Should Workplace Programs be Voluntary or Mandatory? Evidence from a Field Experiment on Mentorship

Jason Sandvik* Richard Saouma † Nathan Seegert‡

Christopher Stanton§

Abstract

There is substantial heterogeneity in whether workplace programs, like training or mentorship programs, are voluntary or mandatory. When programs are voluntary, many workers do not participate. We conducted a natural field experiment on a mentorship program in a sales call center where in one treatment arm, labeled the Mandatory-Condition, all subjects were either randomly assigned a mentor or not. A second treatment arm, the Voluntary-Condition, required subjects to opt into the program before randomization into receiving a mentor. Among those in the Mandatory-Condition, mentored workers generated 17% more revenue in their first two months of tenure, relative to non-mentored workers. In the Voluntary-Condition, treatment gains conditional on opting into the program were negligible. Comparing the conditions indicates that mentorship treatment effects are largest for workers who opt out of participating in voluntary programs. Worker characteristics are weak predictors of treatment gains and participation decisions, suggesting mandatory programs are more effective than alternative targeting schemes.

We thank Emily Beam, Jasmijn Bol, Zoe Cullen, Guido Friebel, Robert Garlick, Jessica Hoel, Mitch Hoffman, Lisa LaViers, John List, Michelle Lowry, Robert Metcalfe, Harish Sujan, Jason Snyder, Brian Waters, Michael Weisbach, and seminar participants at Harvard Business School, MIT, Indian Institute of Management Ahmedabad, the University of Arizona, the Centre for Economic Policy Research, Econometric Society, the 2021 Financial Management Association Recent Research Meeting, the 2022 Labor and Finance Annual Meeting, the Advances in Field Experiments Meeting, and the 2022 Strategy Science Conference for helpful comments.

*Eller College of Management, University of Arizona

†Eli Broad College of Business, Michigan State University

‡David Eccles School of Business, University of Utah

§Harvard Business School, NBER, and CEPR

1 Introduction

Workplace programs, such as training, continuing education, and mentorship programs are thought to improve employee productivity, foster career growth, and build human capital (Fudenberg and Rayo, 2019). Despite the importance of these programs, there are large open questions about their efficacy because of a lack of data and experimental variation. Using a nationally representative worker survey that we conducted, we document that these workplace programs are ubiquitous, that roughly half are mandatory, and that a substantial proportion (roughly 30%) of workers do not participate in programs if they are voluntary.

These stylized facts motivate three questions: (1) How effective are workplace programs for employee productivity and retention? (2) Does a program's efficacy depend on whether it is mandatory or voluntary? In other words, do program treatment effects differ for those who are more or less likely to participate in voluntary programs or by other worker characteristics? (3) How important is selection into participation for the evaluation of workplace programs?

We offer new evidence on the returns to workplace programs through a natural field experiment on a mentorship program in a U.S.-based sales firm. The workers at this firm answer incoming calls to sell digital products like television and internet subscriptions. They have high powered incentives to participate in programs that will improve their performance, as commission pay makes up a significant portion of their total compensation. For treated workers, the mentorship program began immediately after the conclusion of a two-week new hire training program, which all new hires completed. The core novelty of our experiment is that the program was administered to workers doing the same, measurable sales tasks using different treatment conditions that varied the mandatory versus voluntary nature of the programs.

The experiment entailed two levels of randomization. We first randomized cohorts of new hires into one of two groups. In the first group, labeled the Mandatory-Condition, we randomly assigned individuals to have a mentor or to not have a mentor. Although this group is labeled the Mandatory-Condition, this label does not imply perfect participation in the program or universal compliance. In the second group, labeled the Voluntary-Condition, on the first day of training the firm's training staff asked individuals to write a private message indicating whether they wanted to participate in a mentorship program that would begin at the conclusion of training. For those that did, we randomly assigned them to have a mentor or to not have a mentor. The first group simulates a mandatory program by assigning a mentor without first providing the option to opt out (even though individuals could opt out of participating after treatment assignment). The second group simulates a voluntary program by first asking whether individuals wanted to opt into the program and potentially be assigned a mentor. In both cases, mentors were randomly drawn from a pool of established, non-supervisory sales agents at the firm.

All matched mentor-protégé pairs were asked to follow a structured four-week mentorship protocol. This protocol involved a series of brief meetings where protégés were instructed to share their written responses to work-related questions with their mentors, and mentors were tasked with providing feedback on the responses before submitting the written responses to the firm's staff. For all workers, regardless of treatment, we have access to individual performance data at the daily level capturing both extensive margin effort (e.g., up-time) and outcomes (e.g., revenue generation, selling efficiency, and attrition). The experimental design and these matched data allow us to answer the motivating questions about workplace programs.

We first test whether workplace mentorship programs have positive causal effects on productivity and retention when administered without first asking workers to opt in. A recent literature review on mentoring identifies thin causal evidence on program efficacy, as many observational studies have difficulty delineating between the treatment effects of mentorship and the selection effects based on who opts into participating (Allen et al., 2017).[1] Despite this, mentorship programs are ubiquitious. Over 70% of Fortune 500 companies report that they provide their employees with mentorship opportunities (Gutner, 2009), and 45% of the respondents in our nationally representative survey have a mentorship program at work. It is possible that the these programs are widespread but with limited effects, as many firms' do not measure the causal returns to learning and development programs,[2] raising the possibility that sub-optimal practices persist because of a

[1]Given their popularity, mentorship programs have long been of interest both to academics and practitioners (Payne and Huffman, 2005; Mills and Mullins, 2008). A recent literature in economics takes identification seriously, but these studies tend to focus on a single type of program administration rather than comparing how different program characteristics may impact estimated treatment effects. See for example Lyle and Smith (2014), Porter and Serra (2020), and Ginther et al. (2020).

[2]See, for example, https://trainingindustry.com/articles/measurement-and-analytics-how-to-identify-the-right-training-kpis-for-your-learning-and-development-

lack of data or experimentation (Bloom et al., 2019).

We find that workplace mentorship programs have positive causal effects on productivity and retention among new hires in the Mandatory-Condition. Intention to treat estimates show that agents randomized into receiving mentorship generate 17% more daily revenue compared to agents allocated to the non-mentored group during their first two months on the job. Treated agents' higher productivity mostly arises from increased selling efficiency (e.g., higher revenue-per-call (RPC)), but treatment also increases schedule adherence or up-time, increasing agents' availability to take calls.

Over 80% of the productivity treatment effects persist through agents' first six months of tenure, although the estimates become noisier over longer horizons. Mentorship also improves retention. Treated agents are significantly more likely to remain with the firm in the first 30 days (where attrition rates are traditionally highest). Treatment effects on long-term retention are insignificant. The retention benefits from mentorship do not explain the productivity treatment effects, as the effects on productivity remain (i) when accounting for non-random attrition by filling in missing data with average replacement productivity after turnover and (ii) when using Lee (2009) bounds. The net present value to the firm from randomizing 114 workers into mentorship in the Mandatory-Condition was approximately $750,000 after applying the treatment gains over 6 months and accounting for the opportunity cost of mentors' time and the administrative costs of the program.

We next investigate whether the treatment effects of mentorship differ for workers who voluntarily select into the program in the Voluntary-Condition, relative to those in the Mandatory-Condition. From the earliest studies of selection based on program gains (Björklund and Moffitt, 1987), most economists are likely to believe that self-selection will result in larger estimated treatment effects among those in the Voluntary-Condition than the Mandatory-Condition. On the other hand, continuing education and mentoring programs may most benefit workers who do not seek out program opportunities, in which case the individuals who opt into voluntary programs would have the smallest treatment effects. Miscalibrated beliefs or social frictions, like stigma around asking for help when it is most needed (Edmondson and Lei, 2014; Chandrasekhar et al., 2016; Bol

programs-spon-eidesign/.

and Leiby, 2018), are both plausible channels for driving a negative correlation between program participation and treatment gains.

We find that the productivity gains from the mentorship program are substantially smaller when it is voluntary rather than mandatory, although treatment effects on retention are similar across conditions. Sales revenue and selling efficiency treatment effects for those who opt into the program in the Voluntary-Condition are approximately zero, much smaller than the unconditional treatment effects estimated in the Mandatory-Condition. There are two possible explanations for treatment effect differences between the Voluntary-Condition and the Mandatory-Condition. The first is treatment effect heterogeneity that is negatively correlated with program participation. This channel is plausible, as approximately 18% of the new hires in the Voluntary-Condition opted out before randomization into treatment. Under this interpretation, the larger treatment gains arise because workers who otherwise would not have participated in the program were included in treatment and these workers have the largest benefits. A second explanation is that framing the program as voluntary changes subjects' effort or engagement with the program. Auxiliary analyses of program compliance, meeting completion rates, and worksheet contents recorded during mentoring sessions rule out program engagement as an explanation for different treatment effects.

Using a pre-registered procedure to recover heterogeneous treatment effects, we estimate that agents who opted out of the voluntary program would have had the largest treatment gains from mentorship, whereas agents who opted into the program voluntarily had negligible treatment effects. To provide context for these estimates, we note that the larger treatment effects accrue to agents who are much less productive at baseline. A comparison of agents in the Voluntary-Condition who opt out and those that opt in but did not receive a mentor yields a difference in average daily revenue of around 30%, indicating those who opt out are much less productive at baseline. We estimate that treatment effects are sufficiently large for this group to close the initial productivity gap with their peers.

Self-selection may improve program targeting in some cases, but because treatment gains and participation are negatively correlated, the voluntary program has a much less efficient allocation of resources than the mandatory program. Over the first two months of each agents' tenure, average treatment effects for the Mandatory-Condition yielded an additional benefit to the firm averaging

$3,100 per treated agent. Reallocating the 123 treated agents in the Voluntary-Condition to the Mandatory-Condition would have thus improved efficiency by about $381,000 over the first two months of the agents' tenure.

Mandatory programs are blunt instruments, and one might wonder whether targeting could be improved to generate greater gains at lower cost. Our third question deals with this by attempting to understand first whether some types of employees predictably opt out of workplace programs and secondly whether treatment effects from the program vary by observable employee characteristics. In our context, we find that targeting mentorship based on demographics or other factors is likely to offer only modest improvements relative to randomly assigned mentorship. We designed the experiment to test for demographic and personality factors that might correlate with the participation decision. We find little evidence that demographic factors (e.g., age, gender, and marital status), prior work experience (e.g., call center experience and sales experience), or Big 5 personality scores (e.g., degrees of extroversion, agreeableness, conscientiousness, emotional stability, and openness) contribute to opting out. We do find that low hiring scores—assigned to new hires by interviewers during the hiring process—predict the opt-out decision. Overall, however, jointly considering hiring scores with demographic factors, previous work experience, and personality characteristics explains less than 7% of the variation in opt-out decisions across the entire Voluntary-Condition sample.

