “Telework in the European Union,” <http://www.eurofound.europa.eu/eiro/studies/tn0910050s/tn0910050s.htm>
- Foster, L., Haltiwanger, J. and Krizan, C.J. (2000), "Aggregate Productivity Growth: Lessons from Microeconomic Evidence", *New Developments in Productivity Analysis*, NBER, University of Chicago Press.
- Foster, Lucia, John Haltiwanger and Chad Syverson (2008), “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?” *American Economic Review*, 98(1), 394-425
- Gibbons, Robert, and Rebecca Henderson (2012), “What Do Managers Do? Persistent Performance Differences in Seemingly Similar Enterprises,” in Robert Gibbons and John Roberts, eds. *The Handbook of Organizational Economics*, Princeton, NJ: Princeton University Press.
- Griliches, Zvi (1957), “Hybrid Corn: An Exploration in the Economics of Technological Change”, *Econometrica*, volume 25 (4), pp. 501-522.
- Hamermesh, Daniel (1990), "Shirking or Productive Schmoozing: Wages and the Allocation of Time at Work," *Industrial and Labor Relations Review*, January

- Falk, Armin and Andrea Ichino. (2006). "Clean Evidence on Peer Effects". *Journal of Labor Economics* 24(1), pp:39 – 57
- Lazear, Edward, 2000. 'Performance Pay and Productivity', *American Economic Review* 90(5), 1346-1361.
- Lazear, Edward, and Paul Oyer (2012), "Personnel Economics," in Robert Gibbons and John Roberts, eds. *Handbook of Organizational Economics*, Princeton University Press, forthcoming.
- Lee, David (2008), "Training, wages and sample selection: estimating sharp bounds on treatment effects". *Review of Economic Studies*, 76(3), pp. 1071-1102.
- Leibenstein, H. (1966) "Allocative Efficiency vs." X-efficiency," *American Economic Review*, 56(3): 392-415
- Mas, Alexandre and Enrico Moretti (2009), "Peers at Work", *American Economic Review* 99 no. 1:112 - 145
- Maslach, C., & Jackson, S.E. (1981). *Maslach Burnout Inventory*. Research edition. Palo Alto, CA: Consulting Psychologist Press.
- Mateyka, Petr J., Melanie Rapino and Liana Christin Landivar. (2012) "Home-Based Workers in the United States: 2010," United States Census Bureau, Current Population Reports.
- Oettinger, Gerald (2012), "The Incidence and Wage Consequences of Home-Based Work in the United States, 1980-2000", *Journal of Human Resources* forthcoming
- Rossi Hansberg, Esteban, Sarte, Pierre-Daniel and Owens, Raymond (2009), "Firm fragmentation and urban patterns", *International Economic Review*, Vol 50(1), pp. 143-186.
- Syversson, Chad (2011), "What Determines Productivity at the Micro Level?", forthcoming *Journal of Economic Literature*.

DATA APPENDIX

Appendix A1: Table for different types of workers and their key performance measures

Types of Workers	Department	Key Performance Measures	Number of Workers
Order Takers	Airfare	Phone Calls Answered	89
	Hotel	Orders Taken	46
Order Placers	Airfare	Notifications Sent	45
	Hotel	Reservation Phone Calls Made	25
Order Correctors	Hotel	Orders Corrected	36
Night Shift Workers	Hotel	Reservation Phone Calls Made Orders Corrected	11

In the analysis, the Order Takers, Order Correctors and Night Shift Workers were grouped together.

Appendix A2: Comparison between treatment and control groups

	Treatment	Control	SD(treatment+control)	t-stat
Number	131	121		
Age	24.44	24.36	3.53	0.16
Men	0.47	0.46	0.50	0.04
Secondary technical school	0.46	0.48	0.50	-0.34
High School	0.18	0.13	0.36	0.95
Tertiary	0.35	0.36	0.48	-0.07
University	0.02	0.03	0.15	-0.92
Prior experience (months)	18.96	16.54	25.78	0.75
Tenure (months)	26.14	28.17	21.87	-0.73
Married	0.22	0.33	0.45	-1.95
Children	0.11	0.24	0.38	-2.64
Age of the child	0.53	0.69	1.92	-0.69
Rent apartment	0.25	0.20	0.42	0.87
Cost of commute (Yuan)	7.91	8.41	6.96	-0.57
Internet	0.99	0.99	0.09	0.06
Bedroom	0.97	0.99	0.14	-1.26
Base wage (Yuan)	1539.86	1561.79	160.77	-1.08
Bonus (Yuan)	1030.90	1075.76	626.37	-0.57
Gross wage (Yuan)	2949.73	2972.87	797.44	-0.23

Notes: Treatment includes employees in airfare and hotel department in November 2010 who are both willing and eligible to participate in the WFH program and have an even birth date (2nd, 4th, 6th, etc. day of the month). Treatment includes employees in airfare and hotel department in November 2010 who are both willing and eligible to participate in the WFH program and have an odd birth date (1st, 3rd, 5th, etc. day of the month).

Appendix A3: Explanations on the Work Satisfaction Survey

Work Exhaustion: CTrip's in-house psychology counselors used an adapted excerpt from the Maslach Burnout Inventory Survey to measure the emotional exhaustion of the employees from work. The MBI survey was developed by Berkeley psychologist Christina Maslach and Susan Jackson in the 1970s (see Maslach and Jackson, 1981). Each employee was asked to evaluate his or her "emotional exhaustion" at the end of the work week. The survey contained 6 questions. Each employee was asked to report how often he has felt the way described at work during the week: feel this way every day, almost all the time, most of the time, half of the time, a few times, rarely, never. The survey questions are listed below:

1. I feel emotionally drained from my work.
2. I feel used up at the end of the work day.
3. I dread getting up in the morning and having to face another day on the job.
4. I feel burned out from my work.
5. I feel frustrated by my job.
6. I feel I am working too hard on my job.

Positive and Negative Attitudes: CTrip's in-house psychology counselors used an adapted 16-item Positive and Negative Affect Schedule (PANAS) developed by Clark and Tellegen (1988) to measure the positive and negative attitudes of the employees.

The survey comprised two mood scales, one measuring positive affect and the other measuring negative affect. Each item was rated on a 5-point scale ranging from 1 = *very slightly or not at all* to 5 = *extremely* to indicate the extent to which the employee felt this way the day he took the survey. To evaluate the positive affect, psychologists summed the odd items. In cases with internally missing data (items not answered), the sums were computed after imputation of the missing values: # items on scale / # actually answered, multiplied by the sum obtained from the answered items. A higher score indicates more positive affect, or the extent to which the individual feels enthusiastic, active, and alert. The negative affect is evaluated similarly by summing up the even items. The 16 items were (1) Cheerful, (2) Jittery, (3) Happy, (4) Ashamed, (5) Excited, (6) Nervous, (7) Enthusiastic, (8) Hostile, (9) Content, (10) Guilty, (11) Relaxed, (12) Angry, (13) Proud, (14) Dejected, (15) Active and (16) Sad.

Appendix A4: Quality did not change in the experiment

	(1)	(2)	(3)	(4)
Dependent Variable	recording grade	recording grade	conversion (z score)	conversion (z score)
Mean	3.647	3.647	0.098	0.098
SD	(0.139)	(0.139)	(1.050)	(1.050)
Individual FE	No	Yes	No	Yes
Week fixed-effects	Yes	Yes	Yes	Yes
Experiment*Treatment	0.008	-0.005	-0.022	-0.027
	(0.023)	(0.008)	(0.074)	(0.069)
Treatment	0.000		-0.009	
	(0.005)		(0.095)	
Number of Employees	125	125	134	134
Number of Weeks	85	85	85	85
Observations	6264	6264	9483	9483

Notes: Sample in the first two columns includes 89 order takes in the airfare department (for whom we can obtain recording grade information). The sample in the last two columns includes 134 order takers in airfare and hotels (the group for which conversion rate data exists). Individual clustered standard errors. *** denotes 1% significance, ** 5% significance and * 10% significance.

Appendix A5. Panel A: Treatment Effects Seem Homogeneous across Characteristics

Dep. Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Performance z-score	Child	Female	Female w/ Child	Commute >120min	Renter	Young	Short prior experience	Short tenure	Live w/ parents	Live w/ spouse	Live w/ friends	Pre-exper. performance
Experiment _i x treat _i x "characteristic" _i	0.021 (0.169)	0.033 (0.123)	0.035 (0.195)	0.157 (0.142)	-0.198 (0.140)	-0.151 (0.127)	0.050 (0.127)	-0.085 (0.125)	0.038 (0.134)	-0.020 (0.166)	-0.247 (0.245)	0.024 (0.100)
Experiment _i x "characteristic" _i	0.001 (0.130)	-0.061 (0.087)	-0.055 (0.175)	-0.070 (0.090)	0.117 (0.107)	0.025 (0.091)	0.026 (0.092)	0.118 (0.089)	0.009 (0.101)	-0.021 (0.113)	0.266 (0.207)	-0.208*** (0.077)
Experiment _i x Treatment _i	0.208*** (0.066)	0.193** (0.080)	0.204*** (0.065)	0.158** (0.079)	0.256*** (0.074)	0.296*** (0.100)	0.189** (0.093)	0.251** (0.099)	0.186 (0.114)	0.212*** (0.065)	0.226*** (0.065)	0.205*** (0.060)
Observations	18128	18128	18128	18128	18128	18128	18128	18128	18128	18128	18128	18128
R-squared	0.415	0.415	0.416	0.416	0.416	0.415	0.416	0.415	0.415	0.416	0.419	0.415

Notes: The performance z-scores are constructed by taking the average of normalized performance measures (normalizing each individual measure to a mean of zero and standard deviation of 1 across the sample). The sample includes data from January 1, 2010 to August 15, 2011. "Young" equal 1 if an employee is under 24. "Short prior experience" equals 1 if an employee with less than 6 months of experience before joining CTrip. "Short tenure" equals 1 if an employee has worked in CTrip for less than 24 month by December 2010. "Pre-exper. performance" is the average z-score of performance between Jan 1, 2010 and Oct 1, 2010 for each employee. Individual clustered standard errors. *** denotes 1% significance, ** 5% significance and * 10% significance.

