The Value of Office Time*

Cevat Giray Aksoy† Nicholas Bloom‡ Steven Davis§ Victoria Marino¶
Cem Özgüzel‖

February 1, 2026

Abstract

This paper evaluates whether structured in-person contact can improve performance and retention in a fully remote workplace. We study a nine-month randomised controlled trial in the inbound call-centre operations of a large business-process-outsourcing firm in Turkey, conducted in partnership with a multinational telecommunications client. A sample of 248 fully remote customer-service agents was randomly assigned either to remain fully remote (control) or to attend the office one day per month (treatment), holding pay, schedules, tasks, and performance evaluation constant. The monthly office-day policy increases productivity. Calls per hour rise by about 4.4% on average, driven by roughly 5% shorter call durations, with no evidence of declines in service quality (customer ratings and audit scores) or changes in break time. The intervention also reduces attrition over the experimental horizon, with particularly strong retention effects among higher-performing agents. Evidence on mechanisms points to interaction and learning: treated agents report more frequent managerial feedback and stronger team communication, and randomised seating assignments show that brief co-location generates new peer links and productivity spillovers. Overall, predictable in-person contact can improve productivity and retention in fully remote service work.

JEL Codes: M12, M54, J24, L84
Keywords: remote work, productivity, retention, job satisfaction

*We are grateful to the employees of Tempo BPO, especially Cemal Akar, Tugce Hunkar, Nail Karasekreter, and Sedat Pakir, for their significant time and effort in preparing and implementing the experiment. We also thank seminar participants at Cardiff, HEC Liege, London Business School, Paris School of Economics, Sorbonne University, the Tinbergen Institute, Ozyegin University, and the 2025 Remote Work Conference at Stanford for helpful feedback. The project received funding from the European Bank for Reconstruction and Development, King's College London, and Stanford University. It was approved by the King's College London Institutional Review Board (Protocol ID: MRA-23/24-43056) and registered with the AEA RCT Registry (AEARCTR-0013714). The views expressed here are those of the authors and do not necessarily reflect the views of their institutions or the EBRD. Any errors are our own.

†European Bank for Reconstruction and Development (EBRD) and King's College London, aksoyc@ebrd.com
‡Stanford University, nbloom@stanford.edu
§Stanford University, Hoover Institution and SIEPR, stevend5@stanford.edu
¶European Bank for Reconstruction and Development, robinsov@ebrd.com
‖Paris School of Economics and Sorbonne Economics Center, cem.ozguzel@univ-paris1.fr

1 Introduction

Remote and hybrid work expanded sharply during the COVID-19 pandemic and have remained common in many labour markets. By 2025, roughly 100 million employees in Europe and North America worked in schedules that combine home and office days (Barrero et al., 2021; Buckman et al., 2025; Aksoy et al., 2025a; Luca et al., 2025). Remote work can benefit workers through flexibility and firms through access to broader labour pools, with potential inclusion gains for groups facing commuting or scheduling constraints (Aksoy et al., 2025b; Bloom et al., 2026). At the same time, firms frequently express concerns that sustained physical separation weakens mentoring, slows learning, and erodes cohesion.

Evidence on the performance and retention effects of remote and hybrid work remains mixed and context dependent. Hybrid arrangements can reduce attrition without harming performance in some settings (Bloom et al., 2024). In fully remote operations, productivity effects vary with task structure, monitoring technologies, and managerial practices (Emanuel and Harrington, 2024; Gibbs et al., 2023; Battiston et al., 2021; Choudhury et al., 2021). For many organisations, the practical question is one of design. When day-to-day work is remote, what in-person contact, if any, is needed to sustain coordination, learning, and attachment to the firm?

This paper studies a simple policy that many remote-first operations can implement. Agents attend the office one fixed day per month and otherwise remain fully remote. We evaluate this policy in a nine-month randomised controlled trial in the inbound call-centre division of TEMPO BPO, one of Turkey's largest business-process-outsourcing firms, in partnership with a multinational telecommunications client. The study sample comprises 248 fully remote customer-service agents working on the same client account. Agents were randomly assigned either to remain fully remote (control) or to attend the office one day per month (treatment). The intervention was intentionally light-touch. It relied on existing infrastructure and held pay, shift schedules, tasks, call routing, and performance evaluation constant. Office days provided co-location with peers and in-person access to supervisors but introduced no formal training curriculum.

Three findings summarise the results. First, monthly office days raise productivity, but only gradually. We detect little change in the first five months after the policy begins. Effects emerge in later months, when treated agents handle more calls per hour than their fully remote peers. By the end of the intervention, productivity is higher by about 4.4% on average. The increase is driven by shorter calls, with average call duration falling by roughly 5%, and the largest change in talk time rather than administrative or hold components. Second, there is no evidence of a speed–quality trade-off. Customer satisfaction ratings and internal audit scores remain stable, and break time does not decline. Third, retention improves. Attrition is lower in the treatment group over the experimental horizon, the gap persists beyond the final office visit, and the policy disproportionately retains higher-performing agents.

We provide evidence on mechanisms consistent with interaction and learning. Randomised seating assignments during the first office visits generate exogenous variation in peer proximity. Assigned desk neighbours are more likely to communicate afterward, and exposure to higher-performing neighbours predicts larger subsequent productivity gains. Endline surveys point in the same direction. Treated agents report more frequent managerial feedback, stronger team communication, and higher perceived cultural fit, while broader measures of general well-being move little. Taken together, the

evidence suggests that periodic co-location strengthens professional interactions and facilitates the transfer of tacit know-how. From the firm's perspective, the policy is inexpensive and appears cost-effective. The per-visit cost of transport and meals is modest, and a conservative cost–benefit calculation yields a benefit–cost ratio of 1.77, with gains arising from both higher productivity and reduced turnover.

Related literature and our contributions. Our paper speaks to two strands of research. First, we relate to work on structured managerial practices (in particular feedback, coaching, and performance conversations) as determinants of engagement and performance. Bloom et al. (2013) and Lazear et al. (2015) emphasise the productivity returns to better management, while Manthei et al. (2023) show that regular performance conversations can improve both profits and retention in a retail setting. Closest to our mechanisms, Choudhury et al. (2025) documents that temporary co-location shapes communication patterns in an all-remote organisation. We contribute by providing experimental evidence that even infrequent, coordinated co-location can strengthen feedback and peer learning in a fully remote service operation, and we trace these channels using randomised seating assignments and interaction data.

Second, we contribute to the growing evidence on the effects of hybrid and fully remote work on worker performance. Experimental and quasi-experimental studies of hybrid arrangements typically find positive or null impacts on performance (Bloom et al., 2015; Battiston et al., 2021; Choudhury et al., 2021; Bloom et al., 2024). By contrast, evidence on fully remote work spans a wider range of estimates, from negative to positive, reflecting heterogeneity in tasks, monitoring, and organisational practices (Emanuel and Harrington, 2024; Gibbs et al., 2023; Atkin et al., 2023; Aksoy et al., 2025b). Our contribution is to study a remote-first setting, where remote work is already the default, and evaluate a minimal hybrid policy that many remote organisations can implement: one coordinated office day per month. In a semi-routinised service context with tight performance measurement, we show that this policy raises productivity without reducing service quality and improves retention, including by mitigating the selective loss of high performers. More broadly, our results highlight that managerial and peer-interaction structures are central to sustaining performance in fully remote operations.

The remainder of the paper is organised as follows. Section 2 describes the setting and experimental design. Section 3 presents the data and sample characteristics. Section 4 outlines the empirical strategy. Section 5 reports the main results on productivity, service quality, retention, and mechanisms. Section 6 presents the cost–benefit analysis. Section 7 concludes.

2 Setting and experimental design

The company and setting. Our intervention was implemented in the call-centre division of TEMPO BPO (hereafter, TEMPO), in collaboration with one of its clients (a major multinational telecommunications company).[1] TEMPO is one of Turkey's largest business-process-outsourcing (BPO) providers. Founded in 2002, it delivers call-centre and back-office services to a mix of domestic and international clients in sectors such as telecommunications, banking, food delivery, and government services. As of 2026, TEMPO employed more than 3,500 call-centre agents, operated

[1] For confidentiality reasons, we do not disclose the client's name.

in seven provinces, and was headquartered in Istanbul.

The experiment takes place in TEMPO's inbound call-centre operations. This environment is well suited for studying workplace practices because production is highly standardised and routinely monitored. Agents perform a clearly defined task (handling incoming customer calls) under fixed schedules and stable team structures.

Agents are paid the national minimum wage and do not receive performance bonuses or other variable pay. Career progression instead operates through promotions: strong performers can be promoted to team-leader positions, which involve managerial responsibilities and higher pay. Organisationally, each agent is assigned to a team of roughly 20 agents overseen by a team leader. Team leaders manage day-to-day operations and monitor individual performance. TEMPO tracks outcomes continuously through an in-house dashboard that records, among other metrics, call volume, talk time, customer-satisfaction measures, and call-audit scores.