We also note that our experimental design cannot speak to whether the source of non-participation in the Voluntary-Condition is miscalibrated beliefs about the benefits of mentorship or some other factor, as we did not want to prompt agents to consider something unnatural before deciding to join the program. However, responses to our nationally representative survey suggest that miscalibrated beliefs along with time considerations and social frictions are likely responsible for non-participation in voluntary programs. Relatedly, none of the factors that we can access from personnel data or intake surveys predict heterogeneous treatment effects among agents in the Mandatory-Condition. Despite trying to design the experiment to test for the possibility of better targeting rules, these findings suggest that the firm we study cannot improve program targeting based on observable characteristics, at least ex-ante, although targeted follow-up interventions after observing worker productivity may be effective.

Our results have several implications for the evaluation of workplace programs. First, observational comparisons between workers who select into programs and those who do not—like a comparison between mentored workers in the Voluntary-Condition and those who opt out in the Voluntary Condition—would vastly *overstate* the effect of mentorship. Second, even if voluntary programs have randomized assignment to treatment and control, estimated treatment effects from RCTs may not generalize to universal or mandatory programs because of treatment effect heterogeneity that is correlated with voluntary program participation. While RCTs within firms can be powerful tools, our results suggest that their interpretation may depend on how employees are recruited to participate in the experiment.

We have checked a number of different factors that may qualify the results of our experiment. For example, spillovers across agents could confound our estimates. Although we worked closely with the firm's staff to reduce the possibility that non-mentored agents: (i) became discouraged after not receiving a mentor or (ii) sought out internal mentors on their own, we also conduct a variety of tests to check whether these issues are problematic. We test for spillovers by comparing non-treated new hires to hold-out groups of agents who were unaware of the mentorship treatments and to agents hired before the mentorship program began. We find no evidence of discouragement or leakage of the mentorship content/curricula to other agents, nor do we find evidence that the mentorship program crowded out organic mentorship that would have occurred otherwise.

Our findings provide several contributions to the existing literature. First, the traditional focus in personnel economics has been on incentive provision within firms and the alignment of hiring practices given an incentive scheme (Friebel et al., 2019; Oyer and Schaefer, 2011). But hiring is often noisy (Hoffman et al., 2017), and incentives may not be strong enough for all workers to invest in new skills or workplace programs. As a result, many firms face questions about how to deal with lower performers. Our evidence suggests that what might look like a hiring mismatch (the lower tail of new recruits) can be corrected with training interventions like mentoring.[3] It is possible that the wide productivity dispersion in our setting, coupled with workers who do not take advantage of beneficial resources, may be addressed by strong leadership that dictates plans for

[3] The closest related work on training and potential mismatch is likely Hoffman and Burks (2020), who study how worker overconfidence allows firms to provide training because workers are excessively eager to invest in the job. Our results instead indicate that some workers appear to under-invest in seeking out help when it is available, and targeted interventions in settings where selection into training and the treatment effects from training are negatively correlated may substantially alter the productivity distribution.

workers to improve (Lazear et al., 2015; Carter et al., 2019; Hoffman and Tadelis, 2021; Englmaier et al., 2021).

In addition, given the substantial returns to the mandatory but not the voluntary program, we offer evidence that simple tweaks to management practices can have profound consequences, bolstering the literature that has largely focused on across firm variation (Bloom and Van Reenen, 2007; Bloom et al., 2013; Syverson, 2011; Gibbons and Henderson, 2012). Other notable experiments also show that small changes can have profound effects. For example, Gosnell et al. (2020) show that practice changes among airline captains have led to tremendous fuel savings. Work by Bandiera et al. (2005) provides evidence on how social preferences interact with incentives, changing the efficacy of relative performance evaluation, while other experiments have been influential for understanding when group incentives may be effective (Friebel et al., 2017; Bandiera et al., 2013).

Finally, we provide new estimates of the scope and characteristics of workplace programs, complementing other studies of training and mentoring programs, including those in specific industries (Kram, 1988; Jones et al., 2019; Rockoff, 2008; Ginther et al., 2020; Mills and Mullins, 2008; Bruhn et al., 2018; Chatterji et al., 2019). Our findings can also be situated within the broader literature on workplace learning and peer-effects (Mas and Moretti, 2009; Bandiera et al., 2013; Herbst and Mas, 2015). In most contexts, manipulating organizations to capture boss effects is thought to be easier than organizing to capture ephemeral peer effects (Carrell et al., 2013; Lazear et al., 2015), but recent work has demonstrated that relatively simple interventions can unlock substantial peer learning (Sandvik et al., 2020).[4] We find that providing workers with the choice to opt out of mentoring reduces the efficacy of the program, and many workplace programs have features that allow employees to opt out of or into participation.

[4]There are two main differences between this study and our prior work despite the similarity of the protocol used in the mentorship program. First, we test selection into treatment here, whereas our prior paper had no ability to detect whether workers would take advantage of programs if they were voluntary. Second, the mentorship program presented a clear expectation of who would provide information (mentors) and who would receive it (protégés), whereas Sandvik et al. (2020) randomly paired employees together and treated them as equals, with no designated roles within the pairings.

2 Workplace Programs: Prevalence and Participant

Managers can enhance their employees' human capital through different workplace programs, such as training, continuing education, and mentorship programs. The efficacy of these different types of workplace programs is of substantial interest to economists (Acemoglu and Pischke, 1998), but little is known about their effectiveness due to a lack of data availability and experimental variation. Two important questions about these workplace programs are how is participation in the programs determined (e.g., mandatory or voluntary participation) and which workers participate in different programs when they are voluntary?

We conducted a nationally representative worker survey to provide answers to these questions about workplace programs. We administered the survey through the Lucid platform in June of 2022. Respondents were paid between \$1 and \$4 for taking a 7–10 minute survey. The survey began by collecting background information about the employment status of the respondents. Only those who were employed and could pass an attention check continued on to the next set of questions about their experience with workplace programs. Specifically, we asked whether their current employer offers the following programs and if the programs where required/mandatory or optional/voluntary and, if voluntary, the reasons for their participation: (i) a mentorship program, (ii) training for new hires, and (iii) ongoing training or continuing education.[5] We display the results from this survey for mentorship programs, training for new hires, and ongoing training or continuing education in Table 1.[6]

The survey responses provided three main takeaways: (1) workplace programs are ubiquitous; (2) a substantial proportion of programs are voluntary; and (3) many do not participate when the program is voluntary. Specifically, 45% of respondents said their employer offers a mentorship

[5] The survey presented respondents with the following: "Consider your current employer. Which of the following programs does your employer offer to you personally? If offered, are you required to participate (required/mandatory) or can you choose to participate or not (optional/voluntary)?" For each program, respondents chose between "Required or Mandatory," "Optional or Voluntary," or "Not offered". For the three core programs—mentorship, new hire training, and continuing education—if a respondent indicated that a program was voluntary, follow-up questions were asked about their participation and the reasons for their lack of participation, if applicable.

[6] We also included workplace wellness programs as a validation check on these answers. 65% of our respondents indicated that their workplace has a wellness program. This compares favorably to numbers cited in Jones et al. (2019) from a 2016 Kaiser Family Foundation report indicating that 53% of firms with more than 200 employees do biometric screening, 59% assess lifestyle health habits, and 83% have programs that encourage a healthy lifestyle.

program, 87% said they offer new hire training, and 80% said they offer ongoing training or continuing education. About 59% of the mentorship programs and 43% of the continuing education programs offered are voluntary. Not surprisingly, new hire training is much more likely than the other programs to be mandatory. The last column indicates that non-participation rates in voluntary programs are substantial. For example, 27% (28%) of respondents did not participate in their employer's voluntary mentorship program (ongoing training/continuing education program). Even for new hire training, rates of non-participation exceed 20% when the program is optional.

Across all voluntary programs, time, hassle, and doubts about the personal benefits of the program explain non-participation. The most common answer for why someone did not participate was the time or inconvenience of the program (47% of mentorship non-participants, 36% of new hire training non-participants, and 42% of ongoing training non-participants). The next most common answer for all programs was "Didn't believe these programs would benefit me" (26% for mentorship, 28% for new hire training, and 31% for ongoing training). Other options such as, "Didn't plan to stay at the firm, so didn't invest," "Wanted to avoid interaction with coworkers or bosses," and "Felt the program would benefit my employer more than it would benefit me," were selected by 8%–13% of non-participants.

These survey responses motivate the need to understand the administration of workplace programs. For example, whether a business should offer a voluntary versus a mandatory program depends on the expected participation decisions of employees and on the potential heterogeneous benefits of those that opt in or opt out. The remainder of the paper studies these questions in the context of a mentorship program in a sales firm.

3 Firm Setting

The mentoring program occurred in an inbound-sales call center from January to December of 2019, with data collection on protégé performance continuing after the conclusion of the mentoring relationships. The firm markets and sells the services of several companies, most of which are different television, phone, and internet providers. Sales agents answer incoming calls from potential customers and sell digital services according to the customer's needs, with a goal of selling premium services packages. Learning the sales process (e.g., how to run credit checks or determine whether callers qualify for regional promotions) and how best to up-sell premium packages can be difficult for newly hired sales agents, which makes this setting one where mentorship may help employees acclimate to a new job. All work is individual, so it is possible that mentoring or peer effects matter differently in our setting relative to those featuring team production. Throughout the year, 603 newly hired sales agents across 53 hiring cohorts entered the firm and were eligible for mentorship.

When hired, sales agents begin a two-week training program, where they learn the sales process through lectures and by listening in on other agents' calls. Agents receive training that is specific to the particular sales division (i.e., product and brand) in which they will be placed once training ends. Once agents complete their two-week training, they are allocated to a team and begin responding to calls on the sales floor. Teams are typically comprised of 10–15 individuals overseen by a direct sales manager who is responsible for monitoring performance and troubleshooting issues faced by the agents. Mentor-eligible agents were spread across eight different sales divisions.

This setting has a number of attractive features for studying the efficacy of mentorship. Most importantly, the firm provided us with granular, individual-level performance measures for each sales agent. A sales agent's productivity is independent of their coworkers' productivity, as incoming calls are routed to the appropriate division and are allocated to the next available agent within that division (i.e., calls are randomly allocated to agents), and the process is designed so that the same agent works with the caller from start to finish. The three focal productivity measures tracked by the firm are total daily revenue, revenue-per-call (RPC), and revenue-per-hour (RPH). These metrics affect each agent's commission pay, which is a function of their total revenue generation and their selling efficiency, based on RPC and RPH. An agent generates revenue through each sale they make and, at the end of the week, the total amount of revenue generated is multiplied by their commission rate. The commission rate increases in the agent's selling efficiency (RPC and RPH) relative to other agents in their same division.[7] Multiplying the agent's weekly revenue and commission rate determines the amount of commission pay they earn that week. Sales agents also earn an hourly wage, which is above the federal minimum wage and increases with tenure, and

[7] In addition, each agent has a fixed number of calls audited each week, and their commission rate decreases if conduct violations are identified by the auditors.

agents can earn occasional bonuses for doing well during temporary promotional periods.