Table 1. Characteristics of employees who volunteer to join WFH

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	Sample mean
Children	0.123** (0.055)		0.075 (0.081)	0.065 (0.080)	0.084 (0.080)	0.090 (0.079)		0.092 (0.080)	0.09
Married		0.095** (0.044)	0.054 (0.063)	-0.002 (0.064)	0.039 (0.064)	0.037 (0.064)		0.040 (0.065)	0.15
Cost of commute (Yuan)				0.005** (0.002)	0.005** (0.002)	0.005** (0.002)		0.005** (0.002)	5.54
Own Bedroom				0.097*** (0.032)	0.089*** (0.032)	0.090*** (0.033)		0.092*** (0.034)	0.60
Tertiary education and above					-0.087*** (0.032)	-0.090*** (0.032)		-0.089*** (0.032)	0.41
Tenure (months)					-0.002*** (0.001)	-0.003*** (0.001)		-0.003*** (0.001)	24.9
Gross wage (Yuan)						-0.003 (0.021)	-0.019 (0.017)	0.032 (0.023)	2872
Age								-0.001 (0.007)	23.2
Male								0.000 (0.035)	0.32
Number of Employees	994	994	994	994	994	994	994	994	994

Notes: The regressions are all probits at the individual level of the decision to work from home. Marginal effects calculated at the mean are reported. The total sample covers all CTrip employees in their Shanghai airfare and hotel departments. Willingness to participate was based on the initial survey in November 2010. Employees were not told the eligibility rules in advance of the survey (own room, 6+ months tenure, broadband internet connect). Gross wage is calculated as a monthly average of salary from Jan 2010 to Sep 2010 (note that 1 Yuan is about 0.15 Dollars). Robust standard-errors.

Table 2: The performance impact of working from home

	(1)	(2)	(3)	(4)	(5)
Dependent Variable	Overall Performance	Phonecalls	Phonecalls	Phonecalls Per Minute	Minutes on the Phone
Dependent Normalization	z-score	z-score	log	log	log
Period: 11 months pre-experiment and 9 months of experiment					
Experiment _t *Treatment _t	0.232*** (0.062)	0.247*** (0.058)	0.120*** (0.025)	0.034** (0.013)	0.086*** (0.028)
Number of Employees	252	135	135	135	135
Number of Weeks	85	85	85	85	85
Observations	17915	9470	9470	9470	9470

Notes: The regressions are run at the individual by week level, with a full set of individual and week fixed effects. Experiment*treatment is the interaction of the period of the experimentation (December 6th 2010 until August 15th 2011) by an individual having an even birthdate (2nd, 4th, 6th, 8th etc day of the month). The pre period refers to January 1st 2010 until December 5th 2010. The z-scores are constructed by taking the average of normalized performance measures (normalizing each individual measure to a mean of zero and standard deviation of 1 across the sample). Since all employees have z-scores but not all employees have phonecall counts (because for example they do order booking) the z-scores covers a wider group of employees. Minutes on the phone is recorded from the call logs. Standard errors are clustered at the individual level. *** denotes 1% significance, ** 5% significance and * 10% significance.

Table 3: Decomposition of the change in labor supply

VARIABLES	(1) Minutes on the Phone	(2) Minutes on the Phone	(3) Minutes on the Phone/ Hours Worked	(4) Hours Worked/ Days Worked	(5) Days Worked
Sample	All	Airfare	Airfare	Airfare	Airfare
Period: 11 months pre-experiment and 9 months of experiment					
Experiment _t *Treatment _t	0.089*** (0.028)	0.090** (0.044)	-0.017 (0.033)	0.068** (0.028)	0.039** (0.015)
Number of Employees	137	89	89	89	89
Number of Weeks	85	85	85	85	85
Observations	9,503	3531	3531	3531	3531

Notes: The regressions are run at the individual by week level, with a full set of individual and week fixed effects. Experiment*treatment is the interaction of the period of the experimentation (December 6th 2010 until August 20th 2011) by an individual having an even birthdate (2nd, 4th, 6th, 8th etc day of the month). The pre period refers to January 1st 2010 until December 5th 2010. Only employees in the Airfare group provides full holiday and leave data so the breakdown by hours and days in the office is only undertaken for this group of order takers. Standard errors are clustered at the individual level. *** denotes 1% significance, ** 5% significance and * 10% significance. Minutes on the phone is recorded from the call logs. Hours worked is measured by the phone system log-in and log-out data.

Table 4: The treatment performance also looked good benchmarked against non-experimental and Nan Tong employees

VARIABLES	(1) Overall Performance	(2) Overall Performance	(3) Phone calls	(4) Phone calls
Comparison to Nan Tong				
	Treatment Vs. Nan Tong	Control Vs. Nan Tong	Treatment Vs. Nan Tong	Control Vs. Nan Tong
Experiment _t *treatment _i	0.204*** (0.047)		0.283*** (0.048)	
Experiment _t *control _i		-0.026 (0.048)		-0.009 (0.043)
Observations	92235	90908	82541	81861
Comparison to Eligible Non-experiment group				
	Treatment Vs. Non-experiment	Control Vs. Non-experiment	Treatment Vs. Non-experiment	Control Vs. Non-experiment
Experiment _t *treatment _i	0.225*** (0.049)		0.231*** (0.051)	
Experiment _t *control _i		-0.016 (0.050)		-0.067 (0.046)
Observations	48596	47269	31049	30369

Notes: Nan-Tong is CTrip's other large call center, located in Nan-Tong, a city about 1 hour drive outside of Shanghai. This call center also had airfare and hotel departments, and calls were allocated across the Shanghai and Nan Tong call centers randomly. The "Eligible non-experimental group" are the individuals that were eligible for the experiment (own room, 6+ months of tenure and broadband) but did not volunteer to participate in the WFH experiment in the two departments in Shanghai. The regressions are run at the individual by week level, with a full set of individual and week fixed effects. Experiment*treatment is the interaction of the period of the experimentation (December 6th 2010 until August 15th 2011) by a Shanghai based individual having an even birthdate (2nd, 4th, 6th, 8th etc day of the month), while Experiment*control is the interaction of the period of the experimentation by a Shanghai based individual having an odd birthdate. All performance measures are z-scores (constructed by taking the average of normalized performance measures, where these are normalizing each individual measure to a mean of zero and standard deviation of 1 across the sample). Standard errors are clustered at the individual level. *** denotes 1% significance, ** 5% significance and * 10% significance.

Table 5: Selection Effects: Employees that continued to WFH after the experiment outperformed

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	Performance	Performance	Performance	Log(Phonecalls)	Log(Phonecalls)	Log(Phonecalls)
Dependent Normalization	z-score	z-score	z-score	log	log	log
Sample	All	All	Balanced	All	All	Balanced
Experiment _t *WFH _{i,t}	0.232*** (0.062)	0.223*** (0.049)	0.181*** (0.057)	0.120*** (0.026)	0.122*** (0.034)	0.101*** (0.039)
Post-Experiment _t *WFH _{i,t}		0.281*** (0.081)	0.242*** (0.089)		0.208*** (0.056)	0.187*** (0.062)
Number of Employees	252	252	152	136	136	78
Number of Weeks	85	143	143	85	143	143
Observations	17915	25655	18304	9470	13339	8854

Notes: WFH_{i,t} here is defined as working-from home at least one day that week. Post-experiment is the period after August 15th 2011 until end of September 2012. Balanced panel drops anybody that quits before the end of May 2012. Individually clustered standard errors *** denotes 1% significance, ** 5% significance and * 10% significance.

Table 6: Employee switches after the end of the experiment

Dep. Variable is to Switch Status	(1)	(2)	(3)	(4)
	Home to Office	Home to Office	Home to Office	Home to Office
Performance during the experiment	-0.074 (0.057)		-0.171** (0.079)	-0.219*** (0.084)
Performance before the experiment		0.008 (0.067)	0.149 (0.097)	0.206** (0.100)
Married				-0.210** (0.072)
Live with parents				-0.182* (0.100)
Cost of commute				-0.010 (0.008)
Observations	110	110	110	110
	Office to Home	Office to Home	Office to Home	Office to Home
Performance during the experiment	0.053 (0.084)		0.052 (0.110)	0.053 (0.112)
Performance before the experiment		0.053 (0.105)	0.003 (0.141)	-0.008 (0.147)
Married				0.100 (0.132)
Live with parents				0.056 (0.138)
Cost of commute				0.011 (0.011)
Observations	74	74	74	74

Notes: The regressions are all probits at the individual level. Marginal effects calculated at the mean are reported. Sample for returning to the office includes the 110 treatment workers still at CTrip at the end of the experiment in September 2011. Pre-experiment performance is the average of individual weekly performance z-score during the pre-experimental period from January 1st 2010 to December 5th 2010. During experiment performance is the average of individual weekly performance z-score during the post-experimental period from December 6th 2010 to August 15th 2011. The sample for moving home includes the 74 control group employees still in the experiment by September 1st, 2011. Out of 74 control workers, 27 petitioned to work at home, yielding a 36.5% join rate. Robust standard errors: *** denotes 1% significance, ** 5% significance and * 10% significance.

Table 7: Working from home is associated with improved employee self-reported work outcomes

Variables:	(1) Satisfaction	(2) General Satisfaction	(3) Life Satisfaction	(4) Exhaustion	(5) Positive Attitude	(6) Negative Attitude
Data source:	Satisfaction survey			Emotion Survey		
Experiment *treatment	0.155*** (0.052)	0.072*** (0.021)	0.168*** (0.047)	-0.564*** (0.168)	0.160*** (0.040)	-0.183*** (0.058)
Announcement*treatment				-0.102 (0.167)	0.080* (0.042)	-0.095 (0.058)
Treatment	-0.015 (0.048)	-0.012 (0.020)	-0.043 (0.066)			
Observations	855	855	855	5109	5109	5109

Notes: The satisfaction survey used in columns (1) to (3) was conducted five times throughout the experimental period: once in early November before the randomization took place and four times after the experiment had started. Treatment defined as having an even birthdate. See details of survey questions and methodology in Appendix A3. The emotion survey used in columns (4) to (6) was conducted every week. The first week was conducted in late November 2010, before the experiment begun but after the randomization so that individuals had been informed of their status in the treatment or control groups. All the dependent variables are logged values. The regressions are run at the individual level with a full set of time-dummies. Experiment*treatment is the interaction of the treatment group with the period of the experimentation. Announcement*treatment is the interaction with the treatment group with the period of post-announcement but pre-experiment. Standard errors are clustered at the individual level. *** denotes 1% significance, ** 5% significance and * 10% significance.