Work is organised in standard five-day, eight-hour shifts, with two 15-minute breaks and a 30-minute lunch. Calls are routed centrally to the first available agent. Agents handle consumer queries using a headset and a firm-provided computer, following standard protocols set by the client account.

Remote-work transition and motivation for the intervention. In March 2020, following Turkey's first national lockdown, TEMPO moved its entire agent workforce to remote work within two weeks. The firm provided laptops and internet subsidies, while keeping team structures, schedules, monitoring, and evaluation metrics unchanged. Team leaders continued check-ins and performance reviews virtually, and the compensation structure remained the same: agents stayed on the national minimum wage, with unchanged shift lengths and task expectations.

After restrictions eased in late 2021, TEMPO chose to maintain remote work as its default operating model. By mid-2024, more than 90% of agents remained fully remote. While the shift delivered operational benefits, internal reviews increasingly highlighted organisational costs, including weaker cultural cohesion, fewer opportunities for real-time coaching, and persistent retention problems. In practice, day-to-day interactions between agents and team leaders often occurred via messaging applications, with limited informal contact and no in-person interaction. These concerns motivated TEMPO's collaboration on the randomised controlled trial (RCT) evaluated in this paper, which tests whether structured monthly office visits can strengthen managerial support and improve retention and performance.

Client account and implementation site. The study population consists of agents assigned to the collaborating telecommunications client account, working in inbound customer-service roles (handling calls related to billing, packages, and service issues). To implement the intervention, TEMPO relied on its Şanlıurfa office, one of several regional hubs the firm operated before the pandemic. Located in southeastern Turkey, the office is accessible by shuttle and has sufficient capacity to host treatment-group agents once per month. This allowed TEMPO to run the intervention using existing infrastructure and without disrupting other ongoing operations.

Study population and recruitment. As mentioned above, the study focuses on fully remote inbound agents serving the telecommunications client account. In March 2024, TEMPO informed 661 agents in this operation about a pilot program that would require working from the Şanlıurfa office one day per month. Participation in the study was voluntary. Over March–April, 476 agents

(about 72%) expressed interest, while 185 did not volunteer, most commonly citing family or care responsibilities and the burden of commuting.

Eligibility and analysis sample. From the volunteer pool, TEMPO applied pre-specified operational criteria to ensure that monthly office days were feasible and that participants had sufficient experience in the role. Agents were eligible if they (i) had at least three months of tenure at the time of recruitment and (ii) lived within roughly 45 minutes of the Şanlıurfa office (approximately a 16-kilometre / 10-mile radius). These criteria yielded 248 eligible volunteers, who constitute the experimental sample. Relative to non-volunteers, eligible agents tended to have shorter commutes and were more likely to be single, but they did not differ systematically in baseline performance measures.

Random assignment and baseline period. Randomisation was implemented in late April 2024 by TEMPO's HR department using a transparent lottery based on company ID numbers: agents with odd-numbered IDs were assigned to treatment and those with even-numbered IDs to control. This procedure generated two groups of 124 agents each. Team assignments and team leaders were then fixed in June 2024 (approximately 20 agents per team), and agents remained in the same teams throughout the study.

Both groups continued working fully remotely during June and July 2024, providing a clean pre-intervention baseline for measuring outcomes and assessing balance. We discuss the sample characteristics and balance in the next section

Intervention: monthly office day. The intervention ran from August 2024 to April 2025. Agents assigned to treatment were scheduled to work from the Şanlıurfa office on one fixed day per month; control agents remained fully remote throughout (Figure A.1a). On office days, treatment agents travelled together on company shuttles and arrived before their regular shift (Figure A.1b). They performed their usual tasks using TEMPO-provided laptops and the same software systems and protocols as on remote days. Importantly, the production technology and incentives were held constant: shift length, pay, call routing, and performance evaluation criteria were unchanged, and no overtime was permitted.

What changes on office days. Relative to a standard remote day, the monthly office visit shifted three margins that are central to the intervention's hypothesis. First, *managerial input*: office days created opportunities for real-time feedback and mentoring by team leaders that are difficult to replicate through asynchronous messaging. Second, *peer interaction*: the office facilitated informal social contact and on-the-job learning through conversations and peer support. Third, *commuting and the work environment*: treated agents incurred travel time (up to roughly 45 minutes each way), while control agents continued working from home in their usual setting.

To measure interaction patterns and facilitate the study of potential peer spillovers, we also implemented randomised seating assignments during office visits. The office-day schedule (Figure A.1c) also created shared time for informal contact (e.g., breakfast, coffee breaks, and lunch). Importantly, we did not introduce any formal training or curriculum on office days, so the treatment captures the effect of in-person co-location and interaction rather than structured instruction. All direct costs of participation, transport, meals, and refreshments, were borne by the firm.

After each office day, treated agents returned to fully remote work until the next scheduled visit. Compliance was high: TEMPO sent multiple reminders to treated agents ahead of each scheduled

office day and approximately 95% of assigned office days were attended as planned.

Hypotheses. Our pre-analysis plan specifies three primary hypotheses about the effects of monthly in-office visits on workers and the organisation:

H1. **Productivity and service quality.** Monthly office visits improve performance. Through more effective real-time feedback and monitoring on office days, treated agents are expected to handle calls more efficiently (e.g., more calls per hour and/or shorter call duration) and improve service-quality metrics (e.g., customer satisfaction and call-audit scores).

H2. **Retention.** Monthly office visits increase retention (lower quits/turnover) among treated agents.

H3. **Communication and team cohesion.** Monthly office visits strengthen communication between agents and team leaders. Face-to-face interaction is expected to facilitate knowledge sharing, and team cohesion, which may improve both performance and the broader work climate.

These hypotheses reflect the idea that even infrequent in-person contact can "refresh" social ties and support the exchange of tacit knowledge in an otherwise remote work arrangement. At the same time, the magnitude of any effects is ex ante ambiguous: a single day per month may be too limited to generate large changes in behaviour or outcomes. The experiment is designed to quantify these effects.

3 Data, sample characteristics and balance

Data sources and linkage. We combine administrative and survey data to evaluate the effects of the monthly office-day intervention. The analysis draws on three sources: (i) high-frequency operational data from TEMPO's internal performance dashboard, (ii) HR records on employment status and office-day attendance, and (iii) surveys fielded by us to measure worker experiences and interaction patterns. All datasets are linked using anonymised agent identifiers.

Main outcomes and retention data. The core dataset is an agent-level extract from TEMPO's administrative performance management system. It covers the full experimental sample of 248 agents over a 12-month pre-intervention window (August 2023–July 2024) and the nine-month intervention window (August 2024–April 2025). The raw data are recorded at high frequency and can be aggregated to daily, weekly, and monthly panels.

The performance dashboard provides high-frequency measures of agent activity and time use, recorded at the agent-by-hour-by-day level. We observe the number of calls handled and our main productivity outcome, *calls per hour*. To understand the mechanisms behind changes in call volume, we use *average call duration* (seconds per call) and its components: *talk time* (time speaking with the customer), *hold time* (time the customer is placed on hold), and *administrative time* (after-call/on-screen processing), all measured in seconds per call. We also use *break time* (seconds per hour) to capture time away from active call handling within the shift.

Service-quality outcomes are observed at the monthly level and include customer satisfaction ratings, which customers provide voluntarily at the end of calls, as well as internal call-audit scores.

Audit scores are based on 10 randomly selected calls per agent-month and are evaluated by the firm's independent quality-assurance department to assess procedural accuracy (e.g., greeting the customer and addressing them by name) and informational accuracy (e.g., providing correct information and appropriate solutions). These are the same metrics used by TEMPO for routine monitoring and evaluation of agents.

Retention outcomes are constructed from internal HR records that track employment status and start/end dates at the agent level. We define attrition as a permanent exit from the firm. HR records also capture office-day attendance for treated agents, which we use to document compliance with the intervention.

In addition, we observe baseline characteristics and job attributes (e.g., gender, age, tenure, location, team assignment, and shift schedules), which we use for balance checks, heterogeneity analysis, and robustness.

Surveys: worker experience and interactions. To complement the administrative outcomes, we administered two online surveys during working hours. A baseline survey was fielded in July 2024, before the intervention began (and after the study sample was finalised), and an endline survey was fielded in April 2025, at the end of the intervention period. We collected demographics and background information (including prior fully remote work experience), along with attitudes toward remote work and office attendance. The endline survey focused on experiences during the trial, with modules on perceived team cohesion, feedback frequency, satisfaction with management, work–life balance, job satisfaction, and organisational attachment, as well as support for the monthly office-day policy. Response rates were 100% at baseline (248/248) and 90% at endline (223/248).