4 Experimental Design

Cohorts of new agents join the firm on a rolling basis; some weeks two or more cohorts of new hires enter the firm, whereas in other weeks no new cohorts enter the firm. We randomly designated each cohort to either the Mandatory-Condition (probability 40%) or the Voluntary-Condition (probability 60%).[8] Of those in the Mandatory-Condition, assignment to a mentor was contingent on the supply of mentors. Assignment always occurred at random, and half of new all new hires would be given a mentor if adequate supply existed. If there were not enough mentors available (usually because of obligations to mentor other cohorts in the same brand or office), fewer than half of all agents would be assigned a mentor. Agents in the Voluntary-Condition were given the option to opt in or opt out of mentoring. Those who opted out did not receive a mentor. Of those who opted in, approximately half were randomly given a mentor and half were not (conditional on the supply of mentors).

Except for randomization into treatment and analyses of survey and productivity data, the program was administered entirely by the firm's on-boarding staff members. To participants, the mentoring program would have appeared like a normal part of the environment, rather than one imposed by outside researchers. Because of this, the experiment is very close to a natural field experiment, based on the taxonomy of Harrison and List (2004). However, participants did know

[8]Our pre-registration protocol called for the experiment to run between May 27, 2019 and December 20, 2019. The actual data we employ includes cohorts from a pilot period preceding May 27 that we had not planned to use because the mentoring protocol was slightly different (five weeks of meetings, instead of four) and because the Mandatory- versus Voluntary-Condition assignment was not originally randomized across the firm's offices. However, hiring at the firm was slower than indicated by the original projections we were given, and we were unable to extend the experiment to the Spring hiring season to makeup the shortfall because of the onset of COVID-19. When we test for characteristic balance between the pilot cohorts and cohorts arriving after the pre-registered intervention start date, we find no significant differences. While the pilot cohorts in the Mandatory-Condition and Voluntary-Condition were not randomized across offices, this does little to affect our estimates because our pre-registered strategy calls for treatment effect estimation using randomization into mentoring within each cohort (i.e., cohort fixed effects). Within-cohort randomization was no different in the pilot period and the period after pre-registration. Table A.3 shows that our main results are not significantly different between the pilot period when the Mandatory and Voluntary-Conditions were assigned by office compared to when cohorts were assigned to these conditions at random. The pre-registration text is documented in Appendix D, where we note instances in which there were minor deviations between the pre-registration and the implementation.

that outside researchers were analyzing their survey and productivity data, but they were not likely aware of their involvement in an experiment. The main concerns with participants knowing they are part of an experiment are (1) experimenter demand effects and (2) Hawthorne effects, but these are unlikely to be major concerns in our context. First, protégés were *not* told that the researchers' object of interest was the opt-out decision, and they were not told that some cohorts were given the opportunity to opt out, while others were not. Similarly, mentors were not aware that the opt-out decision was the main focus of the research, nor were they told whether their protégé was in the Mandatory-Condition or the Voluntary-Condition. As such, neither protégés nor mentors would have been able to ascertain what the experimenters demanded. Second, Hawthorne effects were not likely strong in our setting. Monitoring in this environment is ubiquitous, as the sales managers are constantly monitoring and providing feedback on sale agents' performance. So it is unlikely that agent behavior was impacted by the knowledge that outsider researchers—with whom they *never* interacted—were tracking their performance. Furthermore, the mentoring program did not change the degree of monitoring relative to the day-to-day performance tracking already in place at the firm.[9]

4.1 Identifying Mentors and Administering Allocations

The firm's internal mentoring staff announced to all incumbent sales agents that a mentoring program would be taking place, and they asked for individuals to volunteer to be mentors. In addition, the staff identified sales agents who they and the sales managers agreed would make good mentors. These agents were directly asked if they would like to participate. If the mentoring staff and sales managers felt a particular sales agent was not a good fit to be a mentor, he or she was excluded from consideration. Mentors were given two main incentives to participate. First, for each pre-scheduled, confirmed meeting they held with their protégé, they received internal currency ("kudos" dollars) worth approximately \$10. This currency could be used to purchase snacks, drinks, accessories, apparel, and other goods. Second, incumbent sales agents were told

[9]Subjects were asked to provide informed consent when responding to an intake survey. The intake survey was framed around understanding employees' preferences, work styles, and personality characteristics so that university researchers could help the firm better serve its workforce. The consent protocol did not specify that selection into or out of the mentoring program was the key metric being studied.

that effective mentoring would become a key indicator for future promotion prospects. While not formally necessary for promotion into a managerial role, being a mentor helped agents demonstrate their potential aptitude for managing a sales team down the line.

The timeline of the mentoring program for a given cohort was as follows. Each new cohort began their two-week training on the Monday of training-week one. All new hires were asked to complete a new hire survey on the first day of training, which asked agents questions about their personality traits, work styles, and previous work experiences (specifically, whether they had (a) previous experience working in a call center and/or (b) previous experience working in sales). The responses to these questions were later used to identify the characteristics of individuals who opted into versus opted out of mentoring. During training-week one, cohorts were allocated to either the Mandatory-Condition or the Voluntary-Condition, and the mentoring staff was informed of this assignment. For cohorts in the Mandatory-Condition, half (or fewer, depending on mentor availability) of the new hires were randomly allocated to receive a mentor and the remainder to not receive a mentor. There was a limited supply of mentors. If agents inquired about their lack of assignment, the mentoring staff told new hires that this was the reason only half (or fewer) of them would receive a mentor and that random allocation by a team of academics was the fairest way to distribute mentors. Those who received a mentor were randomly allocated a mentor from the pool of available mentors. To reduce any feelings of discouragement, those who did not receive a mentor were told that the company provides many other opportunities for new hires to receive help while on the sales floor (e.g., from managers, coworkers, and division leaders). The mentoring staff reported to the authors multiple times that they found no evidence of discouragement among the agents who did not receive a mentor, and we provide evidence consistent with this lack of discouragement in Section 6.1.

For cohorts in the Voluntary-Condition, the mentoring staff described the mentoring program to the newly hired agents and asked them to either opt in or opt out of the program.[10] The agents

[10] The staff members were given latitude in introducing the program, albeit they were asked to insure that the following statement was explicitly read to the new workers: "We have recently begun a mentorship program to help newly-hired sales agents when they begin working on the sales floor. Agents who opt into the program and are chosen by [the research team] will be assigned a mentor. Your mentor will approach you week 1 on the sales floor to initiate the mentoring relationship. The program will run from your first week on the sales floor to your fourth week on the sales floor, and you and your mentor will meet once a week to discuss your progress."

were told that a randomly selected subset of those who opted in would receive a mentor.[11] Agents were then asked to write down on a piece of paper whether they wanted to opt in or opt out of the mentoring program (making their decision anonymous to their peers). Agents who opted out were not given a mentor. Agents who opted in had the chance to receive a mentor. Among agents who opted in, the same procedure to allocate mentors from the Mandatory-Condition was used to assign mentors to new hires.

All of these steps (announcing the program, distributing/completing surveys, collecting opt-in/out decisions from agents, and allocating agents mentors) occurred during the two weeks of formal training. The two-week training process remained exactly the same for all agents, regardless of their mentoring group allocation. The main point of contact between the authors and the firm was the within-firm (internal) mentoring staff, with at least one dedicated employee in each of the firm's two participating locations. The authors did not interface with any of the protégés or mentors, but instead only facilitated the randomization and data collection. After the two weeks of training, sales agents began working on the sales floor. First contact between mentors and protégés was usually initiated by the mentor during the protégés' first week on the sales floor, which coincides with when we first begin tracking their productivity, but sometimes the mentoring staff explicitly facilitated initial introductions. Mentors were incentivized to make sure they met with their protégés, so they were properly motivated to make contact and to conduct scheduled meetings.

4.2 Mentor-Protégé Meeting Protocol

The mentoring relationships lasted for four weeks. Each week, the mentoring staff worked with the firm's workforce management department to build into mentors'/protégés' schedules a specific break-time during which they could meet. This reduced the concern that scheduling conflicts would prevent pairs from meeting. Mentors were told the following (full documentation of the instructions can be found in Appendix B):

[11] It is possible that a new hire's decision to opt in or out of the program may have been impacted by potential aversion to uncertainty as to whether they would receive a mentor even if they chose to opt in. We do not think this was likely a major factor in our setting, as sales agents experience a high level of uncertainty on a daily basis, suggesting that these new hires were unlikely to be strongly uncertainty averse.

> You will meet with your protégé at least once a week. Before meeting, your protégé will complete the Protégé Worksheet. If he/she has not completed it, you will kindly help him/her do so. During your meeting, you and your protégé will discuss his/her responses. You should also take this time to do the following:
>
> - Impart knowledge and skill by explaining, giving useful examples, and demonstrating processes, and asking thought-provoking questions.
> - Discuss actions you've taken to become a successful sales agent.
> - Provide him/her with any tips and sales tactics that have helped you overcome customer concerns and up-sell to better services.
> - Practice the designated sales protocol with them and help them gain a strong understanding of the products, services, and bundles available.
>
> After meeting with your protégé, you will deliver the finished worksheet to [the mentoring staff]. [The mentoring staff] will initial and timestamp the worksheet and make a record that you completed your weekly meeting responsibility.

During the meetings between mentors and protégés, the two were free to discuss whatever they wanted, but they did have to complete the worksheets in order for the mentors to receive credit for having held the meeting. Records of meeting occurrence and completed worksheets were kept by the mentoring staff and given to us. Shortly after the final week of the mentoring intervention, protégés were asked to complete a wrap-up survey that asked them questions about their mentorship experience. We use this data to gauge whether pairs continued to meet after the intervention ended and whether agents felt the program was beneficial, as discussed in Section 7.2.3.

4.3 Treatment Allocations

Figure 1 displays the allocation of cohorts and agents to the different groups and conditions of the experiment. There were 603 treatment-eligible sales agents spread across 53 hiring cohorts.[12] Twenty-

[12] An additional 56 rehired agents—those who had worked at the firm previously—were also mentor-eligible. We exclude these agents from our main analysis because their experience with the company leads to a nuanced interpretation question because of their meaningful prior experience and greater initial productivity upon

two cohorts, and their 281 sales agents, were allocated to the Mandatory-Condition, whereas the other 31 cohorts, along with their 322 sales agents, were allocated to the Voluntary-Condition. Among the agents in the Mandatory-Condition, 114 were randomized to receive mentorship and 167 were not (which we label conditions 1a and 1b, respectively). Among the agents in the Voluntary-Condition, 59 (18.3%) chose to opt out of receiving a mentor (condition 2c). Of the 263 that chose to opt in, 123 were randomized to receive mentorship and 140 were not (conditions 2a and 2b, respectively).[13]

Beyond these 53 cohorts, there were an additional 288 agents who were hired throughout the year in cohorts that were not eligible for the program in either the Mandatory-Condition or the Voluntary-Condition. These cohorts of new hires joined the firm at a time when so few mentors had availability that no agents were offered or made aware of the program. Agents in these cohorts form a hold-out group that was not aware of the mentorship program to the same extent as agents in the program-eligible hiring cohorts. Although these hold out cohorts were not randomly assigned, they have similar characteristics as program eligible cohorts because they tended to arrive on the job right after a prior cohort in the same brand and office received treatment. We later use these cohorts to test for leakage of information between mentored and non-mentored agents in program eligible cohorts.