Table 8. Attrition: Poor performers are more likely to quit, especially in the control group

Dependent variable: Quit Sample	(1) Total	(2) Total	(3) Control	(4) Treatment
Treatment _t	-0.176*** (0.055)	-0.020 (0.064)		
Performance _{i,t}		-0.292*** (0.063)	-0.362*** (0.078)	-0.088* (0.048)
Performance _{i,t} *Treatment _t		0.289*** (0.063)		
Age	-0.032*** (0.010)	-0.030*** (0.010)	-0.025 (0.018)	-0.027*** (0.010)
Men	0.070 (0.056)	0.025 (0.057)	-0.046 (0.099)	0.059 (0.065)
Married	-0.042 (0.10)	-0.109 (0.103)	-0.174 (0.218)	-0.043 (0.103)
Cost of Commute	0.008** (0.004)	0.008** (0.004)	0.01 (0.009)	0.00460 (0.003)
Children	0.175 (0.133)	0.345** (0.171)	0.356 (0.248)	0.356 (0.250)
Observations	250	250	120	130

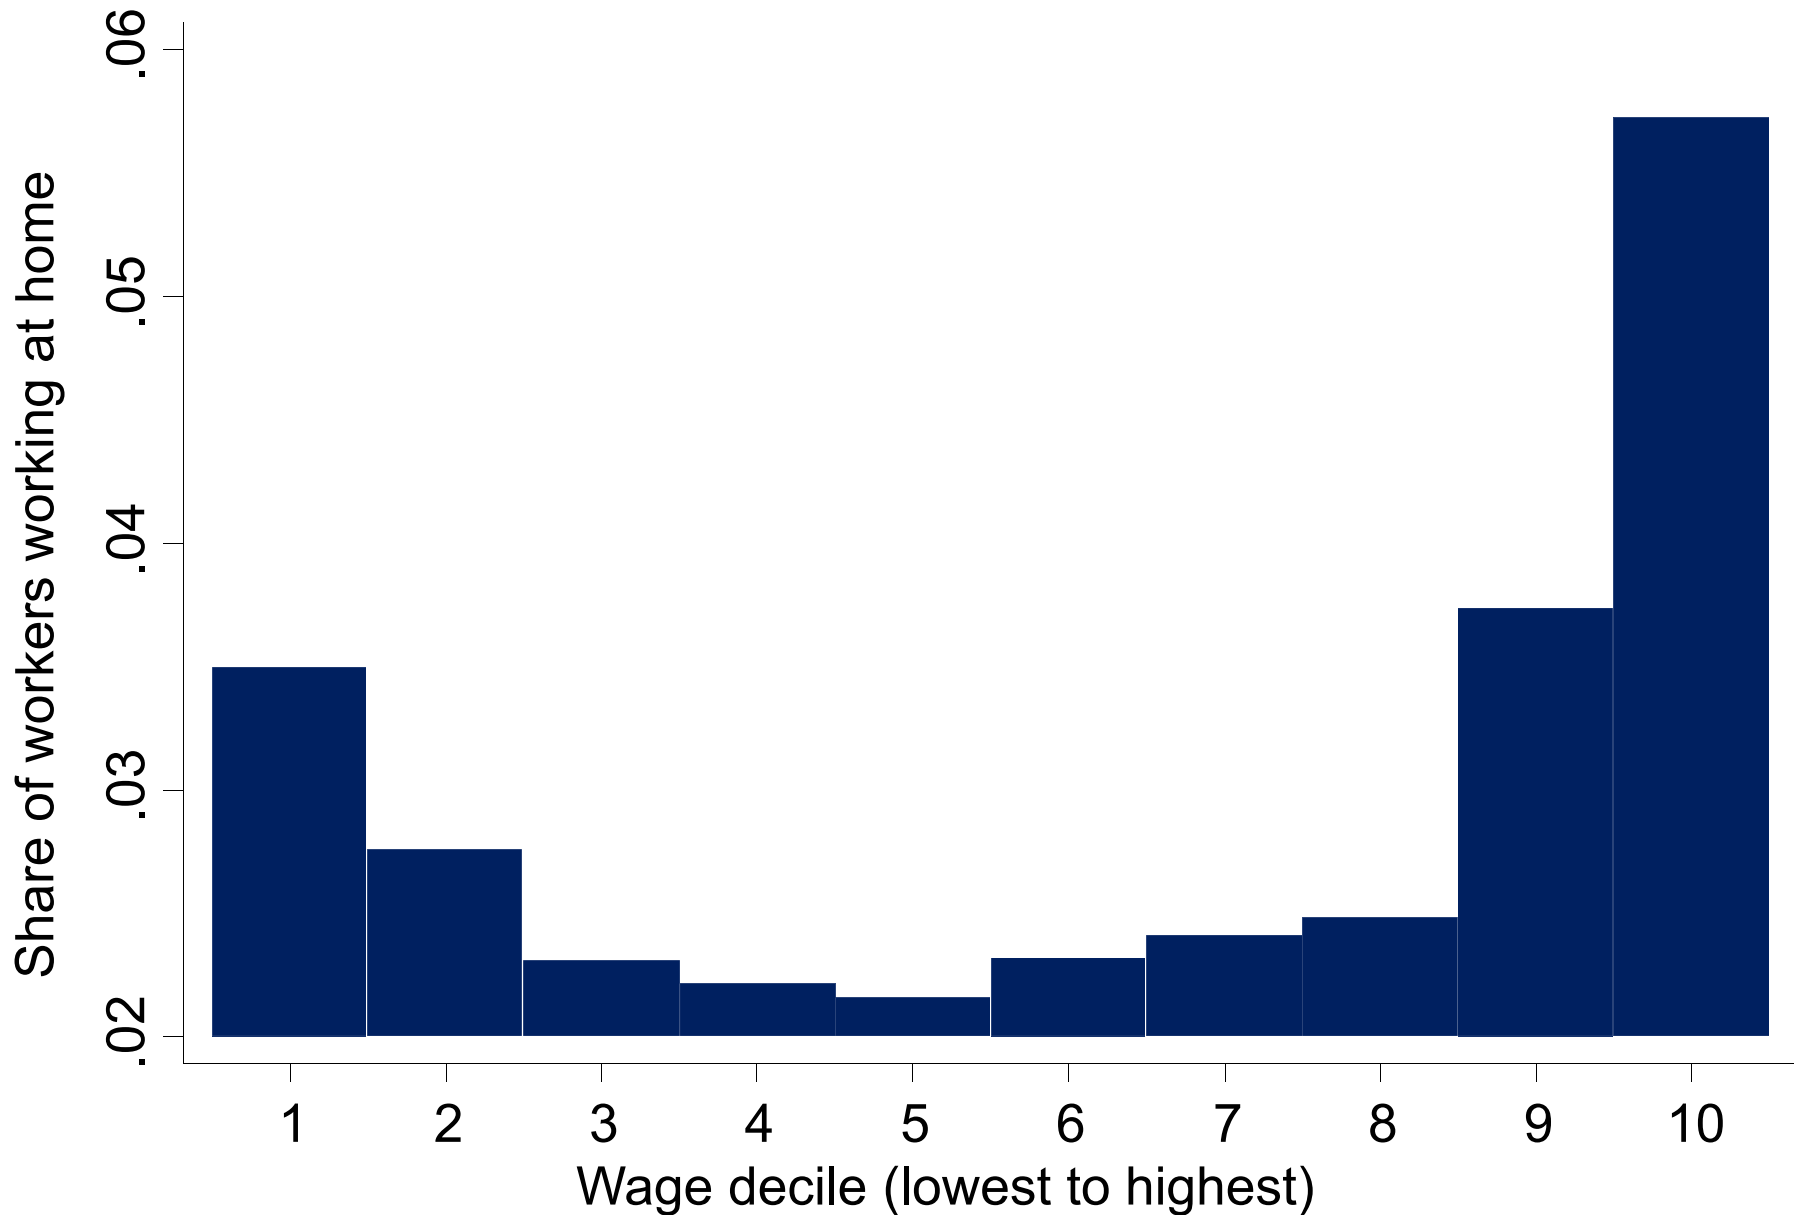
Notes: The regressions are all probits at the individual level. Marginal effects calculated at the mean are reported. The dependent variable is whether the employee quit over the experimental period between December 6th 2010 and August 15th 2011. During this period, control group has an attrition rate of 34.7%, compared to 16.0% in the treatment group. Treatment defined as having an even birthdate. Performance is the average of individual weekly performance z-score during the post-experimental period from December 6th 2010 to August 15th, 2011 (or the week before they quit in the case of quitters). Performance*treatment is the interaction of the performance measure by an individual having an even birthdate. Cost of commute is measured at daily level in Chinese Yuan (note that 1 Yuan is about 0.15 Dollars). Two observations that do not have cost of commute information are excluded. Standard errors are clustered at the individual level. *** denotes 1% significance, ** 5% significance and * 10% significance.

Table 9. Promotion: WFH seems to have a promotion penalty conditional on performance

Dependent Variable: Promoted to team leader or more advanced job function	(1)	(2)	(3)	(4)
Specification	Probit	Probit	Probit	Probit
Treatment	-0.060 (0.046)	-0.087* (0.048)	-0.084* (0.048)	-0.093** (0.047)
Performance during experiment		0.079*** (0.029)	0.083*** (0.031)	0.138*** (0.044)
Performance during experiment*treatment				-0.101* (0.059)
Men			0.079 (0.049)	0.080* (0.048)
Tenure (months)			0.001 (0.001)	0.001 (0.001)
Education: high school and below			0.003 (0.062)	0.009 (0.064)
Observations	252	252	252	252