We also fielded short post-visit interaction surveys for treated agents after each office day. In the week following a visit, agents reported whom they interacted with during the office day (e.g., morning, around lunch, afternoon; shared meals or coffee breaks) and whether they communicated with colleagues in the subsequent week, naming counterparts and indicating the communication channel (e.g., phone or messaging). These data are designed to characterise how in-person contact maps into subsequent communication and potential peer spillovers.

Survey responses are complemented with administrative logistics for office days. We observe seating locations inside the office and shuttle assignments for each visit, which allow us to construct measures of physical proximity and shared travel. As mentioned above, seating was randomised and enforced during the first three office visits; from the fourth visit onward, agents chose their seats. Shuttle rosters identify which agents travelled together on each office day. Taken together, these data provide multiple measures of co-location, interaction, and follow-up communication.

Main sample. Our preferred specifications use the full administrative sample (all 248 agents). Survey-based analyses use the corresponding respondent samples, and when required we restrict attention to balanced panels (e.g., respondents observed at both baseline and endline), which we specify in each table and figure.

3.1 Sample characteristics and balance

Sample description.

Table A.1 reports baseline summary statistics for worker characteristics, productivity, and service-quality measures. At baseline, agents are on average about 23 years old. Women com-

prise the majority of the workforce (81% in control and 74% in treatment). Educational attainment is high for a call-centre setting: roughly 70% of agents hold a bachelor's degree and about 25–28% hold an associate degree. Around 20–23% are married, 11–13% have children, and baseline tenure is just under half a year on average.

Balance on observables and baseline outcomes. The treatment and control groups are similar across all observed covariates: gender, age, education, marital status, parental status, and tenure are closely aligned, and none of the differences are statistically significant. Baseline productivity measures (including calls per hour, average call duration, talk time, administrative time, and break time) also show no meaningful differences. The only variable with a small statistically significant difference is hold time, which is slightly higher in the treatment group. Baseline customer ratings and audit scores are indistinguishable across groups.

Figure A.2 provides a complementary balance check using a linear probability model, plotting the association between treatment assignment and each baseline characteristic. Coefficients are close to zero and confidence intervals span zero throughout, indicating that no observable variable systematically predicts treatment status.

Taken together, Table A.1 and Figure A.2 confirm that the randomisation produced well-balanced groups prior to the intervention. This supports the internal validity of the design: post-intervention differences can be attributed to the monthly office-day policy rather than pre-existing differences across agents.

4 Empirical Strategy

4.1 Main specification

Our empirical strategy exploits the random assignment of agents to monthly office days and follows outcomes over time to capture both the timing and persistence of the intervention's effects. Because the policy was introduced simultaneously for all treated agents and repeated once per month, we use an event-study design as our baseline specification. This approach allows treatment effects to evolve flexibly over time, rather than imposing a constant effect from the outset.

At the agent–day level, we estimate the following specification:

$$y_{iymwd} = \sum_{l=-12}^{-1} \beta_l \left(\text{OfficeVisit}_i \times P_{ym}^l\right) + \sum_{l=1}^{9} \beta_l \left(\text{OfficeVisit}_i \times P_{ym}^l\right) + \mu_i + \gamma_d + \gamma_j + \gamma_{ym} + \epsilon_{iymwd}, \tag{1}$$

where indices i, y, m, w, and d denote the agent, year, month of the year, week of the year, and day of the week, respectively, and j indexes the team leader. The outcome y_{iymwd} is measured at the agent–date level and captures individual performance, including calls per hour, call duration and its subcomponents (talk time, hold time, and administrative time), as well as break time. OfficeVisit_i is an indicator equal to one for agents assigned to attend the office one day per month and zero for agents who remained fully remote. P_{ym}^l is an indicator equal to one if calendar month (y, m) is l months away from the final pre-intervention month, and zero otherwise.

Event time is defined in months relative to July 2024, the final month before the intervention began. We normalise July 2024 ($l = 0$) as the omitted category, so all coefficients are interpreted relative to the last fully remote baseline. Under this convention, August 2024 corresponds to Month 1 ($l = 1$), and the intervention spans nine months in total, through April 2025 ($l = 9$).

The coefficients β_l trace how outcomes for treated agents evolve over time relative to the control group. Estimates for months prior to the intervention ($l < 0$) capture any differential trends before treatment and provide a direct test of the parallel trends assumption. Estimates for months after the intervention ($l \geq 1$) describe the dynamic response to repeated office visits, allowing effects to emerge gradually as agents experience ongoing in-person interaction. This structure is well suited to our setting, where learning, coordination, and feedback may strengthen over time rather than producing immediate changes.

While the event-study specification serves as the baseline model, some results are easier to interpret when post-intervention dynamics are summarised over broader time windows. For this reason, in the main regression tables we group months into two periods: an early experimental phase covering Months 1–5, and a late experimental phase covering Months 6–9.

This grouping is motivated by the gradual patterns observed in the event-study estimates and allows us to assess whether the effects of monthly office days emerge immediately or build up over time with repeated in-person interaction. The coefficients reported in the tables correspond to average effects over these respective periods and provide a compact summary of the underlying dynamic treatment effects captured by the baseline event-study specification.

All specifications include agent fixed effects (μ_i), which absorb time-invariant differences across agents such as baseline ability, experience, and work style. Team leader fixed effects (γ_j) capture persistent differences in management practices and team organisation. Day-of-week fixed effects (γ_d) account for systematic variation in call volumes and customer demand. Year–month fixed effects (γ_{ym}) control for seasonality and common shocks over time. Standard errors are clustered at the agent level throughout to allow for serial correlation in individual outcomes.

For outcomes that are not observed at high frequency or do not lend themselves naturally to a dynamic analysis (specifically, retention and survey-based outcomes), we estimate simpler difference-in-differences specifications that compare treated and control agents before and after the intervention. These models rely on the same random assignment and identifying assumptions but summarise treatment effects over broader post-intervention periods.

5 Results

5.1 Productivity

Figure 1 plots weekly averages of calls per hour (our primary measure of productivity) for treated and control agents over the pre-intervention, intervention, and post-intervention periods with 95% confidence bands. Although purely descriptive, the figure is informative about the timing and persistence of the intervention's effects. In the months prior to August 2024, the two series track each other closely, providing evidence that there are no differential pre-trends.

After monthly office days begin, the treatment group gradually departs from the control group. The divergence is not immediate: in the initial months, differences in calls per hour are small.

Over subsequent months, however, treated agents increasingly outperform their fully remote peers, and the gap widens as exposure to repeated in-person contact accumulates. By the end of the experimental period, the difference is sizeable and clearly visible. The post-intervention period provides suggestive evidence on persistence. Following the end of monthly office days, treated agents continue to exhibit higher calls per hour than the control group, with no sharp convergence in the immediate months after the program ends.

Next, we formalise and quantify these dynamics in the regression analysis. Figure 2 plots dynamic treatment effects in calls per hour, reporting both a *total* treatment–control difference and a *within-agent* (agent fixed effects) estimate. The two series tell a consistent story: effects build gradually after the first office visit and remain positive after the last visit, with no sharp reversal. Table A.2 quantifies these patterns using specifications with agent fixed effects and standard errors clustered at the agent level.

The table shows that productivity gains emerge only with repeated exposure. In the early phase (Months 1–5), the treatment effect on calls per hour is close to zero (-0.02). In the late phase (Months 6–9), calls per hour rises by 0.43 (4.0% of the pre-treatment mean), and in the post-experiment period (Months 10–14) it increases further to 0.63 (5.8% of the pre-treatment mean), suggesting that gains do not dissipate immediately once office days end.

The increase in calls per hour is driven by shorter calls. Average call duration is essentially unchanged in Months 1–5 (-0.04 seconds), but falls by 13.21 seconds in Months 6–9 (3.8% of the pre-treatment mean) and by 17.60 seconds in Months 10–14 (5.0% of the pre-treatment mean). Columns 5–10 decompose call duration into talk time, administrative time, and hold time. The reduction is driven primarily by shorter talk time, with little systematic movement in administrative or hold time. Finally, break time is broadly unchanged, indicating that higher output is not achieved by compressing breaks.

We also examine heterogeneity in treatment effects across observable worker characteristics. Figure A.3 shows broadly similar effects by gender, age, baseline performance, marital status, and parental status; differences across subgroups are small and statistically indistinguishable from zero, suggesting the productivity gains are not concentrated in a particular group.