As would be expected with successful randomization to treatments, agent characteristics are well balanced across the different treatment groups and conditions of the experiment. In Panel A of Table 2, we consider the balance across observable characteristics for agents in the Mandatory-Condition compared to those in the Voluntary-Condition (the top level of randomization). Comparing the cohort-level averages of agent age, gender, marital status, and hiring score, we find no significant differences between conditions. The average agent age in both groups is 23 years old. Women make up 44% of the agents in the Mandatory-Condition and 40% of agents in the Voluntary-

re-entry with the company compared to the average newly hired agent. Our results and conclusions do not change when we include these rehired agents into any analysis, but doing so requires us to present a large number of fully saturated parameters when we analyze the opt out decision.

[13]Limits on the supply of available mentors prevented us from achieving an exact 50/50 allocation. In addition, there were some newly hired agents who were initially randomly assigned a mentor but who, due to scheduling conflicts or their mentor leaving the firm, were not able to receive mentorship because there was never overlap in their and their mentor's schedule. These individuals are coded as not having received mentorship.

Condition; 14%–15% of agents are married in the two groups. The average hiring scores (which have a maximum value of 1) were 0.84 and 0.85, respectively. Formal tests of mean differences never reject the null of equality between the Mandatory-Condition and Voluntary-Condition.[14] Panel B performs a similar test, but in this panel we consider agents who were not eligible for mentorship because they were already working in the divisions that newly hired mentor-eligible agents would join. Non-mentor-eligible agents' average productivity levels in the divisions to which cohorts were assigned did not differ between the Mandatory-Condition and the Voluntary-Condition. The average agent who was not mentor-eligible generated $760–$794 in daily revenue, between $46.80 and $49.27 in revenue-per-call (RPC), between $115.61 and $120.21 in revenue-per-hour (RPH), while taking taking approximately 17 calls each day and spending a little over 6.5 hours at work.[15] In addition, adherence to schedule (i.e., *Adherence*, a measure of uptime and overlap of time on/off with planned breaks) does not differ between groups, nor does the conversion rate (an indicator of making any sale on a particular call). Across all productivity measures, difference-in-means tests fail to reject equality between groups.

Panel C of Table 2 considers the second level of randomization used in our experimental design, the random allocation of mentors to new hires *within* the Mandatory-Condition and Voluntary-Condition. Columns (1) and (2) show the agent-level average characteristics of agents in the Mandatory-Condition who did and did not receive a mentor, respectively. Difference-in-means tests show that these two groups of agents are similar in age, gender, marital status, and in their hiring scores. They are also similar in their referral status (whether or not they were referred to the firm by an existing employee), in their prior experience working in call centers, in their prior sales experience, and in their personality characteristics. Columns (3) and (4) and the associated p-values show that mentored and non-mentored agents who opted in among the Voluntary-Condition are also similar across all these observable characteristics, with the exception of prior call center

[14] We report similar statistics for the demographics of mentors in Table A.1, which shows that mentors of agents in Mandatory- and Voluntary-Conditions were similar in age, gender, and tenure. The mentors of agents in the Voluntary-Condition were more likely to be married than were the mentors of agents in the Mandatory-Condition. Mentors were not designated exclusively to either of the mentoring groups. In other words, a mentor's first protégé could have been from the Mandatory-Condition, whereas their second protégé could have been from the Voluntary-Condition. Mentors were never informed as to whether their protégés were in the Mandatory-Condition or Voluntary-Condition.

[15] The *Hours* measure that we use captures the amount of time agents are at work and not on unpaid breaks (i.e., total time "clocked-in").

experience (39% of non-mentored agents have prior call center experience, whereas 25% of mentored agents do). Column (5) shows that agents who opt out of mentoring in the Voluntary-Condition have worse average hiring scores and are less likely to be women. Formal tests to compare the demographics of agents across all three Voluntary-Condition cells show that all three are similar in age and marital status. Taken together, the summary statistics in Table 2 suggest that agent characteristics are, ex ante, well balanced across the different treatment groups and conditions of our experimental design.[16]

5 Results

The first question we consider is whether mentoring has a positive effect on the productivity of employees during their first months of tenure. To answer this question, we compare the performance of mentored and non-mentored agents in the Mandatory-Condition (conditions 1a and 1b, respectively), who were not given the option to opt in or opt out of the program. The comparison between conditions 1a and 1b is shown in the two leftmost bars of Figure 2, which report agents' average daily natural logarithm of revenue in their first two months on the sales floor. Mentored and non-mentored agents in the Mandatory-Condition generate average daily log revenues of roughly 6.18 and 6.03, respectively, capturing a difference in productivity of approximately 15%. We can reject that productivity is the same between conditions 1a and 1b at the 1% level, as reported in the figure note.[17] These results suggest that mentorship, when it is not made optional, has a significant positive effect on productivity.

The second question we consider is whether mentoring has a similar productivity effect when new hires can voluntarily select into or out of the program. To answer this question, we compare the performance of agents who opted in and then did or did not receive a mentor via randomization (conditions 2a and 2b, respectively). The comparison between 2a and 2b, shown in the third and fourth bars of Figure 2, shows that mentorship had no effect for agents who opted into the program. The difference between conditions 2a and 2b is small, and we fail to reject that productivity is the

[16]In Figure A.1 we display plots to illustrate the comparisons in agent characteristics across different treatment groups.

[17]We display similar bar charts in Figure A.2 to show agents' productivity during months 3–6 on the sales floor. These long-term effects are similar to the short-term, 1–2 month, effects, albeit they are less precise.

same with a p-value of 0.331. Said differently, in contrast to the large positive effect of having a mentor in the Mandatory-Condition, we find no effect for those that opt into mentorship in the Voluntary-Condition.

The difference between the Mandatory-Condition effect and the Voluntary-Condition effect suggests that those who stood to gain the *most* from mentorship were likely those who opted out. This motivates asking: what type of agents opt out? To answer this question, we compare agents who opted in and did not receive a mentor to agents who opted out (conditions 2b and 2c, respectively). The comparison of 2b and 2c, in Figure 2, shows that agents who opted in and did not receive a mentor have higher performance than agents who opted out, suggesting that opt-out agents were, on average, worse performers by at least 20%. The difference between conditions 2b and 2c is large and we can reject that productivity is the same at the 1% level. Said differently, we find a large positive selection effect on the decision to opt into mentorship.[18]

In summary, we find a large, positive productivity treatment effect of receiving a mentor in the general population (conditions 1a vs. 1b), find no effect of receiving a mentor for those that opt into the mentorship program (conditions 2a vs. 2b), and find a positive selection effect, such that agents that select into the mentorship program have higher average daily revenues even if they do not receive a mentor (conditions 2b vs. 2c). These results have profound implications for understanding the returns to mentoring programs as well as for evaluating research designs intended to estimate returns to workplace programs. Our design highlights a negative correlation between voluntary enrollment (participation) in mentorship programs and the gains from those programs. These findings suggest that observational estimates of the gains to mentoring programs are likely biased, but the direction of bias is unclear. For example, even potential designs that look at randomized mentoring among those who volunteer for a program may understate the effect of mentorship because of the first-level selection into participation. On the other hand, studies that fail to randomize and only look at those who select into mentorship might otherwise conclude that the effects of mentorship are substantial given underlying performance differences between those who opt in versus those who opt out. Our design allows us to separate treatment from

[18] We discuss the characteristics of agents who chose to opt out in Section 5.2. In general, demographic information, hiring scores, previous work experience, and personality traits explain very little of the variation in the opt-out decision.

selection, showing heterogeneous treatment effects that are largest for those who are least likely to participate. Said another way, agents who chose to opt out likely would have meaningfully benefited from mentorship. We defer discussion of why these agents opted out in the face of large potential gains to Section 5.2. We now formalize these insights with statistical tests and analyze additional measures of performance that provide context for these gains.

5.1 Treatment and Selection Effects of Mentoring

In this section, we estimate treatment and selection effects on revenue together in a single model and show that they extend to a host of other measures of agent productivity. To do this, we use a sample of agent-day productivity data for all mentor-eligible agents in their first two months on the sales floor and estimate the following model using ordinary least squares:

$$\mathrm{y}_{i,t} = \alpha + \beta_1 \mathrm{Mentored}_i + \beta_2 \mathrm{Mentored}_i \times \mathrm{Voluntary}_i + \beta_3 \mathrm{Voluntary\ Opt\text{-}Out}_i + \gamma_j + \varepsilon_{i,t}, \tag{1}$$

where $Mentored_i$ equals one if agent i was randomly assigned to receive a mentor, and zero otherwise, $Voluntary_i$ equals one for agents in the Voluntary-Condition, and zero otherwise, and $Voluntary\ Opt\text{-}Out_i$ equals one for agents who chose to opt-out of possibly receiving a mentor, and zero otherwise. t denotes the date of the observed productivity. All models include cohort fixed effects, γ_j, which control for differences in productivity that are specific to cohorts, and an idiosyncratic error term, $\varepsilon_{i,t}$.[19] The productivity outcome variable, $y_{i,t}$, differs by specification and is one of the following: ln(Revenue), ln(RPC), ln(RPH), or Adherence, where RPC is revenue-per-call, RPH is revenue-per-hour, and Adherence captures how closely agents adhere to their pre-set schedule (e.g., taking breaks and eating lunch at the correct time).[20]

The results of these estimations are reported in Columns (1)–(4) of Table 3. The point estimates

[19] All of our pre-registered specifications include cohort fixed effects, as we expected ex-ante that between cohort variation would significantly reduce detectable effects because cohorts face different demand conditions based on the products they sell and the timing of their hiring. Note that with cohort fixed effects, calendar time and elapsed time since hire are co-linear, so cohort fixed effects absorb time fixed effects when looking at fixed windows of time after the agent's hiring date. We show in Section 6.4 that our results are robust to the inclusion of date fixed effects. In addition, we show that our results are robust when we exclude cohort fixed effects.

[20] We use the natural log of one plus Revenue, RPC, or RPH. Adherence is bounded by 0 and 1.

on *Mentored* capture the treatment effect of mentorship among agents in the Mandatory-Condition. Column (1) shows evidence of a positive and significant effect of mentorship on ln(Revenue). The point estimate suggests that mentored agents generate 17% more in daily revenue than do their non-mentored peers.[21] In Columns (2)–(4) of Table 3, we show that the positive treatment effect for those in the Mandatory-Condition is apparent across the other performance metrics captured by the firm. The significant effects of *Mentored* on both selling efficiency measures—ln(RPC) and ln(RPH)—and on the main measure of labor supply—Adherence—suggest that mentored agents are both allocating more time to revenue-generating tasks and are generating more revenue per unit of time. Table A.2 reports the long-term treatment effects of mentorship, at months 3–6. While the point estimates on *Mentored* are statistically insignificant, they are at least 80% of the magnitude of the 1–2 month effects, suggesting that mentorship helps newly hired agents get up-to-speed more quickly, while also possibly leading to a shift in performance that persists over time.