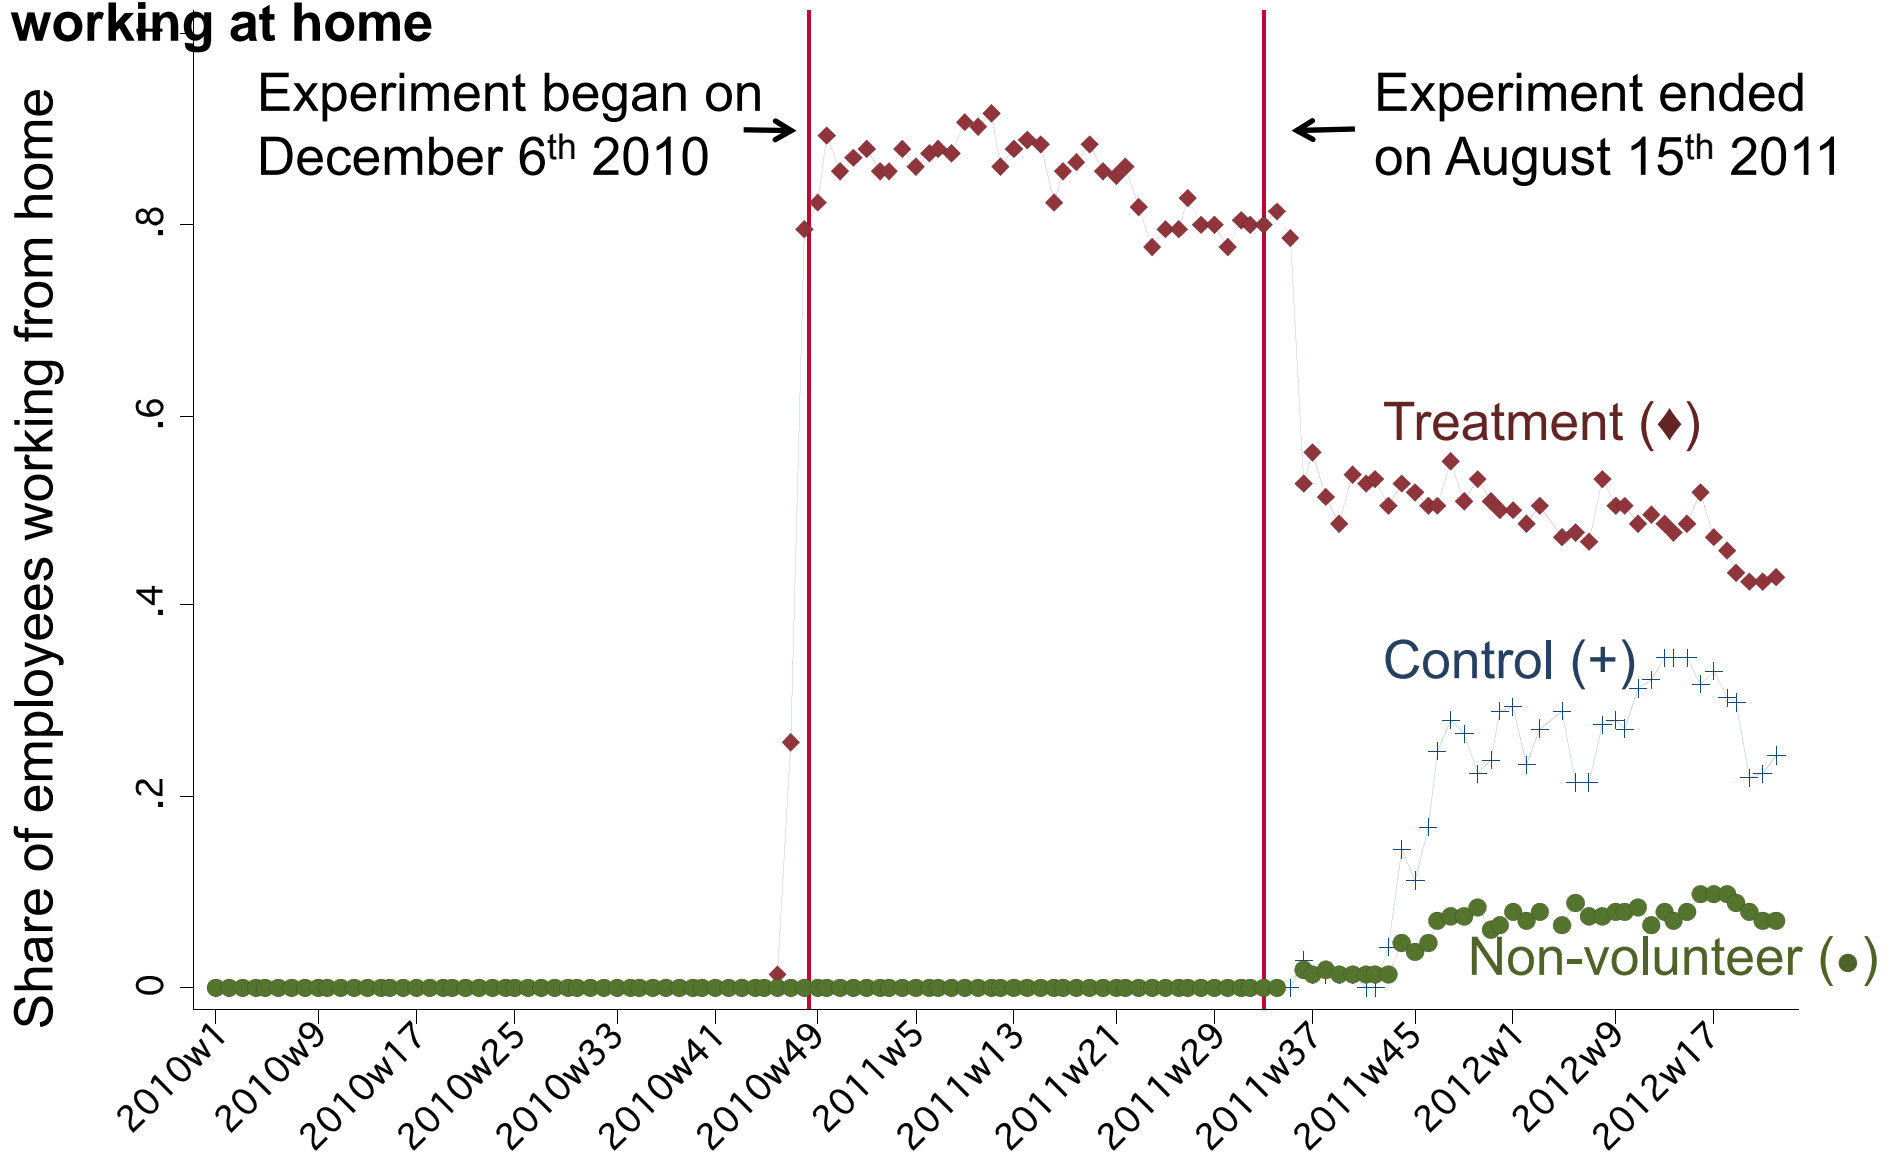
Notes: All regressions are with the dependent variable being whether promoted to team leader or more advanced job function during a 22-month period between Dec 6th, 2010 and Sep 30th, 2012. During this period, a total of 40 employees were promoted, resulting in a 15.9% promotion rate. Performance during the experiment is the average of individual weekly performance z-score during the post-experimental period from December 6th 2010 to August 15th, 2011. For probits, marginal effects evaluated at the mean are reported. Performance*treatment is the interaction of the performance measure by an individual having an even birthdate (2nd, 4th, 6th, 8th etc day of the month). Robust standard errors. *** denotes 1% significance, ** 5% significance and * 10% significance.

Figure 1: Home working is more common for low and high wage groups



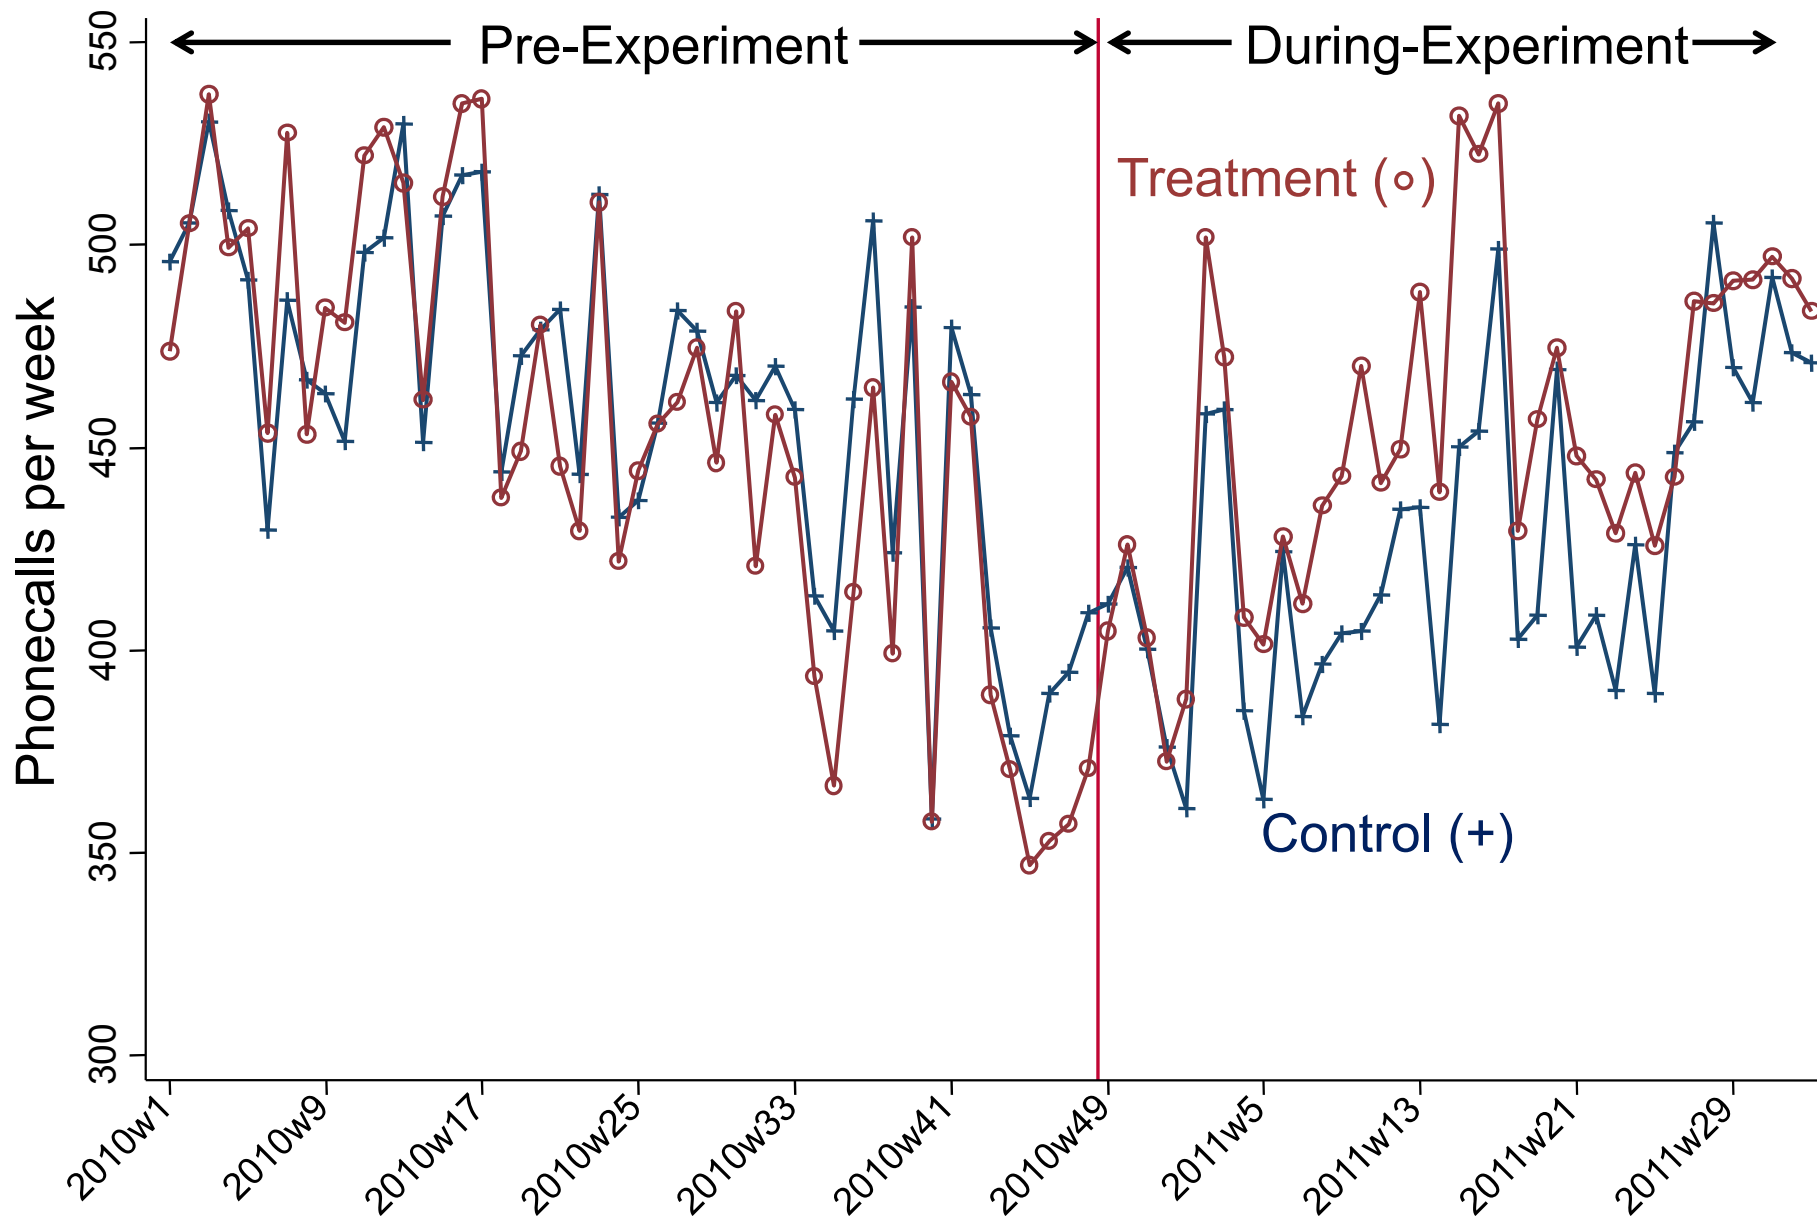
Note: PUMS census data for 2010. We classify workers as mainly working from home if they answer “work from home” to the census question “How did you get to work last week?” All employees are divided into 10 deciles by wage. Share of workers working at home is then calculated within each wage decile.