5.2 Service quality

A natural concern is that the productivity gains documented above reflect a speed–quality trade-off: treated agents may handle more calls by cutting corners, reducing the quality of customer interactions, or deviating from protocols. Table A.3 addresses this concern by reporting treatment effects on two service-quality outcomes measured at the agent–month level: customer satisfaction ratings and internal random audit scores. Across specifications, we find no evidence that monthly office days affected service quality. In the agent fixed-effects specifications (columns 2 and 4), the estimated treatment effects are small and statistically indistinguishable from zero in both the early and late phases of the intervention. For customer ratings, the treatment–post interactions are 0.93 (s.e. 1.01) in Months 1–5 and -0.76 (s.e. 1.40) in Months 6–9; for audit scores, the corresponding estimates are -0.03 (s.e. 0.02) and 0.01 (s.e. 0.03).

Taken together, these results suggest that the increase in calls per hour reflects greater efficiency in handling customer requests rather than a deterioration in customer experience or procedural

compliance.

5.3 Attrition

Fully remote organisations (including call-centre operators such as TEMPO and many firms worldwide) often cite elevated turnover as a key operational challenge. We therefore examine whether occasional, predictable in-person contact can reduce attrition in an otherwise fully remote setting.

Figure 3 plots cumulative attrition rate for treated and control agents. Once monthly office days start, attrition rises more slowly in the treated group and the gap widens steadily over time. The gap also persists after the last office visit, implying the policy does more than postpone exits.

Figure 4 shows that these retention gains are concentrated among higher performers. In the control group, higher baseline productivity is associated with a higher probability of exit, consistent with stronger outside options. Under monthly office days, this relationship largely disappears.

Overall, the figures show that monthly office visits reduce turnover and, in particular, help retain top performers, an economically important margin in settings where training is costly and service delivery depends on experienced agents.

5.4 Mechanisms

The productivity gains (higher calls per hour driven by shorter calls, without deterioration in service quality) and the retention effects documented above point to channels operating through interaction, learning, and feedback rather than short-run increases in effort. In this section, we provide evidence on these mechanisms using (i) random variation in peer proximity generated by seating assignments on office days and (ii) endline survey measures of communication and managerial support.

Office days create new peer links and sustained communication. We start by testing whether in-person co-location translates into subsequent interaction. During the first three office visits, treated agents were assigned seats according to a randomised seating plan. Table 1 shows that agent pairs who were randomly allocated as desk neighbours were substantially more likely to report communicating in the week after the office visit: the probability rises by 11 percentage points relative to a mean of 9%. Self-reported interactions during the office day are similarly predictive of follow-up communication: pairs that interacted during the morning, lunch, or afternoon breaks are 18–19 percentage points more likely to communicate afterward. Because desk neighbours were assigned mechanically rather than chosen, these patterns isolate the causal role of physical proximity in generating new peer links that persist beyond the office day itself.

Peer exposure generates productivity spillovers, especially from high performers.

We next ask whether this proximity translates into learning spillovers. Table 2 relates post-visit productivity to the baseline productivity of randomly assigned neighbours. Agents seated next to more productive peers experience larger subsequent increases in calls per hour: the interaction between the post period and neighbours' baseline calls per hour is positive and statistically significant. Importantly, neighbours' baseline productivity does not predict whether pairs communicate in the following week (columns 2–4 and 8–10), and the effect of being assigned as a desk neighbour on communication remains large (columns 5–7). This combination of results suggests that spillovers are not driven by selective matching into follow-up communication with high performers; instead, brief

exposure on the office day itself appears sufficient to transmit work practices and problem-solving approaches.

Treated agents report more feedback and stronger organisational attachment. Survey evidence is consistent with these interaction-based channels. Figure A.4 shows that treated agents are more supportive of the monthly office-day policy and report more frequent performance feedback from managers, more effective team communication, and a stronger perceived fit with the firm's culture at endline. By contrast, broader measures of overall life satisfaction and work–life balance move little. This pattern points to mechanisms operating through workplace relationships and professional support (coaching, communication, and cultural integration) rather than a general improvement in well-being.

Taken together, the evidence supports an interaction-and-learning mechanism: monthly office days create opportunities for peer exposure and in-person feedback that strengthen communication networks and generate spillovers from high performers. These channels are consistent with the gradual build-up of productivity (driven by shorter calls) and the improved retention outcomes observed in the main analysis.

6 Cost–benefit analysis

This section provides a back-of-the-envelope cost–benefit calculation from the firm's perspective. We focus on the two margins where we observe clear treatment effects (productivity and retention) and monetise benefits using conservative assumptions. All figures are expressed in USD.

Programme costs. The intervention required treated agents to attend the Şanlıurfa office one day per month. The incremental cost per visit was \$9.63 per employee, covering transport, lunch, refreshments, and use of existing office space. Over nine visits and 124 treated agents, total programme expenditure was \$10,749. The firm did not incur additional fixed costs (e.g., new space, equipment, or staffing), as the office and shuttle services were already in place.

Turnover savings. The intervention reduced attrition by 7 percentage points relative to the control group. Applied to the 124 treated agents, this implies 8.7 fewer separations over the nine-month intervention window.[2] We value each avoided separation at \$714, an estimate provided by the firm that covers recruitment, initial training, and the productivity ramp-up of new hires. This yields turnover savings of \$6,198. This estimate is conservative because it excludes managerial time spent interviewing and onboarding, IT and equipment costs, and productivity losses borne by teams while vacancies are filled.

Productivity gains. By the end of the intervention period, treated agents handled 3.4% more calls per hour. To avoid overstating benefits given the gradual ramp-up in effects, we assume that productivity gains increase linearly from zero to 3.4% over the nine months, implying an average gain of 1.7% over the intervention window. We value this gain using an employer labour cost of 130 TRY per hour (approximately \$4) and assume each treated agent worked 1,560 hours over the nine-month period. Under these assumptions, the implied productivity benefit is \$12,786.[3] This

[2] We compute avoided separations as $0.07 \times 124 = 8.68$.

[3] This is computed as $124 \times 1{,}560 \times \$4 \times 0.017$.

calculation captures only direct output gains and excludes any additional benefits from potential reductions in supervisory burden or improved customer experience.

Net benefit and benefit–cost ratio. Total monetised benefits equal \$18,984 (= \$6,198 turnover savings + \$12,786 productivity gains). Netting out programme costs of \$10,749 yields a net benefit of \$8,235. The implied benefit–cost ratio is 1.77, meaning that each dollar of programme spending generated about \$1.8 in measurable benefits under conservative assumptions.

Taken together, the monthly office-day policy appears cost-effective. Roughly one-third of the monetised gains come from reduced turnover, with the remainder from higher output. Because several relevant cost components (e.g., managerial time, vacancy costs, and other operational frictions) are excluded from the turnover calculation, and because we do not monetise any potential benefits to service quality or team functioning, the true return to the programme may be higher. Overall, the results suggest that occasional in-person contact can generate meaningful value in fully remote settings, particularly when firms can leverage existing office infrastructure.

7 Conclusion

This paper asks whether occasional, coordinated in-person contact can improve outcomes in an otherwise fully remote workplace. We study a nine-month randomised controlled trial in a large Turkish call centre in which treated agents worked from the office one day per month and otherwise remained fully remote. The intervention was intentionally light-touch: it leveraged existing infrastructure and kept pay, schedules, tasks, and performance evaluation unchanged.

Three findings stand out. First, monthly office days raise productivity. Calls per hour increase gradually, with effects emerging only after repeated exposure and remaining positive after the final visit. The gains are driven by shorter calls (primarily reduced talk time) rather than by reductions in breaks or shifts in administrative or hold time. Second, service quality is unchanged. Customer ratings and internal audit scores do not deteriorate, suggesting that higher output reflects greater efficiency rather than a speed–quality trade-off. Third, retention improves. Attrition accumulates more slowly among treated agents and the gap persists beyond the intervention window; moreover, the policy mitigates the selective loss of high performers.

Mechanism evidence points to interaction and learning. Treated agents report more frequent managerial feedback, more effective team communication, and a stronger perceived cultural fit. Randomised seating assignments on office days create new peer links that extend beyond the visit itself, and exposure to more productive neighbours generates positive spillovers in subsequent performance. Together, these results indicate that even infrequent face-to-face contact can strengthen feedback, coordination, and the transfer of tacit know-how in ways that are difficult to replicate through remote communication alone.

From the firm's perspective, the policy is inexpensive and appears cost-effective. Modest spending on transport and meals is outweighed by monetised gains from higher productivity and lower turnover, yielding a conservative benefit–cost ratio of 1.77.

More broadly, the results emphasise that the performance of remote work depends not only on location but also on organisational design. In fully remote service settings, a small and predictable amount of in-person contact can deliver durable gains without undermining service quality or sacrificing the flexibility that makes remote work attractive.

References

Cevat Giray Aksoy, Jose Maria Barrero, Nicholas Bloom, Steven J. Davis, Mathias Dolls, and Pablo Zarate. The global persistence of work from home. *Proceedings of the National Academy of Sciences of the United States of America*, 122(27):e2509892122, 2025a. doi: 10.1073/pnas.2509892122.