The negative point estimates on *Mentored* × *Voluntary* in Columns (1)–(4) of Table 3 indicate that mentorship had a much smaller effect on the productivity of agents in the Voluntary-Condition compared to those in the Mandatory-Condition. The penultimate row of Table 3 reports the p-values from post-estimation tests of the null that the sum of the coefficients on *Mentored* and *Mentored* × *Voluntary* equal zero. The large p-values in Columns (1)–(4) suggest that the effect of mentorship among agents in Voluntary-Condition was negligible. This finding is consistent with the results in Figure **??**, which show that the mentored agents after opting in do not generate more revenue than do non-mentored opt-in agents in their first months on the sales floor. The point estimates on *Voluntary Opt-Out* captures the selection effect, comparing the performance of those who opt out to those who opt it but do not receive a mentor. The negative, significant estimates in Columns (1)–(3) indicate that agents in the Voluntary-Condition who chose to opt out of the mentorship program performed significantly worse than did non-mentored agents who had signaled their interest in the program.[22][23]

[21]When we use level revenue as the dependent variable, we find that mentorship increases daily revenue generation by over $55.

[22]We follow Anderson (2008) and report sharpened q-values, which is like a p-value after adjusting for the False Discovery Rate (FDR), in brackets in Table 3. These q-values indicate that the statistical significance of our point estimates are robust to accounting for multiple hypothesis testing concerns. The sharpened q-value is conservative in our case because it does not account for the positive correlation across tests.

[23]We report the effect of mentorship on the number of calls answered, the number of hours worked, and

We estimate the impact of the mentoring program on agent retention in Columns (5) and (6) of Table 3. Call centers have notoriously high levels of attrition (Hoffman et al., 2017). Executives at the study firm talk about the importance of new hires hitting multiple important milestones with tenure. The most important is the initial transition from training to the sales floor, where new hires begin working independently.

To estimate retention effects, we use data with a single observation per unique hired agent, and we set the dependent variable to be $Tenure_t$, which equals one for agents who achieve at least t months of tenure at the firm, and zero otherwise. The point estimate on *Mentored* in Column (5) indicates that agents in the Mandatory-Condition who were mentored were 10.2 percentage points more likely to achieve one month of tenure, relative to non-mentored agents. The baseline retention rate for these non-mentored agents is reported in the bottom row of the table, and implies that 78.4% of non-mentored agents make it to the one-month benchmark. The retention effect in month two decreases to 8.7 percentage points relative to a baseline retention rate of 54.5%. The point estimate on *Mentored* $\times$ *Voluntary* in Column (5) is small and insignificant, and the p-value on the sum of the coefficients on *Mentored* and *Mentored* $\times$ *Voluntary* is less than 0.05, suggesting that mentored agents in the Voluntary-Condition were also more likely to remain with the firm than their non-mentored counterparts.

These retention effects are visibly present in Figure 3. This figure displays Kaplan-Meier survival rate estimates, which take agents' hire dates as a starting point and then display the fraction of agents who remain at the firm at various points in event time, relative to their hire dates. This measure allows us to depict retention over time for mentored and non-mentored agents separately, illustrating if and where their retention patterns diverge. In Figure 3a, we see that over 20% of non-mentored agents in the Mandatory-Condition exit the firm within one month of their hire date. Among mentored agents, however, only about 8% depart in their first month. This gap in retention persists in month two, and then the gap shrinks in months three and four. The survival rates for both groups of agents then follow a similar pattern, with a flattening of the curve in months five and six at around 23%–29% retention. The survival rates of agents in the Voluntary-Condition are depicted in Figure 3b. This figure shows that about 17% of non-mentored opt-in agents leave the

the number of calls answered per hour worked in Table A.4.

firm in the first month, whereas less than 3% of mentored opt-in agents do so. By the end of month three, however, the overall retention rates between mentored and non-mentored opt-in agents only differ by 4 percentage points.

We defer a comprehensive discussion of robustness to Section 6, but it is important to note that these results do not arise from the discouragement of agents who are randomized out of receiving a mentor. Nor does the program appear to crowd out organic mentoring that may have occurred in the absence of a program. In addition, the estimated effects on productivity are not simply driven by differential retention between mentored and non-mentored agents or by different compliance with the program.

5.2 How do Opt-out Agents Differ from Opt-In Agents?

The differential treatment effects between the Mandatory-Condition and the Voluntary-Condition suggest that opt-out agents stood to gain from mentorship. Before we formally estimate these potential gains, we consider how this subset of agents differs from those who chose to opt into the mentorship program. We restrict our sample to only include the 322 agents in the Voluntary-Condition, and then we estimate the following logistic model:

$$\text{Opt-Out}_i = \alpha + \beta_1 \text{Age}_i + \beta_2 \text{Female}_i + \beta_3 \text{Married}_i + \varepsilon_i, \tag{2}$$

where *Opt-Out* equals one if the agent opts out of receiving a mentor, and zero otherwise. *Age* captures the age of the agent at the time they were hired. *Female* equals one if the agent is a woman, and zero otherwise. *Married* equals one if the agent is married, and zero otherwise, and ε_i is an idiosyncratic error term. We display the result of this estimation in Column (1) of Table 4, which indicates that neither age, gender, nor marital status are significant predictors of the opt-out decision.

In Column (2) we include *Hiring Score* into the logistic regression, which is the score given to the agent by the recruiter who interviewed them for the job. We are missing hiring score data for 25 agents, so we set their hiring scores equal to zero and include a dummy variable into the model that indicates that they had missing data. The results in Column (2) show that *Hiring Score* is a significant predictor of *not* opting out. Said another way, agents with low hiring scores are significantly more likely to opt out of mentorship when given the choice.

In Column (3), we control for several other observable characteristics that may influence an agent's opt-out decision. Specifically, we control for their location, whether they were referred by an existing agent, and their personality traits (which we gathered via the new hire survey). None of these additional factors have significant predictive power for the decision to opt out. In Column (4), we incorporate two final variables into the model: *Call Center Exp.*, which equals one if the agent indicated in the hiring survey that they had previous experience working in a call center, and zero otherwise, and *Sales Experience*, which equals one if the agent indicated in the hiring survey that they had previous experience working in sales, and zero otherwise. The results in Column (4) show that neither *Call Center Exp.* nor *Sales Experience* have significant explanatory power for opting out.

All of these observable characteristics, together, explain less than 7% of the variation in the opt-out decision. This suggests that unobservable characteristics largely drive the decision to forgo mentoring. We also know that opt-out agents perform worse on average than agents that opt in, and this too could be driven by unobservable characteristics.

5.3 Would Opt-Out Agents Have Benefited from Mentorship?

We can use the estimated Mandatory-Condition and Voluntary-Condition productivity treatment effects, along with the data on the fraction of Voluntary-Condition agents who opt out of receiving a mentor, to estimate the treatment effect of mentorship among agents who opted out. We pre-registered the following procedure for this purpose. Using output measure Y, we define the average treatment effect of mentoring given selection as:

$$\underbrace{\beta_{OptInMentored}}_{ATE|Selection} = E(Y_{OptInMentored}) - E(Y_{OptIn\sim Mentored}).$$

We can then write the unconditional average treatment effect of mentorship as the weighted average of heterogeneous effects with shares π:

$$\underbrace{E(Y_{MandatoryMentored}) - E(Y_{Mandatory\sim Mentored})}_{ATE} = \beta_{OptInMentored} \times \pi_{OptIn} \; + \beta_{OptOutMentored} \times \pi_{OptOut}.$$

Rearranging terms, we get,

$$\beta_{OptOutMentored} = \{ATE - \beta_{OptInMentored} \times \pi_{OptIn}\}/\pi_{OptOut}.$$

We use the estimated treatment effect in the Mandatory-Condition as the estimated ATE, and we use the estimated treatment effect in the Voluntary-Condition as the estimated $ATE|Selection$.[24] The values of π come from the proportion of agents who opted out versus opted into mentoring among the Voluntary-Condition. We show the estimated treatment effect for opt-out agents in Table 5, where standard errors come from 500 block-bootstrap iterations by cohort. The point estimate of 1.212 in Column (1) of Panel A implies that opt-out agents would have been able to more than double their overall revenue generation, on average, had they received mentorship. It appears likely, therefore, that opt-out agents in this setting were actually those who would have benefited the most from receiving additional help from a mentor.

We also implement a GMM estimator of heterogeneous effects as an alternative to making cross-condition comparisons. This analysis was not pre-registered, but is motivated by the desire to account for observables in light of the large effect sizes in the pre-registered approach. Specifically, the GMM approach builds in additional flexibility for estimating the opt-out probability, which may have differences across cohorts due to sampling variation. In particular, we form moment conditions that correspond to two equations. The first equation accounts for the opt-in probability among agents in the Voluntary-Condition based on the moment conditions,

$$E[x_i \times (OptIn_i - \exp(x_i\delta)/(1 + \exp(x_i\delta))]. \tag{3}$$

Denoting $\pi_{OptIn}(x_i, \delta) = \exp(x_i\delta)/(1 + \exp(x_i\delta))$, the second equation moment conditions come

[24] We include cohort and day-of-week fixed effects in estimating these treatment effects.

from

$$
\begin{aligned}
E[z_i \times (Y - Cohort_i - \beta_{OptIn}(Mandatory \times \pi_{OptIn}(x_i, \delta) + VoluntaryOptIn) \\
-\beta_{OptInMentored} Mentored(Mandatory \times \pi_{OptIn}(x_i, \delta) + VoluntaryOptIn) \\
-\beta_{OptOutMentored} Mandatory \times (1 - \pi(x_i, \delta)))].
\end{aligned} \tag{4}
$$

This is a 2-equation system GMM estimator that recovers an interaction between a latent opt-out probability (estimated in equation (3) and treatment effects in a equation (4). The procedure iteratively guesses the parameters of the opt out probability at baseline and as a function of agent characteristics for the Voluntary-Condition and applies this functional for the latent probability in the Mandatory-Condition. This enables some regularization of the estimates based on allowing the opt-out probability to vary with agent and cohort characteristics.

The instruments z_i for the second equation are indicators for *Mandatory* and *Mentored*, *Voluntary Opt-In*, and *Voluntary Opt-In* and *Mentored*, along with cohort dummies (in practice we use the within transformation). We estimate the system jointly and allow x_i to include Age, Hiring Score, Female, and a time trend that captures when the cohort was hired and onboarded.

We continue to find substantial heterogeneous treatment effect estimates for non-mentored agents in the Mandatory-Condition who were deemed likely to opt out if provided the choice. Using this approach, the estimated treatment effect on log revenue for those who would opt out is 0.794 with a standard error (clustered by cohort) of 0.359, as reported in Panel B of Table 5. If we exclude cohort fixed effects, as we do in Panel C, the treatment effect is 0.531 and the standard error is 0.290.