Figure 2: Share of treatment, control and non-volunteer employees working at home



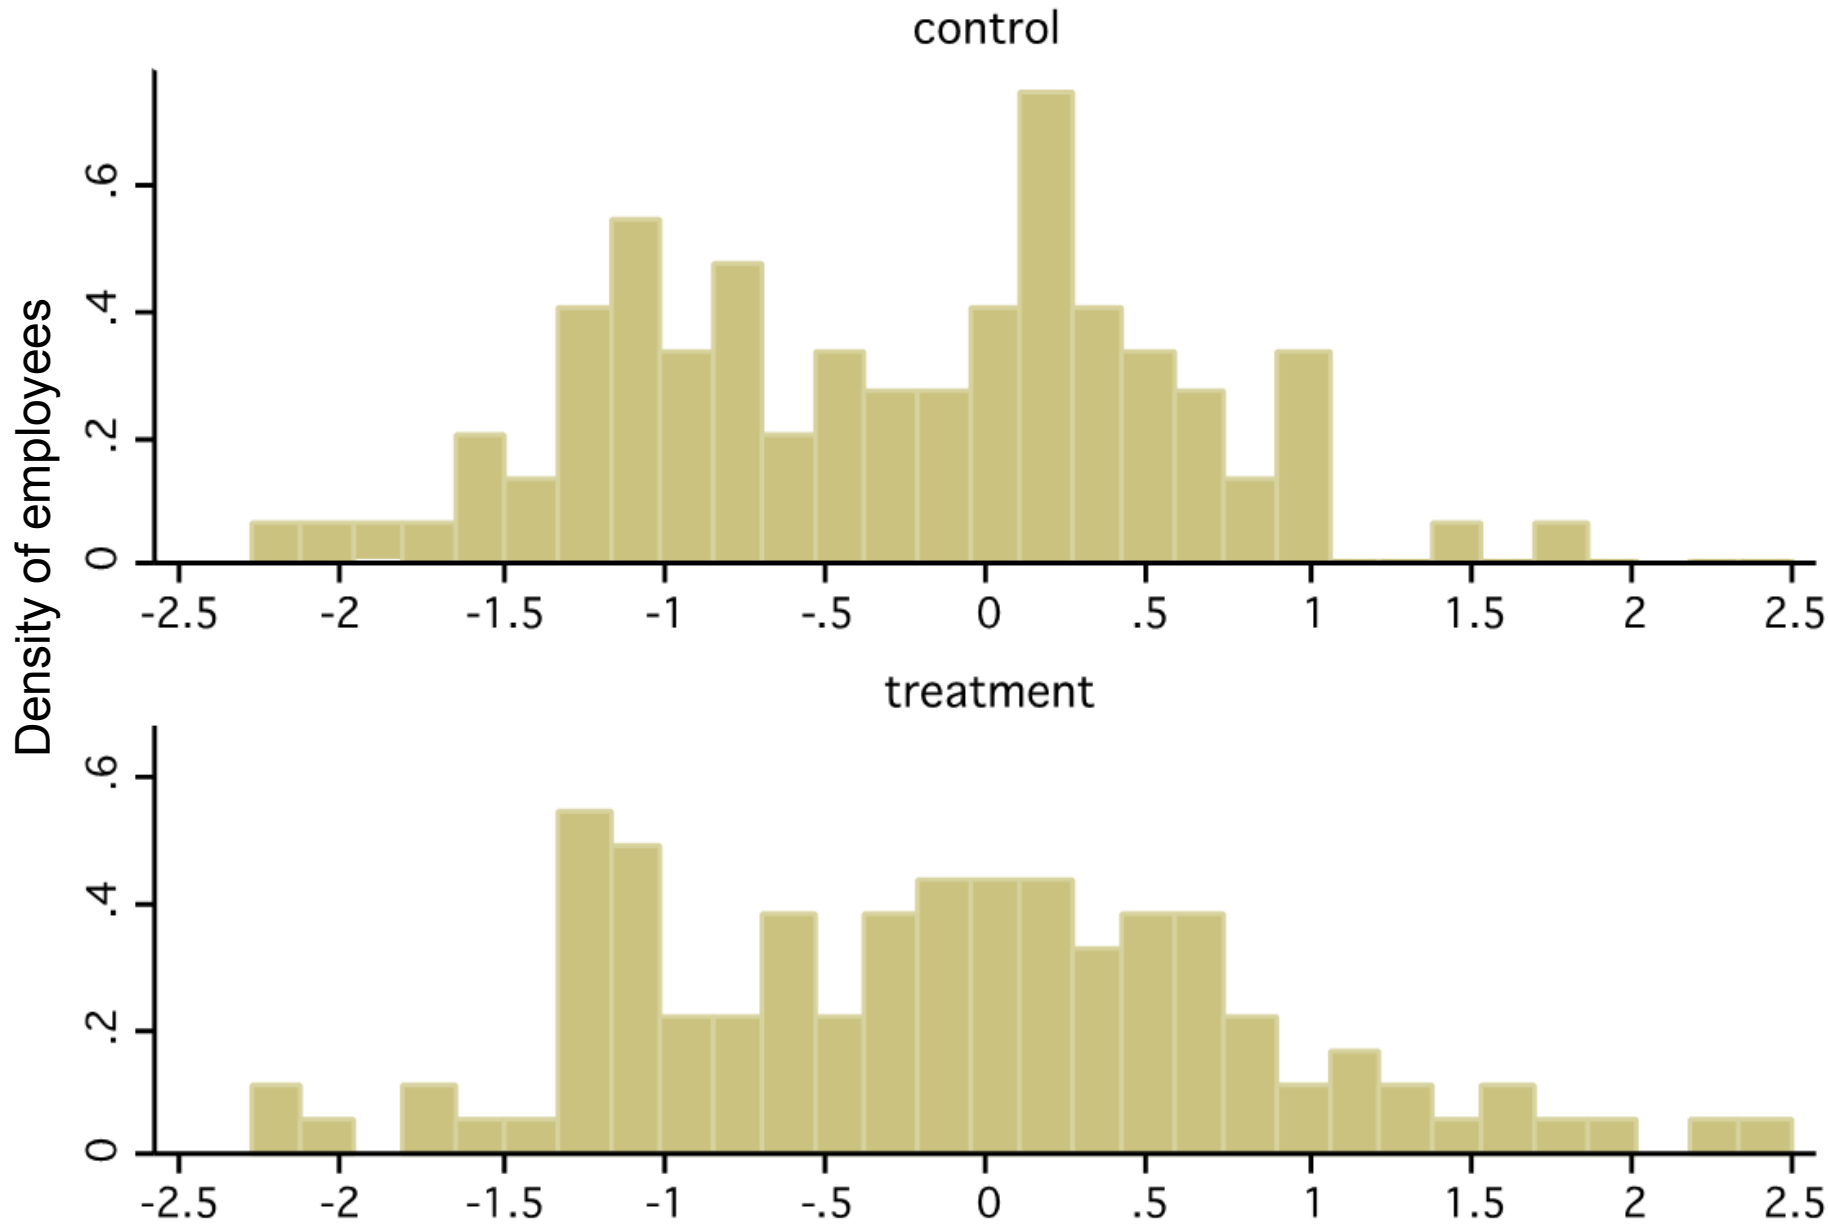
Note: Data from January 4th 2010 until June 1st 2012. Percentage of workers working at home = (number of workers working at home / number of workers still employed) calculated separately for treatment (even-numbered birthdays), control (odd-numbered birthdays) and non-volunteer workers (those did not volunteer to work from home). First red line indicates the beginning of the experiment on December 6, 2010 and the second red line indicates the end of the experiment on August 15, 2011, after which the option to work from home was available to all employees. There is about a two week transition lag in returning to the office. There is also some WFH before the start of the experiment due to the installation period (the IT department had to set up computer terminals in every employee's home). The sample is all employees in airfare and hotel departments in Ctrip's Shanghai headquarters who were employed on December 6, 2010.