Cevat Giray Aksoy, Nicholas Bloom, Steven J. Davis, Victoria Marino, and Cem Özgüzel. Remote work, employee mix, and performance. NBER Working Paper 33851, National Bureau of Economic Research, 2025b. URL https://www.nber.org/papers/w33851. Working paper.

David Atkin, Antoinette Schoar, and Sumit Shinde. Working from home, worker sorting and development. NBER Working Paper 31515, National Bureau of Economic Research, 2023. URL https://www.nber.org/papers/w31515.

Jose Maria Barrero, Nicholas Bloom, and Steven J. Davis. Why working from home will stick. NBER Working Paper 28731, National Bureau of Economic Research, 2021. URL https://www.nber.org/papers/w28731.

Daniele Battiston, Jordi Blanes i Vidal, and Tom Kirchmaier. Face-to-face communication in organizations. *Review of Economic Studies*, 88(2):574–609, 2021. doi: 10.1093/restud/rdaa060.

Nicholas Bloom, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. Does management matter? evidence from India. *Quarterly Journal of Economics*, 128(1):1–51, 2013. doi: 10.1093/qje/qjs044.

Nicholas Bloom, James Liang, John Roberts, and Zhichun Jenny Ying. Does working from home work? evidence from a chinese experiment. *Quarterly Journal of Economics*, 130(1):165–218, 2015. doi: 10.1093/qje/qju032.

Nicholas Bloom, Ruobing Han, and James Liang. Hybrid working from home improves retention without damaging productivity. *Nature*, 630(8018):920–925, 2024. doi: 10.1038/s41586-024-07500-6.

Nicholas Bloom, Gordon Dahl, and David-Olivier Rooth. Work from home and disability employment. *American Economic Review: Insights*, 2026. doi: 10.1257/aeri.20240538. Forthcoming.

Shelby R. Buckman, Jose Maria Barrero, Nicholas Bloom, and Steven J. Davis. Measuring work from home. NBER Working Paper 33508, National Bureau of Economic Research, 2025. URL https://www.nber.org/papers/w33508.

Prithwiraj Choudhury, Cirrus Foroughi, and Barbara Zepp Larson. Work-from-anywhere: The productivity effects of geographic flexibility. *Strategic Management Journal*, 42(4):655–683, 2021. doi: 10.1002/smj.3251.

Prithwiraj Choudhury, Miguel Espinosa, Tarun Khanna, Christos Makridis, and Kyle Schirmann. Internal communication and remote work. *SSRN Electronic Journal*, 2025. doi: 10.2139/ssrn.5106503. URL https://ssrn.com/abstract=5106503. Available at SSRN.

Natalia Emanuel and Emma Harrington. Working remotely? selection, treatment, and the market for remote work. *American Economic Journal: Applied Economics*, 16(4):528–559, 2024. doi: 10.1257/app.20230376.

Michael Gibbs, Friederike Mengel, and Christoph Siemroth. Work from home and productivity: Evidence from personnel and analytics data on information technology professionals. *Journal of Political Economy Microeconomics*, 1(1):7–41, 2023. doi: 10.1086/721803.

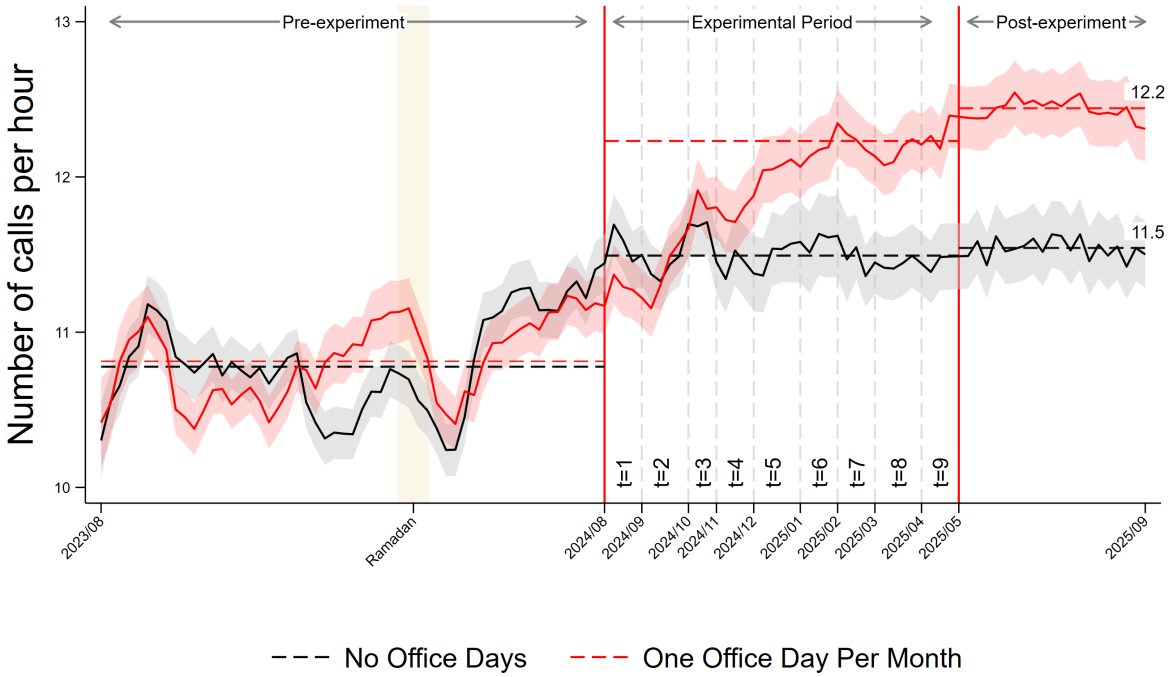
Edward P. Lazear, Kathryn L. Shaw, and Christopher T. Stanton. The value of bosses. *Journal of Labor Economics*, 33(4):823–861, 2015. doi: 10.1086/681097.

Davide Luca, Cem Özgüzel, and Zhiwu Wei. The new geography of remote jobs in Europe. *Regional Studies*, 59(1):1–19, 2025. doi: 10.1080/00343404.2024.2352526.

Jan Manthei, Dirk Sliwka, and Timo Vogelsang. Talking about performance or paying for it? a field experiment on performance reviews and incentives. *Management Science*, 69(4):2198–2216, 2023. doi: 10.1287/mnsc.2022.4431.