Also reported in the table are the baseline effects for agents who opt in compared to those that opt out and the effect of mentorship for those who opt in. Across Panels B and C, agents who opt into the program but are who are not randomized to a mentor are found to generate 33 to 44 percent more daily revenue than agents who opt out.

6 Robustness

In this section we rule out several alternative explanations for our results, and we discuss a host of robustness tests. First, we show that our results are not driven by discouragement or leakage. We then show that the mentoring program does not appear to have crowded out organic mentorship that would have occurred in the program's absence. Finally, we test for differences in framing effects that may have changed compliance and we show that the effects on productivity are not simply driven by non-random attrition.

6.1 Effects Are Not Driven By Discouragement or Leakage

We first consider the possibility that discouragement or leakage may have affected our estimates of the treatment effects of mentorship. After a number of exercises meant to detect these issues, we find no evidence that these potential concerns alter our findings. However, to understand these tests, it is useful to detail how these different channels may have interfered with our estimates. First, agents who did not receive a mentor may have become discouraged, reducing their performance as a result. Discouragement would cause our estimates to overstate the actual benefits of receiving a mentor because it would negatively affect non-mentored agents, who form the control group for our estimates. Second, agents that did not receive a mentor via our random allocation process may have sought out their own mentor, leading to treatment leakage. A different source of leakage may be where non-mentored agents query mentored agents about the information received from mentoring. Any treatment leakage would increase the performance of non-mentored agents. Although the staff implementing the program reported no evidence of leakage in either the Mandatory-Condition and Voluntary-Condition, it is possible that agents who opted into the program in the Voluntary-Condition would be most likely to seek out help, as the framing of this condition may have made agents more aware that they should seek out program resources.

We implemented the mentoring program in a way that was meant to limit discouragement and leakage. First, we worked with the company to reduce the chance that non-mentored agents in the Mandatory-Condition received information about the mentorship program. Specifically, we asked the internal mentoring staff to privately notify new hires who would receive a mentor about their involvement in the program—dampening the potential for discouragement among non-mentored agents. Second, the firm's internal mentoring staff told agents in the Voluntary-Condition, all of whom were made aware of the program, that ample opportunities for receiving help were available to those who did not end up receiving a mentor—again, reducing discouragement and the desire to independently seek out a mentor. We asked the staff to monitor potential discouragement and leakage throughout the study, including any complaints or concerns over not being matched to a mentor, albeit no such feedback was reported back to us.

We test the net effect of discouragement and leakage by comparing the performance of three groups of agents: (1) new hires who were in mentor-eligible hiring cohorts; (2) new hires who were not in mentor-eligible hiring cohorts during the time of the experiment; and (3) seasoned veterans who began working at the firm before the onset of the mentorship program (whose tenure exceeds 18 months). Unlike the experimentally identified results within cohorts and across the Mandatory- and Voluntary-Condition treatment cells, this exercise requires a comparison to agents who were not part of the experiment. Our approach is to compare the productivity of new hires relative to veteran agents in these different conditions. Under the null of no discouragement or leakage, we would expect that new hires who are not mentored but who are in eligible cohorts should have indistinguishable productivity differences (relative to veterans) compared to new hires entering the firm in non-mentor eligible cohorts or prior to the existence of the mentoring program. Because these results are not experimentally identified, we use regression adjustments to make conditions comparable between new hires and veterans across the different treatment or hold-out groups.[25] We estimate the following model using ordinary least squares:

$$\begin{aligned}\ln(\text{Revenue})_{i,t} = \alpha &+ \beta_1 \text{New Hire}_i + \beta_2 (\text{New Hire} \times \text{Mandatory})_i \\ &+ \beta_3 (\text{New Hire} \times \text{Voluntary})_i + \beta_4 (\text{Mentored} \times \text{Mandatory})_i \\ &+ \beta_5 (\text{Mentored} \times \text{Voluntary})_i + \zeta_{j,l,t} + \gamma_{n,l} + \varepsilon_{i,t},\end{aligned} \tag{5}$$

where *New Hire* equals one if the agent has tenure of two months or less and zero for seasoned

[25]Part of this adjustment is recognizing that divisions in the firm have different levels of revenue per agent, and our approach requires a comparison of relative revenue differences between new hires and veteran agents. As a result, we restrict to divisions with 5 or more mentor-eligible agents and 5 or more new hires that were not mentor-eligible.

veterans, *Mandatory* equals one for agents in the Mandatory-Condition and zero otherwise, *Voluntary* equals one for agents in the Voluntary-Condition and zero otherwise, *Mentored* equals one for agents who receive mentorship and zero otherwise. Divisions may experience their own idiosyncratic shocks that can be partially smoothed with different staffing choices that may vary across office, so we include division-by-location-by-date fixed effects, through $\zeta_{j,l,t}$, to provide flexibility. To capture potential differences in new hires across offices for all cohorts, $\gamma_{n,l}$ removes a location-by-new-hire fixed effect. Finally, $\varepsilon_{i,t}$ is an idiosyncratic error term.

The results of this estimation are displayed in Panel A of Table 6. This table only contains data from the period during the experiment, and is designed to test for discouragement and leakage. The negative and statistically significant coefficient on *New Hire* in Column (1) suggests that newly hired agents generate approximately 42% less in daily revenue, relative to seasoned veterans. The small and insignificant coefficients on *New Hire* × *Mandatory* and *New Hire* × *Voluntary* suggest that newly hired non-mentored agents in mentor-eligible cohorts perform similarly to newly hired agents in non-mentor-eligible cohorts. This result suggests neither discouragement nor leakage are likely an issue in our setting. In addition, we fail to reject that these estimates are jointly equal to zero and that they are equal to each other, suggesting that newly hired, non-mentored agents in the Mandatory-Condition performed similarly to newly hired, non-mentored agents in the Voluntary-Condition. The positive and statistically significant effect on *Mentored* × *Mandatory* and the insignificant effect on *Mentored* × *Voluntary* align with our main mentoring treatment effects discussed in Section 5.1. Column (2) shows similar results when we use an alternative combination of fixed effects that add additional flexibility to capture the possibility that newly hired agent performance relative to veteran agent performance may vary throughout the year (through the inclusion of New Hire-by-Date fixed effects) or by division and office. Column (3) shows that our results are robust when controlling for agent characteristics, which is important given that randomization of agents into treatments did not occur for seasoned veterans and non-mentor-eligible agents.

Columns (4)–(6) repeat this exercise while using the natural log of RPC as the dependent variable. The negative point estimate on *New Hire* in Column (4) shows that seasoned veterans significantly outperform less tenured agents on this productivity metric as well. The small, insignificant coefficients on *New Hire* $\times$ *Mandatory* and *New Hire* $\times$ *Voluntary* in most of the columns further support the notion that discouragement and leakage are not likely driving our estimated mentorship treatment effects.

6.2 Did the Mentoring Program Crowd Out Organic Mentoring?

It is possible that our mentoring program may have crowded out organic mentoring that would have occurred in the program's absence. Specifically, if our program absorbed all of the potential mentors, then non-mentored agents would perform less well because they did not have the opportunity to naturally find a mentor. As a result, we would expect that non-mentored agents in treatment-eligible cohorts would perform less well than new agents at other times.

To test the possibility that the mentorship program crowded out organic mentoring that would have occurred in the program's absence, we compare the productivity of mentor-eligible new hires to the productivity of new hires who joined the firm prior to the onset of the mentorship program. We continue to use the performance of veterans as a basis for comparison. If the mentorship program in our study crowded out organic mentorship opportunities for non-mentored agents—for example, by occupying all potential mentors with formal mentorship duties—then we would expect non-mentored new hires who were mentor-eligible to perform worse than new hires from prior years. We find no such evidence. Panel B of Table 6 shows the results of estimations that are similar to those in Panel A, albeit the comparison group is new hires from before the mentorship program, rather than those from contemporaneous hiring cohorts (contemporaneous cohorts are not a good comparison group because they would be subject to the same limited supply of mentors). The positive point estimates on *New Hire* $\times$ *Mandatory* and *New Hire* $\times$ *Voluntary* in Columns (1)–(3) suggest that non-mentored, mentor-eligible new hires are no less productive than new hires from previous hiring seasons relative to veterans (in fact point estimates are positive), suggesting that crowd-out is not an issue in our setting. The positive coefficients become smaller and statistically insignificant in Column (3) when we adjust for agent characteristics, suggesting there may be composition differences in new hires over time. It is not surprising to find no evidence of crowd-out, as the company at baseline had relatively little organic peer-to-peer mentoring (see (Sandvik et al., 2020)). This further reduces the concern that our mentorship program crowded out mentoring that would have occurred naturally.

6.3 Robustness to Attrition and Lee Bounds Estimates

The results in Section 5 provide evidence that mentorship significantly increases productivity and early stage retention. The higher productivity due to mentoring captures benefits both from the intensive margin (within-agent increases from agents learning more) and the extensive margin (due to lower/differential attrition). To understand the impact of retention on our estimates of productivity, we estimate the treatment effects of mentorship after creating something akin to a balanced panel by agent slot. What this entails is filling in the productivity of agents who leave the firm with the expected productivity of a replacement. The total productivity gain to the firm for the slot is then the productivity of the treated agent over the days until a replacement arrives. Similarly, for non-treated agents who turnover, we consider their replacements. That is, for all mentor-eligible agents who leave the firm before the two-month mark, we extend the time series of their productivity through two months and replace their post-termination productivity values with the average productivity that would be expected of a newly hired replacement agent. We do this by computing the average productivity of newly hired non-eligible agents in the same location-division-year-quarter as the departed agent. We then re-estimate our main regression models using Equation (1). The results are displayed in Table 7, and they are largely consistent with our main results. Specifically, the positive, significant point estimates on *Mentored* indicate that mentorship improves agent productivity, and the negative, significant coefficients on *Voluntary Opt-Out* imply that the weakest performers are those who opt out of receiving help.

As an additional approach to delineating between the extensive and intensive margins, we estimate treatment effect bounds that account for non-random attrition as proposed by Lee (2009). The key assumption when implementing this approach in our setting is that some mentored agents would have left the firm in the absence of mentorship, but no mentored agents left the firm *because* they were mentored (a traditional monotonicity assumption). Table A.5 reports upper and lower bounds of the estimated treatment effect of Mandatory-Mentoring on productivity in months 1–2 and months 3–6 in Panels A and B, respectively. In Panel A, the lower bound in Column (1) of 0.084 is about half the size of the main effect in Table 3, and the upper bound is over double the main effect. Both upper and lowers bounds are statistically significant, suggesting our estimated treatment effects are largely attributable to the intensive margin of agents becoming more productive.

6.4 Further Robustness to Controls and Design Decisions

We consider several additional specifications that highlight the robustness of our main findings. We report these specifications in the coefficient plots in Figure 4, where Figure 4a considers the treatment effect of mentorship on log revenue and Figure 4b considers the effect on log RPC. The baseline coefficients and 95% confidence intervals from Table 3 are labeled "Baseline." We report the coefficients on *Mentored*, *Mentored* + *Mentored* × *Voluntary*, and *Opt-Out*.