Figure 3. Performance of treatment and control employees: phone calls



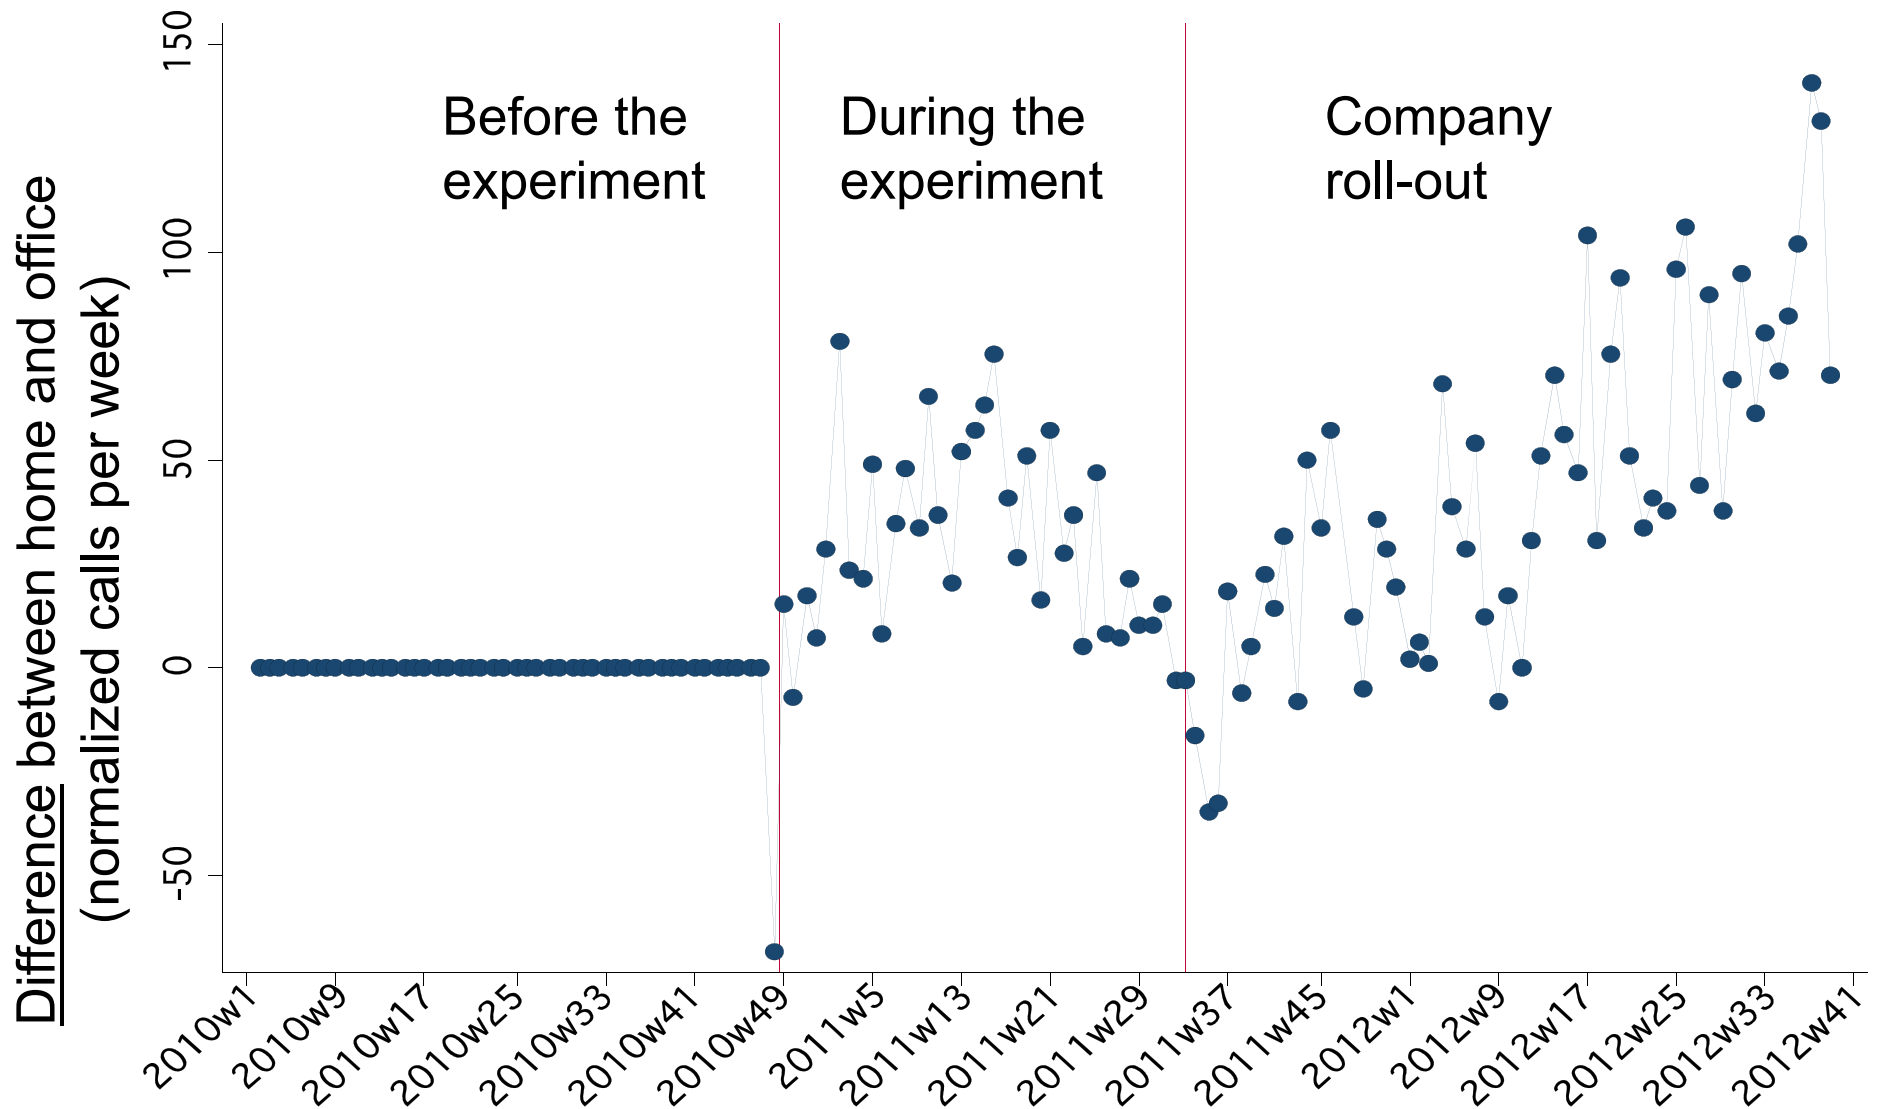
Note: Data from January 4th 2010 until August 15th 2011. Number of phone calls made for order-takers (the group for whom number of phone calls taken is a performance metric) calculated separately for treatment (even-numbered birthdays) and control (odd-numbered birthdays).