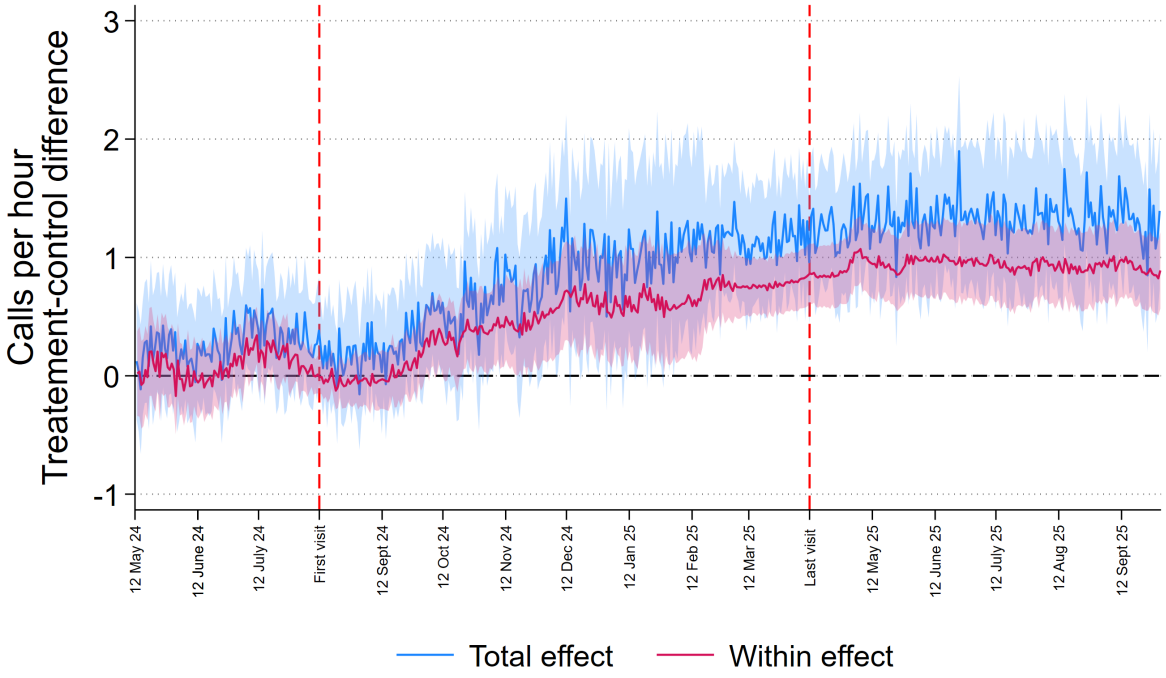
Figures and Tables

Figure 1: Calls per hour before, during and after the RCT



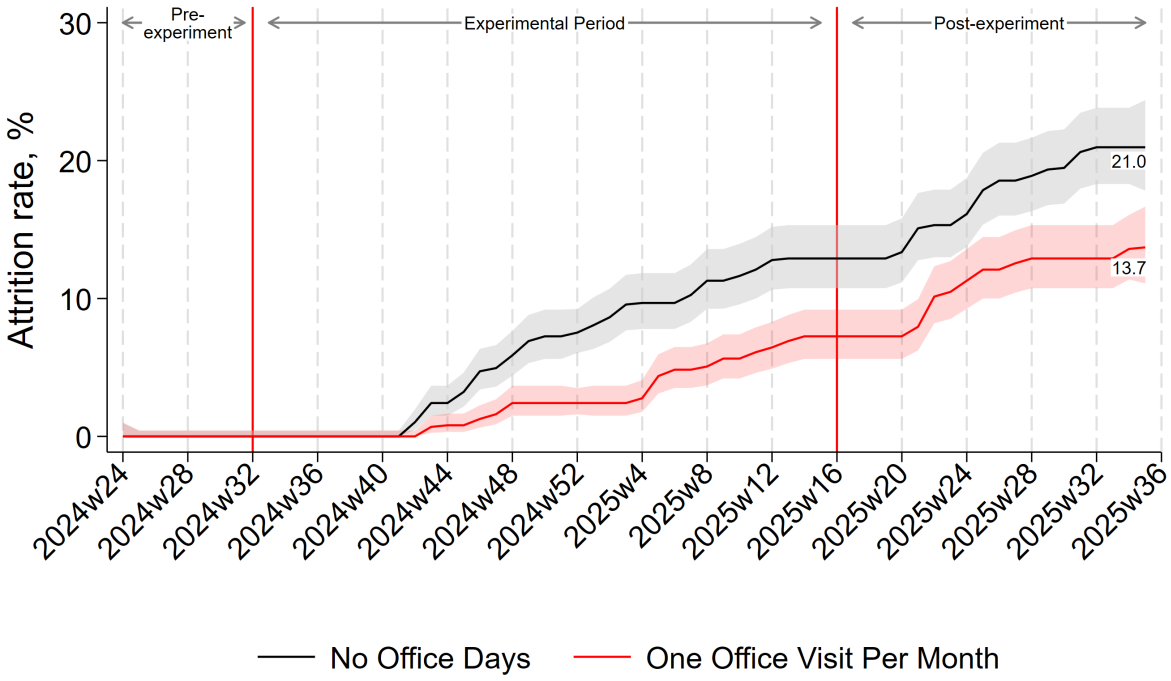
Notes: This figure shows weekly means of *Call duration in seconds per call* and *Calls per hour* for control (No Office Days) and treatment (One Office Day Per Month) groups. Weekly means are calculated using agent-day level data. Horizontal dashed lines represent means in the *pre-experiment* and *post-experiment* periods among the *No Office Days* (black dashed line) and *One Office Day per Month* (red dashed line) groups. Means of *Number of calls, per hour* in pre-treatment period for control and treatment groups are 10.8 and 10.8, respectively, while the means in post-experimental period for control and treatment groups are 11.5 and 12.2, respectively. The vertical red lines at August 2024 and May 2025 indicates the start and end of the treatment period, respectively. Shaded red and grey areas depict 95% confidence intervals around the mean.

Figure 2: Agent and total effects



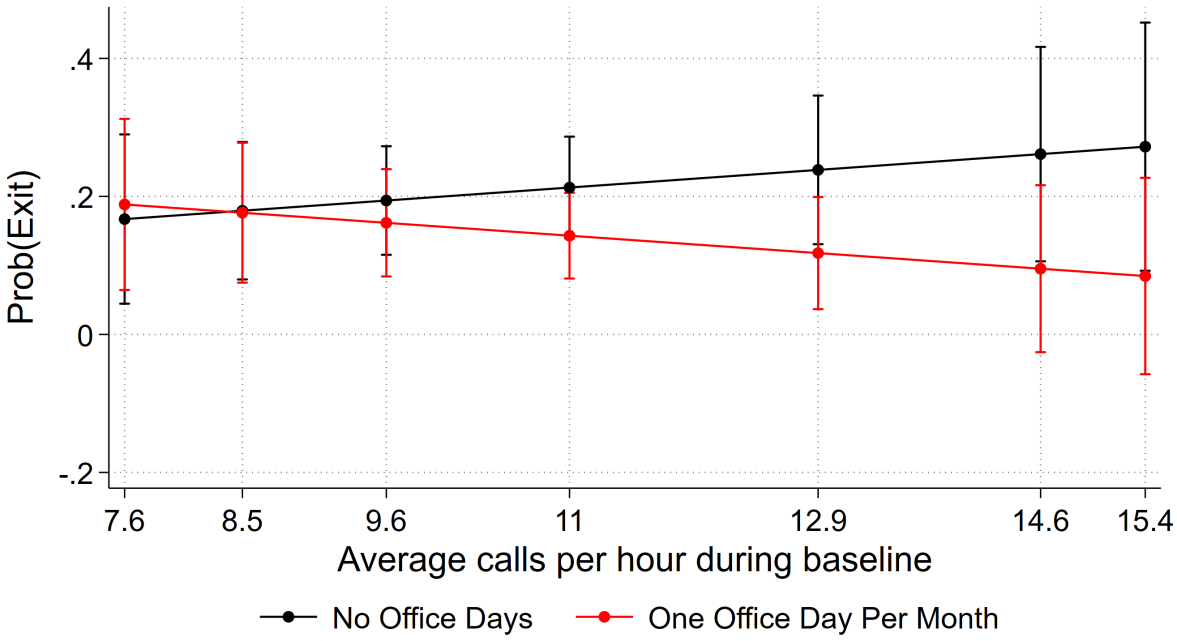
Notes: Figure showing point estimates from two regressions. The *Total effect* (blue line) corresponds to a regression of calls per hour on a treatment dummy, date time dummies and the interaction of the treatment dummy with each of the date time dummies. The omitted category is the day before the first office visit. For these estimations calls per hour is first smoothed using a five period moving average (two lagged and two forward terms). To estimate the *Within effect* (red line) the regression is further saturated with agent fixed effects. Point estimates shown correspond to the interactions. Standard errors are clustered at the agent level and 95% confidence intervals are shown.

Figure 3: Attrition is lower among agents who went to the office



Notes: This figure shows the attrition rate among agents in the treatment group (red line), who attended the office, compared to those in the control group (blue line), who remained fully remote. These lines plot weekly means calculated using agent-day level data. The vertical red lines mark the weeks of the first and last office visits for the treated group. Shaded areas represent 95 percent exact binomial confidence intervals for proportions.

Figure 4: Evidence that visiting the office once per month helps retain higher performers



Notes: Figure plotting predicted values of *Prob(Exit)* (vertical axis) at specific values of *Average calls per hour during baseline* (horizontal axis). Based on a regression of *Exit* dummy on a *One Office Day Per Month* dummy, *Average calls per hour during baseline* and the interaction of these two variables. The sample is the agent level. Standard errors are heteroscedasticity robust and whiskers depict 95% confidence intervals.

Table 1: Agents were more likely to report communicating with a colleague after an office day when that colleague had been randomly assigned to sit next to them

	(1)	(2)	(3)	(4)
Seating plan neighbour	0.11***			
	(0.02)			
Morning interaction		0.18***		
		(0.01)		
Lunch interaction			0.19***	
			(0.01)	
Afternoon interaction				0.18***
				(0.01)
Adj. R-squared	0.00	0.04	0.04	0.04
Number of observations	7,976	7,976	7,976	7,976
Outcome mean	0.09	0.09	0.09	0.09

Notes: This table reports estimates from linear regressions. The dependent variable is a dummy equal to 1 if an agent pair communicated in the week after the office visit and 0 if they did not communicate. *Seating plan neighbour* is a dummy variable equal to 1 if an agent pair were (randomly allocated via a seating plan) to be neighbours in the office in the during the first three visits and 0 otherwise. *Morning interaction* is a dummy equal to 1 if an agent pair interacted during a morning coffee break together. *Lunch interaction* is a dummy equal to 1 if an agent pair interacted during a lunch break together. *Afternoon interaction* is a dummy equal to 1 if an agent pair interacted during an afternoon coffee break together. Follow-up communication, morning, lunch and afternoon interactions are self-reported in surveys at the end of each week following an office visit. The data includes all possible agent pairs and focuses on the first three office visits and subsequent weeks where the seating plan was randomised and enforced. Standard errors are clustered at the agent level. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 2: Being seated next to highly productive agents has is positive spillovers but neighbour productivity is not associated with the likelihood of subsequent communications