First, we consider specifications that do not include cohort fixed effects (second line). In the third specification, we report estimates where we include date fixed effects, to capture variation in productivity that is idiosyncratic to a particular day on the sale floor. The fourth estimation includes controls for the agent's demographic characteristics: age, gender, marital status, and referral status. The fifth estimation includes additional controls for the agent's hiring score, previous call center experience, and previous sales experience. The sixth estimation removes observations in which agents are no longer working in the division in which they were initially hired. The point estimates and confidence intervals are quite stable across these different specifications, highlighting the robustness of our results. One exception is a muted selection effect when cohort fixed effects are omitted from the regressions.

7 Additional Analysis

Next we consider several extensions of our main analysis to better understand the possible mechanisms at play. First, we discuss heterogeneity in the observed treatment effect of mentorship across observable characteristics. We consider meeting completion rates, worksheet content, and post-mentorship survey data to investigate additional potential mechanisms. To end, we consider the net present value to the firm of the mentorship program, and we address the external validity of our results.

7.1 Observable Characteristics and Heterogeneous Productivity Treatment Effects

An important question is whether the factors that determine the opt-out decision also drive gains to mentorship, which may explain why opt-out agents stand to benefit from being mentored. In this section, we consider whether there are heterogeneous mentorship treatment effects based on observable sales agent characteristics. When we estimate an adaptive lasso regression that penalizes over-fitting, we find that the only significant heterogeneous effects of mentorship are based on age and marital status. The penalized coefficient on mentored from the lasso is 0.16, and the coefficients on the interaction of mentored with above-median age and married are 0.16 and -0.12, respectively. Note that there are only 5 married agents who are below median age, so the negative coefficient largely offsets the positive effect of age for about 23 percent of the older sample. This heterogeneity would not have have been known ex-ante without the experiment, and ex-post it appears treatment effect heterogeneity by observable characteristics is less important than whether the program allowed self-selection in participation.

7.2 Evidence on Potential Mechanisms

Here we explore additional possible differences between mentored agents in the Mandatory- versus Voluntary-Condition that may have contributed to the observed differences in mentorship treatment effects. We specifically consider differences in mentor meeting completion rates and worksheet content. We then briefly discuss anecdotal evidence from the wrap-up surveys.

7.2.1 Meeting Completion Rates

The first potential mechanism that we explore is the amount of time protégés spent interacting with their mentors. To capture this, we assess the mentorship meeting completion rates of mentor-protégé pairs. We tabulate these in Table A.6. Of the 114 agents assigned to mentorship in Mandatory-Condition, four never completed a recorded meeting with a mentor, whereas nine of the 123 treated agents in Voluntary-Condition never met with their mentor. Mandatory-Condition protégés completed both a greater number of their scheduled meetings (2.58 vs. 2.27) and had a higher meeting completion ratio (83% vs. 72%). This difference potentially suggests that the opt-in/opt-out protocol of the Voluntary-Condition caused agents to see meeting completion as more optional than mandatory. It may be the case that program framing effects were partially at play in our setting (Hossain and List, 2012; Hong et al., 2015; Englmaier et al., 2017). These framing effects are unlikely to explain the entire difference in treatment effects between Mandatory-Condition and Voluntary-Condition, however, as our main results in Table 3 are robust when controlling for agents' meeting completion rates. Furthermore, the 72% completion rate in Voluntary-Condition means that the vast majority of scheduled meetings still took place.

As an additional way of considering the role of meeting completion rates on the efficacy of mentorship, we perform an instrumental variables estimation. We use assignment to a mentor, *Mentored*, as an instrument for mentor-protégé meeting completion. We report the results in Table A.7, where Columns (1)–(4) consider only Mandatory-Condition agents, and Columns (5)–(8) consider only Voluntary-Condition agents who did not choose to opt out of the mentorship program. The dependent variable in Columns (1) and (5) is *Meeting Completion Ratio*, the fraction of possible (i.e., scheduled) mentor-protégé meetings that the protégé completed. The dependent variable in Columns (3) and (7) is *Number Recorded Meetings*, the number of mentor-protégé meetings that the protégé completed. Then in Columns (2), (4), (6), and (8), we regress log revenue on the predicted values of these meeting completion numbers. The results in the mandatory condition in suggest completing all meetings yielded a treatment effect among the compliers of 0.2. Agents in the Mandatory-Condition who complete only half of their meetings would then have a treatment gain of about 10%. The same estimate for those in the Voluntary-Condition indicates that full meeting completion would yield a treatment effect of -7%, which is not due to noise given the substantial first-stage coefficient on being assigned to treatment. That is, these estimates indicate that the effect of a meeting is different conditional on the Mandatory or Voluntary Condition among compliers. Because treatment effects differ for compliers, overall intention to treat estimate differences are not driven by differences in the compliance rate.

7.2.2 Worksheet Content

The second test we undertake to understand potential mechanisms is to examine the worksheet content of mentored agents. We do this two different ways. First, we consider the amount of content transcribed on each agent's worksheets by counting the total number of words written. While this is not a perfect measure of the quality of the mentor-protégé meetings, it allows us to proxy for the level of engagement the agents had towards the mentorship protocol. In our second approach, which is motivated by the worksheet analysis in Sandvik et al. (2020), we use a bag-of-words to determine how much of a response's content is focused on job-specific skills and knowledge and how much is focused on receiving support or encouragement.[26] Specifically, we count up the number of "skill" words an agent uses in their responses, and we do the same thing for the number of "support" words. Words that do not get classified as either skill or support words are categorized as "other," including stop words.

We then compare the worksheet content of Mandatory-Condition agents and Voluntary-Condition agents. We have completed worksheet data for 159 out of the 224 mentored agents, as some worksheets that were turned in to the internal mentoring staff were never returned to us. For each agent, we compute the number of words written on all of their completed worksheets, and we divide this by the number of worksheets received. We do the same thing to create variables for the number of skill words per worksheet, support words per worksheet, and other words per worksheet. We then regress these agent-level word count per worksheet variables on the indicator *Mandatory-Condition*, which equals one for agents in Mandatory-Condition and zero for agents in Voluntary-Condition. Table A.8 reports the results. We do not find meaningful differences in the number of total words, skill words, or other words recorded by agents in the two mentorship groups. We find some evidence that Mandatory-Condition agents use more support words than do Voluntary-Condition agents. The point estimate represents a 30% greater use of support words, but the effect is only significant at the 10% level. Taken together, differences in worksheet content do not appear large enough to explain the heterogeneous treatment effects of mentorship. As a result, we conclude that the most likely reason for treatment effect heterogeneity is that different agents benefited from similar program features, rather than the possibility that program features or characteristics differed by

[26] We list the words in each category in Appendix C, along with multiple examples of responses.

agent. That is, agents who were most likely to opt out of mentorship appeared to benefit more from the same types of mentorship that agents with a low opt-out propensity received.

7.2.3 Wrap-Up Survey Responses

As a final way to understand how mentorship drove the observed productivity benefits, we consider evidence from survey responses. Two weeks after mentors and protégés completed their final mentorship meeting, the protégés were asked by the internal mentoring staff to complete a post-mentorship wrap-up survey. This survey asked them questions about how beneficial they felt their mentoring relationship was and whether they had continued to have contact with their mentor after the meetings ended. The completion rates for this survey were quite low (less than 10%), as the firm did not provide any explicit incentives for completing it. As such, we treat these responses as anecdotal evidence of the feelings that some had towards the mentorship program.

We display the protégés' responses to ten different questions in Figure 5, all of which were presented as statements to which respondents were asked to indicate their level of agreement. The responses indicate that protégés, on average, felt like they and their mentors both benefited from the mentoring relationship. We do not find evidence that the relationships between mentors and protégés extended beyond the workplace, nor do we find responses that suggest the mentoring relationship distracted agents from reaching their potential. Most importantly, the responses indicated that mentors and protégés continued to interact after the formal relationship ended. Specifically, the average respondent said they continued to seek out help/advice from their mentor and that their mentor continued to teach them skills after the four-week protocol ended. This suggests that one benefit of the mentoring program was that it provided new hires with an additional resource—their mentor—to receive help in the future. The average respondent also said that mentorship helped them incorporate important selling tactics into their sales process and that mentorship increased their day-to-day satisfaction at work. This suggests that protégés likely benefited from both an enhanced knowledge of the sales process and a greater level of social support in the office.

7.3 Net Present Value of Mandatory Mentorship and Misallocation

Next we estimate the returns to mentorship for the firm by considering productivity out to the six month horizon, net of the administrative and opportunity costs associated with the program. The average agent, among the 114 mentored agents in the Mandatory-Condition, experienced over a 16% increase in daily revenue, which results in \$2,100–\$2,800 more in revenue per agent-month, depending on the tenure of the agent. The firm earns this additional revenue net of an 8% commission rate that is paid to sales agents. We multiply these monthly net-revenue amounts by the number of mentored sales agents still present in the firm each month, over these six months.[27] We discount future cash flows using a 12.5% discount rate, which gives us a present value of the additional revenue earned by mentored agents equal to \$861,664.

We then subtract the estimated time costs of taking the mentors and protégés off the phone and administrative costs to calculate the net present value of the mentorship program. Mentors and protégés spent 30 minutes in the mentorship meetings each week. The average revenue-per-hour for mentors and protégés are \$146.25 and \$92.77, respectively. Mentors were also paid an additional \$10 of "kudos" points for completing each meeting. Together this implies a cost of \$129.51 per meeting. We include the administrative costs of the two internal mentorship staff members who oversaw the program in the two locations, estimated to be approximately \$33,750, (generously) assuming that mentoring administration accounted for 50% of their workload. This leads to a total estimated administrative cost of about \$92,806, and a net present value of the mentorship program equal to approximately \$770,000. It is possible that this is a lower bound on long-term value from the program, as increasing frontline sales worker productivity may allow sales managers to have larger spans of control (Espinosa and Stanton, 2021).

If instead the 123 agents in the Voluntary-Condition who were mentored had instead been under the Mandatory-Condition (which would have given mentoring to those agents who opted out), the gains would have been substantial. Over the first two months alone, the average revenue

[27] We refer to the survival rate trends in Figure 3a to determine the numbers of agents who remained in the firm at different points in time, and we take the average between the month-begin and month-end employment numbers to capture the number of agents present throughout a given month.

gain for treated agents in the Mandatory-Condition was \$3,100. Applying this to the 123 slots for mentorship that went to agents in the Voluntary-Condition suggests misallocation cost the firm around \$380,000 relative to a program with administrative costs of under \$100,000. Taken out to the 6 month horizon, the loss from choosing the wrong type of program to administer exceeds \$800,000.

7.4 External validity

Our study contributes to a first wave of empirical evidence on mentorship and workplace programs more broadly, and our setting has several strengths in terms of external validity. As part of the first wave of empirical evidence, we made decisions to give us the best chance to empirically test the theory, ensuring high internal validity (List, 2020). Despite these choices, performing our experiment in the field provides several strengths in terms of applying our findings to other settings.