Figure 4. Cross-sectional performance spread during the experiment



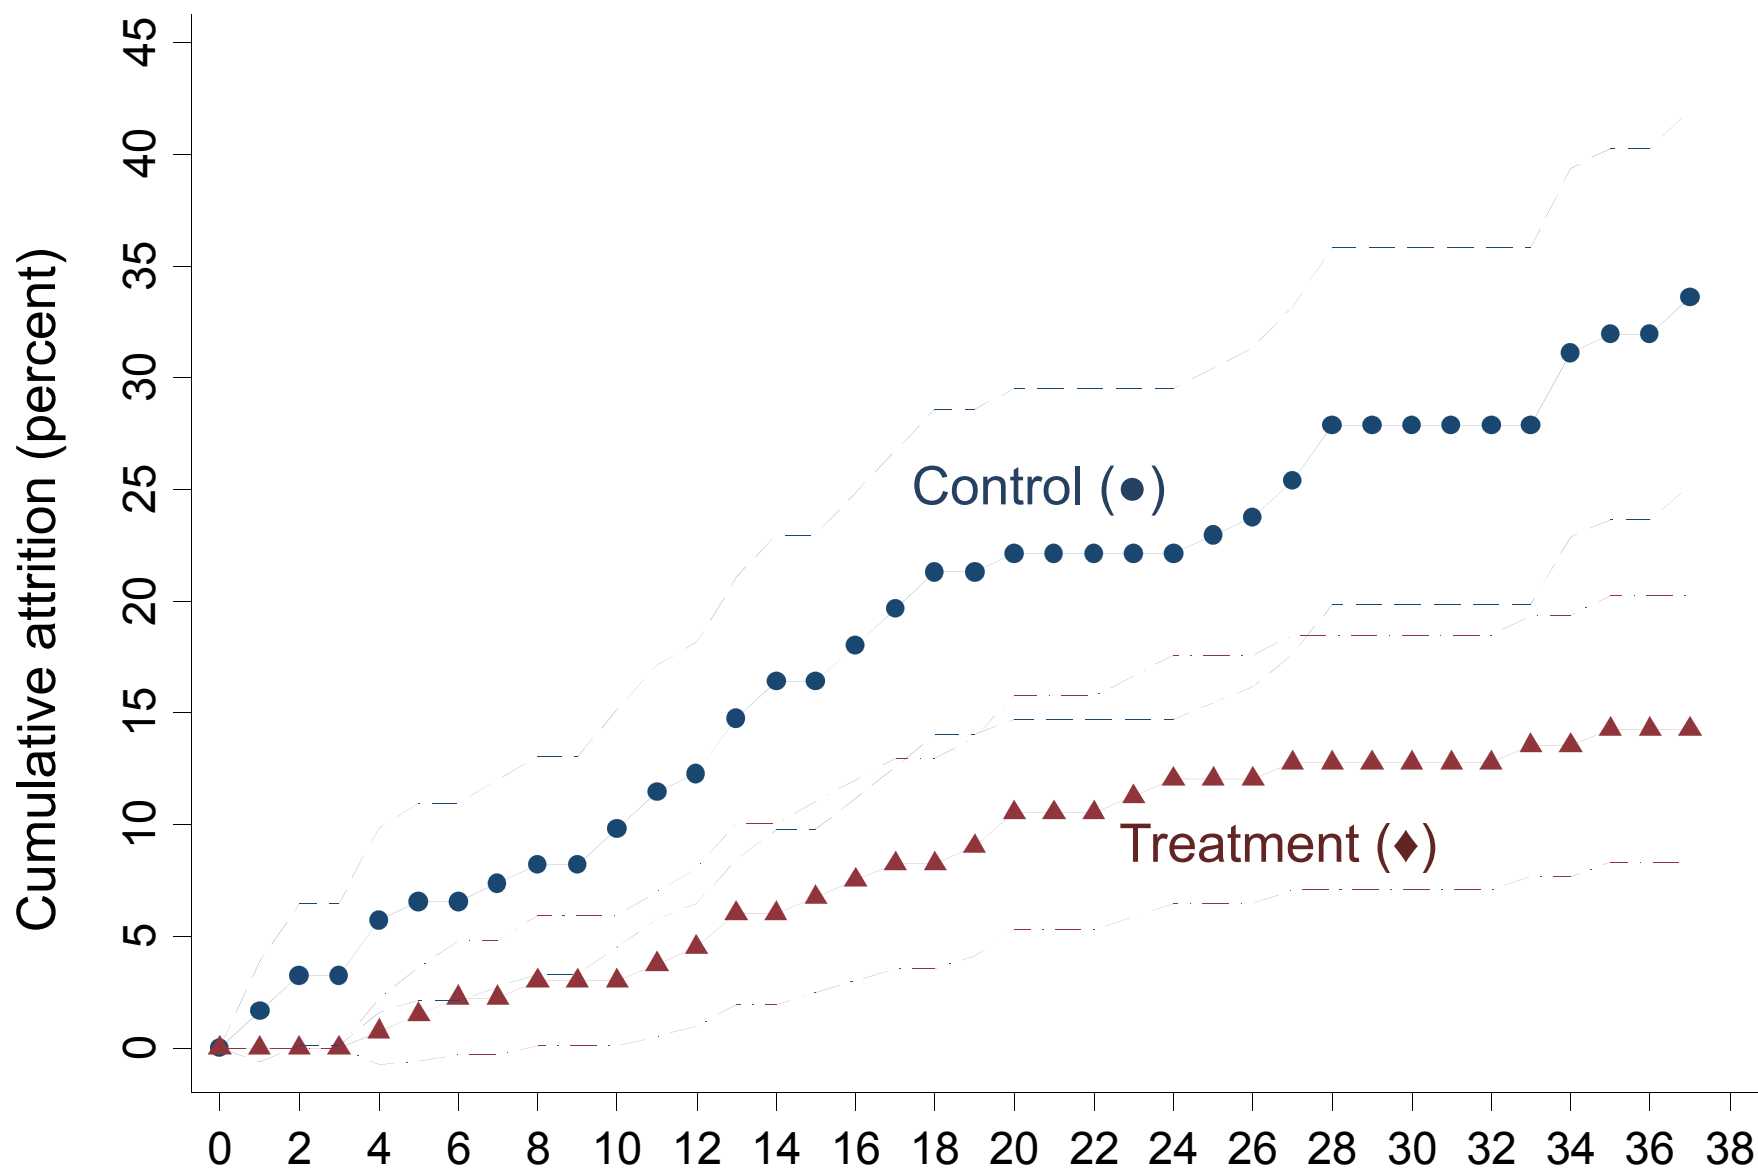
Note: Histograms of the performance z-score for the treatment and control groups after 3 months into experiment (SD=1 across individuals in pre-experimental data)

Figure 5. Selection further increased the performance impact of home working during the company roll-out



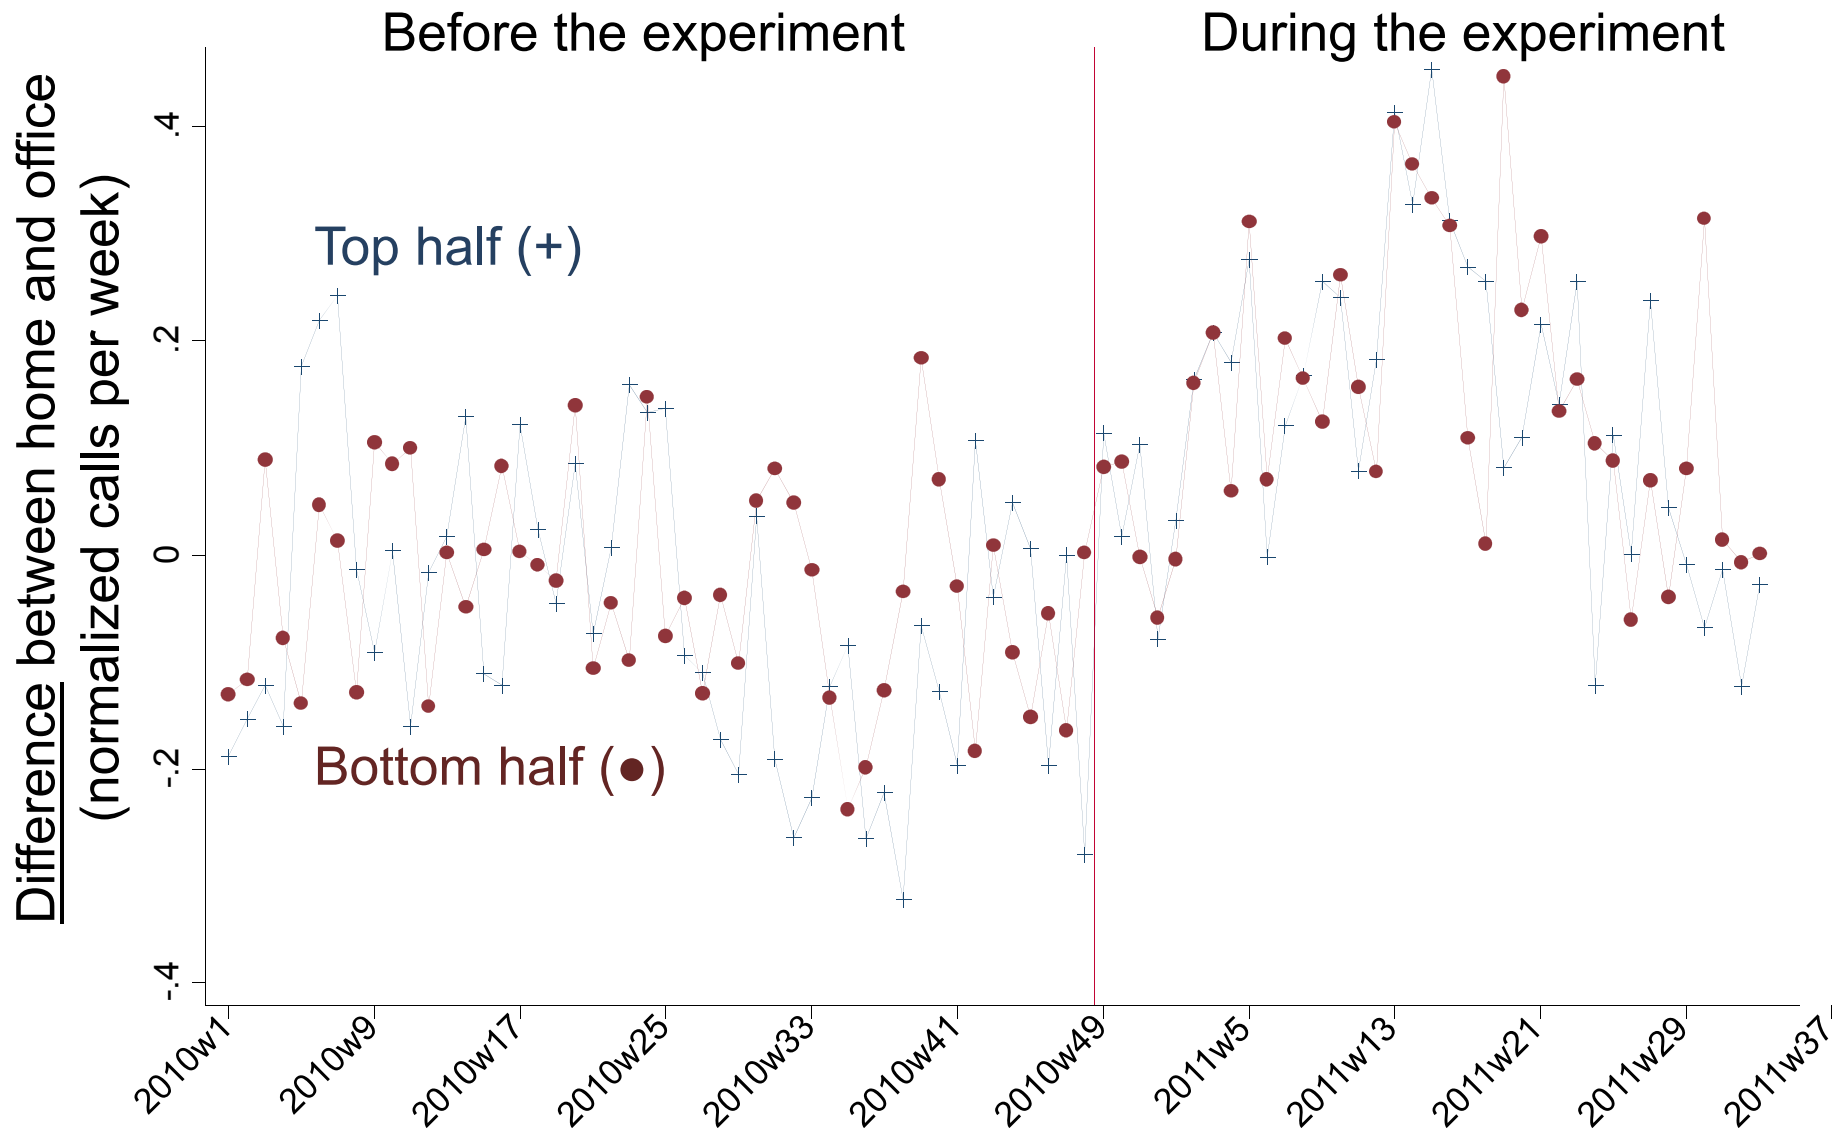
Note: Data from January 4th 2010 until October 1st 2012. Phone calls in z-scores (normalized so the pre-experiment values are mean zero and standard deviation 1) shown as the difference between home and office workers. The drop in performance before the start of the experiment is due to the disruption from the transition from office based to home based working for the treatment group. The dip at the end is similarly the disruption for home based employees that are moving back to the office (who until they are fully office based are coded as home workers).