Dependent variable:	Calls per hour	Dummy equal to 1 if an agent pair communicated in the week after the office visit and 0 if they did not communicate.								
Sample:	Production panel	All agent pairs						Seating plan agent pairs		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Average baseline calls per hour of neighbours	0.00									
	(0.02)									
Post	0.75***									
	(0.14)									
Post × Average baseline calls per hour of neighbours	0.04**									
	(0.02)									
Follow contact has high productivity		-0.01			-0.01			0.03		
		(0.01)			(0.01)			(0.05)		
Follow contact has very high productivity			-0.00			-0.00			0.05	
			(0.01)			(0.01)			(0.06)	
Follow contact: baseline average calls per hour				-0.00			-0.00			-0.00
				(0.00)			(0.00)			(0.01)
Seating plan neighbour					0.12***	0.12***	0.12***			
					(0.02)	(0.02)	(0.02)			
Adj. R-squared	0.04	0.00	0.00	0.00	0.01	0.01	0.01	0.00	0.00	0.00
Number of observations	144,769	7,852	7,852	7,852	7,852	7,852	7,852	248	248	248

Notes: This table reports estimates from linear regressions. The dependent variable as indicated in column headings. *Average baseline calls per hour of neighbours* is calculated using odd calendar dates in the pre-experiment period. Odd calendar dates are dropped from the same before the regression in estimated. *Follow contact has high productivity* is dummy variable equal to one if the potential follow-up agent has above median baseline calls per hour and 0 otherwise. *Follow contact has high productivity* is dummy variable equal to one if the potential follow-up agent has above 75th percentile baseline calls per hour and 0 otherwise. *Seating plan neighbour* is a dummy variable equal to 1 if an agent pair were (randomly allocated via a seating plan) to be neighbours in the office in the during the first three visits and 0 otherwise. Follow-up communication is self-reported in surveys at the end of each week following an office visit. The data includes all possible agent pairs and focuses on the first three office visits and subsequent weeks where the seating plan was randomised and enforced. Standard errors are clustered at the agent level. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Appendix

Figure A.1: During the intervention

(a) Working at home environment



(b) Transport to and from the office

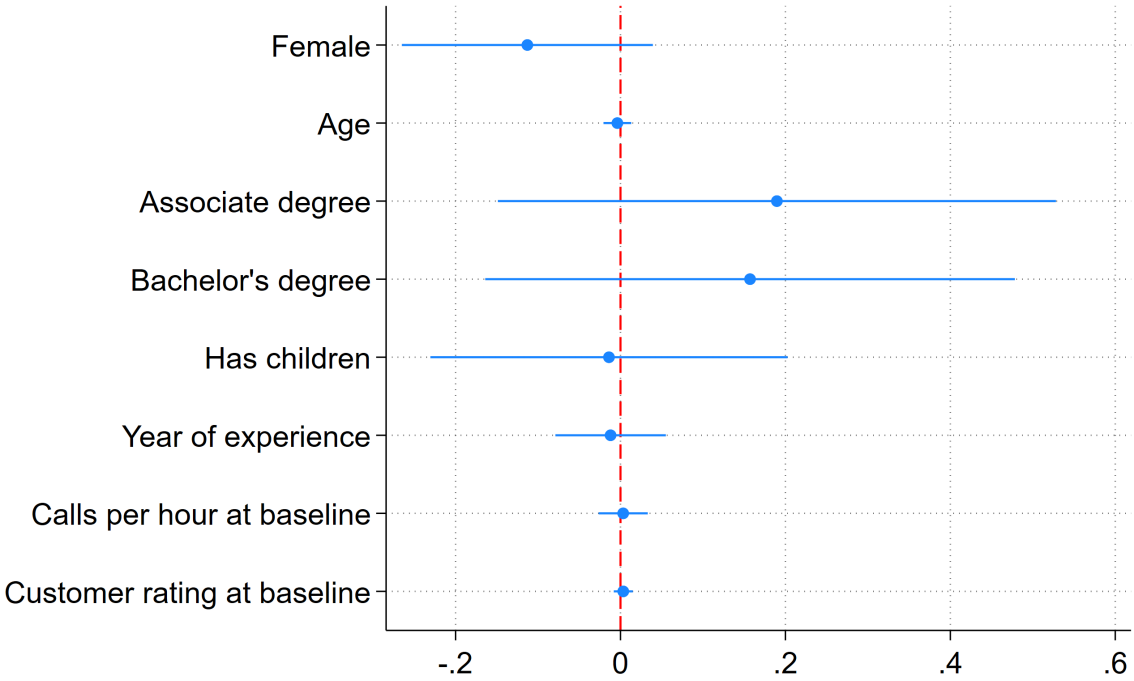


(c) Office-day activities



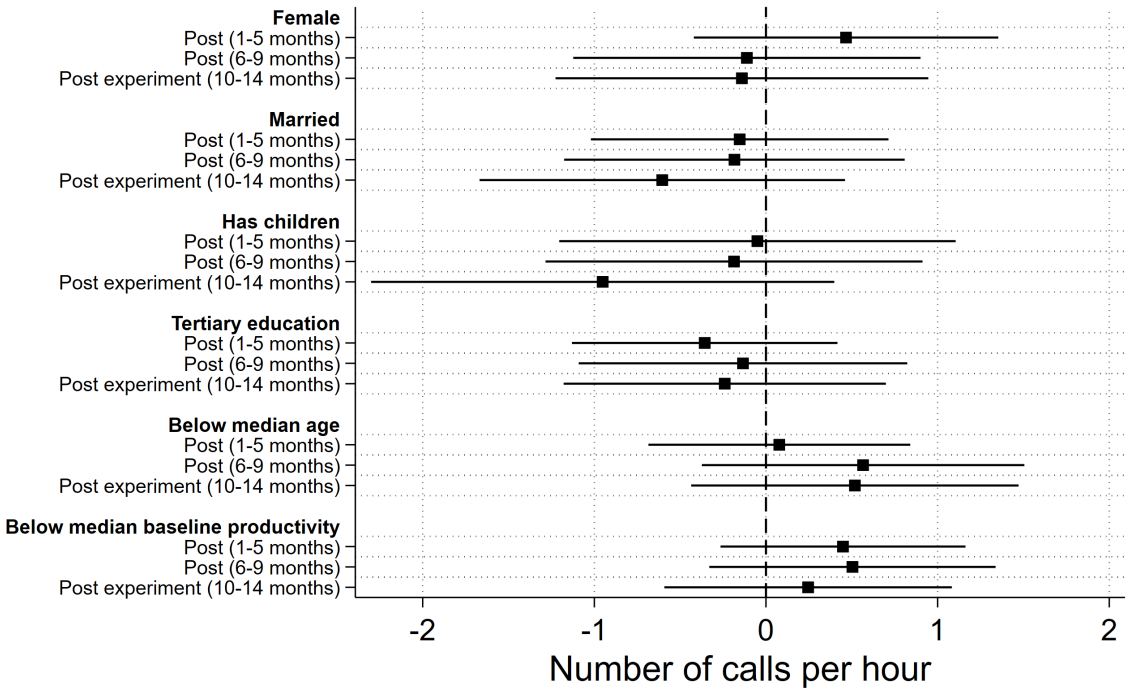
Notes: Illustrative photos from the intervention period showing: employees working from home (Panel (a)); transport arranged by the company and group activities held during monthly office days (Panel (b)); and the working environment on office days and team social activities (Panel (b)).

Figure A.2: Balance test for treatment assignment



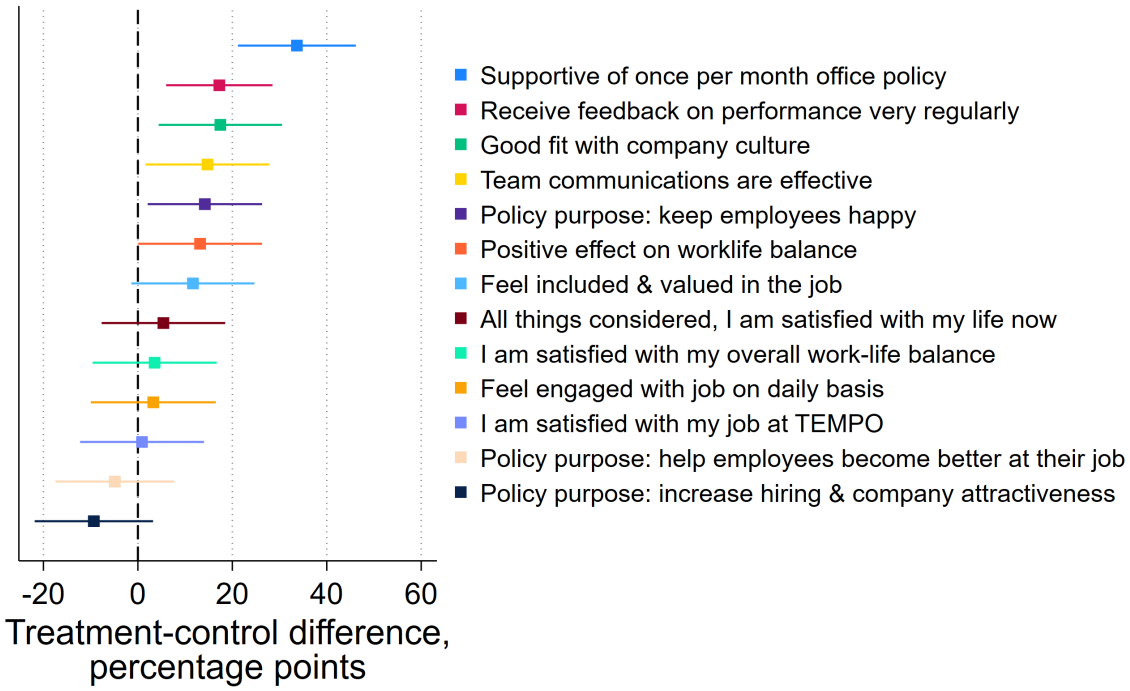
Notes: Figure shows coefficient estimates from a linear probability model. The dependent variable equals 1 if an agent is treated. The data include one observation per agent ($N = 248$). Standard errors are robust to heteroskedasticity, and whiskers show 95 percent confidence intervals.

Figure A.3: No statistically significant evidence of heterogeneity in productivity effects



Notes: This figure shows point estimates from six separate linear regressions. The point estimates shown correspond a triple interaction between *One Office Day Per Month*, a post period and individual characteristic as indicated on the vertical axis. *Below median productivity* is calculated using odd calender days in the pre-experiment period, these observations are dropped before estimating the regression. Standard errors are clustered at the agent level and whisker represent 95% confidence intervals.

Figure A.4: At endline treated agents were more supportive of the policy and reported more frequent manager feedback, stronger cultural fit



Notes: Figure showing linear probability regressions, each coefficient corresponds to a separate regression of the outcome (listed in the legend) on a treatment dummy. All outcomes are dummy variables. Coefficients have been rescaled to percentage points. The sample corresponds to the endline sample where N=223. Whiskers depict 90% and 95% confidence intervals and standard errors are heteroskedasticity robust.

Table A.1: Observables are balanced at baseline

	(1)	(2)	(3)	(4)	(5)	(6)
	N		Mean		Mean	p-value
	Control	Treatment	Control	Treatment	Difference	Difference
A. Worker characteristics						
Female	124	124	0.81	0.74	-0.07	0.17
Age	124	124	23.60	23.31	-0.29	0.62
High school	124	124	0.05	0.02	-0.02	0.31
Associate degree	124	124	0.25	0.28	0.03	0.57
Bachelor's degree	124	124	0.70	0.69	-0.01	0.89
Married	124	124	0.23	0.20	-0.03	0.54
Has children	124	124	0.13	0.11	-0.02	0.70
Year of experience	124	124	0.47	0.43	-0.04	0.77
B. Productivity variables at baseline						
Number of calls per hour	36,435	34,846	10.79	10.82	0.03	0.91
Call duration in seconds per call	36,435	34,846	349.57	349.37	-0.20	0.98
Break time in minutes per day	36,402	34,831	33.27	35.07	1.80	0.20
Talk time in seconds per call	36,435	34,846	345.27	344.13	-1.13	0.90
Admin time in seconds per call	36,435	34,846	2.90	2.93	0.04	0.53
Hold time in seconds per call	36,435	34,846	1.22	2.07	0.85	0.03
C. Service quality variables at baseline						
Customer rating	1,243	1,187	44.02	44.31	0.29	0.67
Random audit rating	1,243	1,187	0.17	0.17	0.00	0.86

Notes: The table reports summary statistics of agent characteristics (Panel A), productivity outcomes (Panel B), and performance outcomes (Panel C). Panel A presents worker characteristics at the first observation. Panels B and C report productivity and performance variables at the agent-date level.

Table A.2: Early, late and post experiment treatment effects

	Calls per hour		Average call duration (seconds per call)		Talk time (seconds per call)		Admin time (seconds per call)		Hold time (seconds per call)		Break time (minutes per day)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
One Office Day Per Month	0.03		-0.27		-1.19		0.04		0.84**		1.85	
	(0.27)		(8.83)		(8.82)		(0.06)		(0.39)		(1.40)	
Post (1-5 months)	0.71***	0.78***	-19.13***	-21.37***	-19.42***	-21.56***	0.08*	0.07*	0.13	0.06	8.84***	8.81***
	(0.13)	(0.12)	(3.82)	(3.60)	(3.80)	(3.58)	(0.04)	(0.04)	(0.15)	(0.13)	(0.76)	(0.75)
Post (6-9 months)	0.69***	0.85***	-18.62***	-23.77***	-19.04***	-23.98***	0.02	0.02	0.35	0.18	10.33***	10.87***
	(0.17)	(0.14)	(4.77)	(4.14)	(4.77)	(4.14)	(0.05)	(0.05)	(0.22)	(0.18)	(0.92)	(0.84)
Post experiment (10-14 months)	0.74***	0.85***	-20.66***	-24.10***	-21.12***	-24.40***	-0.03	-0.01	0.45**	0.28	11.58***	12.38***
	(0.17)	(0.14)	(4.85)	(4.12)	(4.84)	(4.11)	(0.06)	(0.05)	(0.21)	(0.19)	(1.02)	(0.91)
One Office Day Per Month	0.07	-0.02	-3.35	-0.04	-3.01	0.20	0.05	0.06	-0.29	-0.21	0.94	0.96
× Post (1-5 months)	(0.19)	(0.18)	(5.51)	(5.07)	(5.51)	(5.08)	(0.06)	(0.06)	(0.30)	(0.29)	(1.17)	(1.15)
One Office Day Per Month	0.67***	0.43**	-21.34***	-13.21**	-20.70***	-12.80**	-0.09	-0.09	-0.43	-0.26	0.83	0.36
× Post (6-9 months)	(0.23)	(0.20)	(6.51)	(5.64)	(6.53)	(5.68)	(0.06)	(0.06)	(0.36)	(0.33)	(1.34)	(1.28)
One Office Day Per Month	0.87***	0.63***	-25.10***	-17.60***	-24.28***	-16.84***	-0.10	-0.10	-0.50	-0.47	2.07	1.16
× Post experiment (10-14 months)	(0.23)	(0.20)	(6.51)	(5.53)	(6.52)	(5.54)	(0.07)	(0.07)	(0.37)	(0.36)	(1.51)	(1.42)
Adj. R-squared	0.05	0.81	0.04	0.81	0.04	0.82	0.01	0.45	0.01	0.71	0.13	0.46
Number of observations	144,769	144,769	144,769	144,769	144,769	144,769	144,769	144,769	144,769	144,769	144,769	144,769
Outcome pre-treatment mean	10.81	10.81	349.21	349.21	344.46	344.46	2.91	2.91	1.63	1.63	34.25	34.25
Agent FE		Yes		Yes		Yes		Yes		Yes		Yes

Notes: Table showing estimates from linear regressions. Dependent variables are shown in the column headings. *One Office Day Per Month* equals 1 for agents assigned to monthly office visits and 0 for fully remote agents. Standard errors are clustered at the agent level. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A.3: Effects on service quality outcomes

	Customer rating		Random audit rating	
	(1)	(2)	(3)	(4)
One Office Day Per Month	0.69		-0.01	
	(0.60)		(0.01)	
Post (1-5 months)	-0.67	-0.61	0.01	0.01
	(0.63)	(0.64)	(0.02)	(0.02)
Post (6-9 months)	0.55	0.76	-0.00	-0.00
	(0.98)	(0.99)	(0.02)	(0.02)
One Office Day Per Month × Post (1-5 months)	0.88	0.93	-0.03	-0.03
	(1.02)	(1.01)	(0.02)	(0.02)
One Office Day Per Month × Post (6-9 months)	-0.84	-0.76	0.01	0.01
	(1.39)	(1.40)	(0.03)	(0.03)
Adj. R-squared	-0.00	0.02	-0.00	0.02
Number of observations	4,919	4,919	4,919	4,919
Pre-treatment mean	44.24	44.24	0.17	0.17
Agent FE		✓		✓

Notes: Dependent variables appear in the column headings. *One Office Day Per Month* equals 1 for agents assigned to the monthly office-visit group and 0 for fully remote agents. Standard errors are clustered at the agent level. ***, **, and * denote significance at the 1, 5, and 10 percent levels.