Figure 6. Attrition is halved by working from home



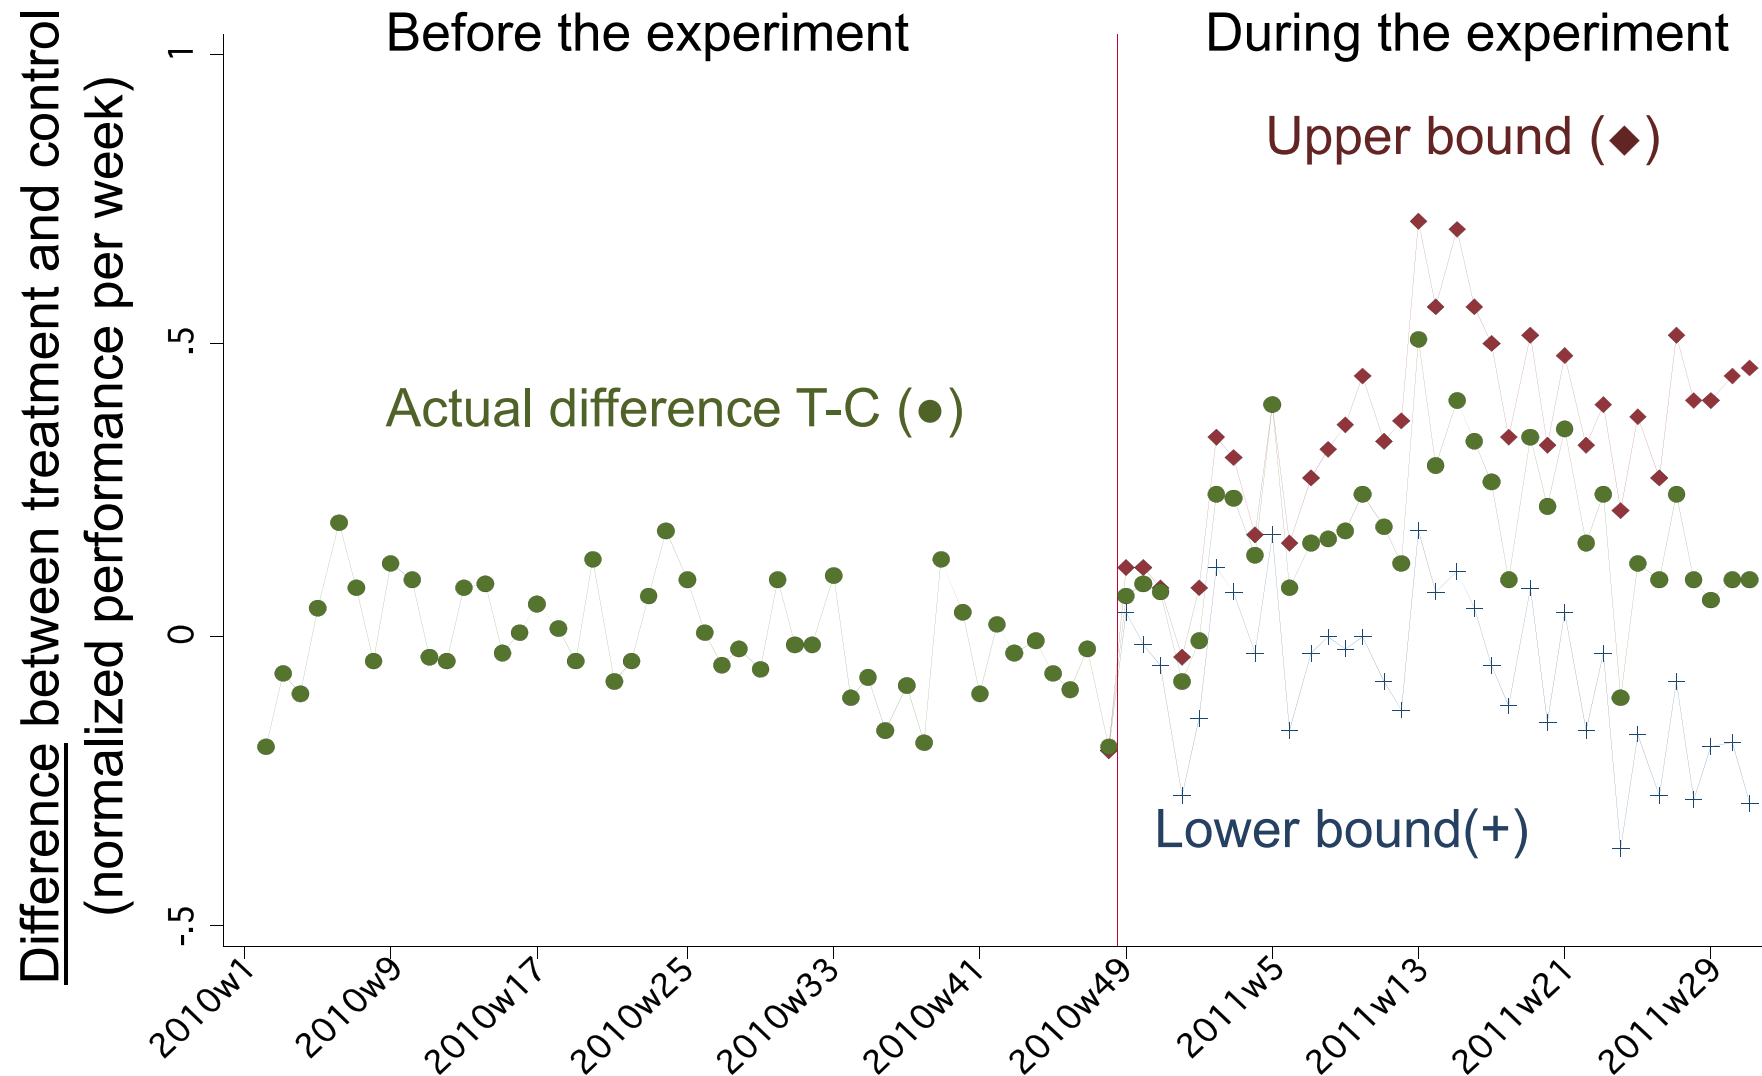
Note: Cumulative attrition rate equals number of employees attrited by week x of the experiment divided by total number of employees at the beginning of the experiment, calculated separately by treatment and control group. Dashed lines represent 95% point-wise confidence intervals calculated by bootstrap. Experiment started on week1 and ended on week 39.

Figure 7. The top and bottom half of employees by pre-experiment performance both improved from working at home



Note: Data from January 4th 2010 until August 15th 2011. Phone calls in z-scores (normalized so the pre-experiment values are mean zero and standard deviation 1). Calculated separately for the difference between the top half of the treatment and control groups and the bottom-half of the treatment and control groups, where performance halves are based on pre-experiment performance.

Appendix Figure A1. The Lee (2008) bounds for estimating the impact of non-random attrition



Note: Data from January 4th 2010 until August 15th 2011. Performance in z-scores (normalized so the pre-experiment values are mean zero and standard deviation 1; performance measures used depend on the type of employees). Upper bound is calculated assuming the worst performers of the treatment group would quit to create a balanced panel (between treatment and control) each week. Lower bound is calculated assuming the best performers for the treatment group would quit.

Exhibit A: Ctrip is a large and modern Chinese firm



Headquarters in Shanghai



Main Lobby



Call Center Floor



Team Leader Monitoring Performance

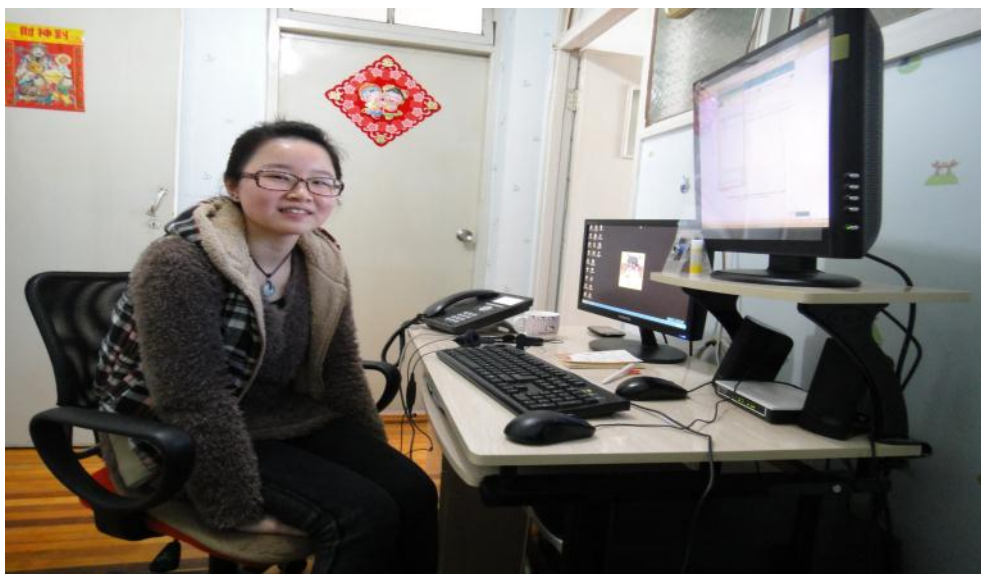
Exhibit B: The randomization, and examples of home-workers



Treatment groups were determined by a lottery



Working at home



Working at home



Working at home