First, the participants in our study appear approximately representative of the population they are drawn from on multiple dimensions, and they are representative of the broader population of workers in the United States. The participants in our study are new workers at a representative call center in Utah. In terms of hourly earnings and gender composition, our study participants are similar to the national- and state-level workers in similar occupations (customer service representatives in SOC code 43405, telemarketers in SOC code 41904, and miscellaneous sales representatives in SOC code 41309), based on data from the 2015–2019 5-year American Community Survey. A comparison of hourly earnings shows that mentor-eligible agents' earnings are quite similar to those in similar roles nationally and at the state level. The average hourly earnings among agents in our sample was about \$21. The average worker in similar roles at the national-level earns approximately \$23 per hour, and the average state-level worker in these occupations earns about \$20 per hour.[28]

As shown in Table 2, approximately 42% of the mentor-eligible agents were female, which is below the national- and state-level averages of 61% and 59%, respectively. However, if we exclude customer service call center jobs and only focus on telemarketers and sales representatives, our sample is more gender diverse than the national- and state-level averages (36% and 32% female,

[28]To construct hourly earnings in the ACS data, we divide total individual income by the product of weeks worked last year and usual hours per week.

respectively). The average age of the mentor-eligible agents was about 23 years old. The average age of seasoned veterans was approximately 26 years old. Thus, the agents in our sample are younger than the average age of workers in similar occupations nationally (37 years old) and in the state in which the call centers are located (33 years old). The study firm's workforce management group makes a concerted effort to recruit and hire workers from a broad range of demographic and socioeconomic groups. In general, our study sample appears to be fairly representative of the workforce in general, although because this is an entry-level job, ages are lower than average. To understand how our setting compares to broader occupations, note that among all individuals working for wages in any occupation in the state, average and median hourly earnings are about \$26 and \$19, respectively, with an average age of 37. Forty-six percent of those working for wages in the ACS are female.

Second, the task that we asked agents to perform—reflecting on their work, sharing these thoughts with their mentors, and acting on their mentors' advice—was a natural extension of their day-to-day activities. Whereas other experiments might ask participants to perform tasks that are wholly unrelated to their job and current knowledge set, our mentorship protocol provided mentors and protégés with a structured approach to discuss and learn from job-related successes and struggles.

Third, our intervention was done at scale relative to the firm, as all new hires within our study firm participated in this large-scale mentorship program, subject to the supply of mentors. Pilot programs often have different features than large-scale programs. For example, a pilot program may benefit from using only the best mentors.These scaling complications are less of a concern in our setting because the mentorship program was implemented as part of the firm's regular hiring process. Hence, our two main takeaways—(1) that mentorship is beneficial and (2) that those who benefit the most from mentorship might opt out if given the choice—provide clear implications for organizations considering large-scale mentorship programs.

As a final point regarding external validity, we consider our findings in relation to those in Sandvik et al. (2020). We estimate that mentorship increases newly hired agents' revenue-per-call by 10% (see Column (2) of Table 3). Sandvik et al. (2020) implemented a "structured-meetings" protocol that called for a structured conversation early in the week, followed by a second, unstructured conversation over lunch at the end of the week. They find revenue gains in excess of 20% across the entire population of seasoned agents. While the protocols in the two experiments are similar—e.g., each randomly paired sales agents together and asked them to discuss job-specific struggles and successes over four weeks—they differ in several aspects that likely contribute to the difference in effect sizes. Whereas Sandvik et al. (2020) paired together two experienced sales agents, the mentorship program studied here pairs together a newly hired sales agent with a seasoned sales agent, and new employees potentially benefit from peer-effects differently than do existing employees. For example, mentorship may help newly hired sales agents learn how to close entry-level deals (i.e., selling a baseline product), whereas knowledge exchange among veterans may help agents learn how to close high-level deals (i.e., up-selling to premium services packages). In addition, the intervention in Sandvik et al. (2020) was approximately twice as intensive as that described here due to the second, unstructured conversation at the end of each week of treatment. Finally, the mentoring protocol clearly specifies who will be providing information (the mentor), and who will be receiving said information (the protégé), whereas Sandvik et al. (2020) randomly paired employees and treated them as equals. The most important difference, however, is that the prior study did not have a channel to test voluntary participation in the experiment and compliance rates in the prior study were high. Making the decision to opt out salient in the Voluntary-Condition here was absent in the prior work and is the main object of interest in this paper.

8 Conclusion

Employees gain human capital through learning on the job. The importance of firm-specific human capital means that workplace programs are often important tools for workers and firms. The evaluation of these programs, however, is often difficult because of lacking data and selection concerns in who participates. We provide new evidence on the extent and nature of workplace programs, showing that many are voluntary and have substantial rates of non-participation. We then offer a novel evaluation of a mentorship program for new hires by running an experiment to (1) separate treatment and selection effects and (2) identify heterogeneous treatment effects that vary with the propensity to participate in the mentorship program.

We find that mentorship programs can have a large positive effect on productivity. Specifically, in one group of our experiment, labeled the Mandatory-Condition, individuals in a sales firm who were randomly assigned a mentor had revenues that were 17% higher than agents randomly not assigned a mentor in the first two months on the job. In contrast, new hires in a second group, labeled the Voluntary-Condition, who were given the opportunity to opt in or out of the program before random mentor assignment, had very different treatment effects. Agents who opted into the mentorship program and received a mentor did not have higher revenues than agents that also opted in but did not receive a mentor. This finding underscores the practical importance of the potential for selection bias in program recruitment to alter inference about program efficacy. The direction of in a RCT depends on the correlation between treatment gains and participation propensity, which we find is negative in our mentorship program application. This selection effect relates to the wellness program participation patterns found in Jones et al. (2019) and the site selection bias identified in Allcott (2015), though the latter estimated a positive relationship between selection into participation and treatment effects, whereas we estimate a negative relationship.

Our design also allows us to conclude that the mentoring program would have had the largest effect for agents who would opt out if given the opportunity. On-the-job training programs, like this one, may have the largest impact when they are implemented broadly rather than delivered to a subset of workers that select into participation. In our setting, the on-the-job training delivered through mentorship is a substitute for ability. Said differently, training that leverages help from coworkers can lift those that struggle the most, but these workers might be the *least* likely to seek out the resources for improvement.

There are many reasons why workers may choose to opt out of programs that could improve their productivity. First, those that are struggling may not want to admit their difficulties and may be less likely to ask for help. Second, those that are struggling may not know what they do not know—i.e., they may not realize they are struggling or understand how training/mentorship programs could help them. Finally, those that are struggling may be the least engaged, which could jointly explain the low performance and low uptake of the program in our setting. Taken together, our results suggest that productivity-improving information may be hidden from some workers, but managerial interventions can help trigger the dissemination of it.

Given these possible reasons for why workers might opt out, we attempted to identify the demographic and personality factors that drove some workers to self-select out of the program. Unfortunately, sales agents' demographic information, personality traits, and prior work experiences do little to explain variation in program participation. Combined with little treatment effect heterogeneity based on worker characteristics, improved program targeting in our setting would be difficult. As a result, we suspect examining why some workers do and do not participate in workplace training, mentoring, or other potentially beneficial programs is a fruitful area for future research.

Implementing an on-the-job randomized control trial in a call center allows us to provide novel and generalizable insights. Some findings, however, may be context specific and warrant future investigations. For example, we find that the on-the-job training provided through the mentoring program is a substitute for ability. This may not be a general phenomenon, and in other contexts training may lift productivity most for the best workers. Similar experiments can help to determine how the returns to training programs vary over the distribution of worker productivity. Finally, an important question for the COVID-19 era is whether interventions done in-person will translate to remote environments. Some emerging evidence highlights the power of virtual meetings for workers' careers, suggesting our results will likely scale beyond face-to-face interactions (Bojinov et al., 2021).

References

Acemoglu, Daron, Jörn-Steffen Pischke. 1998. Why do firms train? theory and evidence. *The Quarterly journal of economics* **113**(1) 79–119.

Allcott, Hunt. 2015. Site selection bias in program evaluation. *The Quarterly Journal of Economics* **130**(3) 1117–1165.

Allen, Tammy D, Lillian T Eby, Georgia T Chao, Talya N Bauer. 2017. Taking stock of two relational aspects of organizational life: Tracing the history and shaping the future of socialization and mentoring research. *Journal of Applied Psychology* **102**(3) 324.

Anderson, Michael L. 2008. Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American statistical Association* **103**(484) 1481–1495.

Bandiera, Oriana, Iwan Barankay, Imran Rasul. 2005. Social preferences and the response to incentives: Evidence from personnel data. *The Quarterly Journal of Economics* **120**(3) 917–962.

Bandiera, Oriana, Iwan Barankay, Imran Rasul. 2013. Team incentives: Evidence from a firm-level experiment. *Journal of the European Economic Association* **11**(5) 1079–1114.

Björklund, Anders, Robert Moffitt. 1987. The estimation of wage gains and welfare gains in self-selection models. *The Review of Economics and Statistics* 42–49.

Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, John Van Reenen. 2019. What drives differences in management practices? *American Economic Review* **109**(5) 1648–83.

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

Bloom, Nicholas, John Van Reenen. 2007. Measuring and explaining management practices across firms and countries. *The quarterly journal of Economics* **122**(4) 1351–1408.

Bojinov, Iavor, Prithwiraj Choudhury, Jacqueline N Lane. 2021. Virtual watercoolers: A field experiment on virtual synchronous interactions and performance of organizational newcomers. *Harvard Business School Technology & Operations Mgt. Unit Working Paper* (21-125).

Bol, Jasmijn C, Justin Leiby. 2018. Subjectivity in professionals' incentive systems: Differences between promotion-and performance-based assessments. *Contemporary Accounting Research* **35**(1) 31–57.

Bruhn, Miriam, Dean Karlan, Antoinette Schoar. 2018. The impact of consulting services on small and medium enterprises: Evidence from a randomized trial in mexico. *Journal of Political Economy* **126**(2) 635–687.

Carrell, Scott E., Bruce I. Sacerdote, James E. West. 2013. From natural variation to optimal policy? The importance of endogenous peer group formation. *Econometrica* **81**(3) 855–882.

Carter, Susan Payne, Whitney Dudley, David S Lyle, John Z Smith. 2019. Who's the boss? the effect of strong leadership on employee turnover. *Journal of Economic Behavior & Organization* **159** 323–343.

Chandrasekhar, Arun G., Benjamin Golub, He Yang. 2016. Signaling, stigma, and silence in social learning. *Working Paper* .

Chatterji, Aaron, Solène Delecourt, Sharique Hasan, Rembrand Koning. 2019. When does advice impact startup performance? *Strategic Management Journal* **40**(3) 331–356.

Edmondson, Amy C., Zhike Lei. 2014. Psychological safety: The history, renaissance, and future of an interpersonal construct. *Annual Review of Organizational Psychology and Organizational Behavior* **1**(1) 23–43. doi:10.1146/annurev-orgpsych-031413-091305.

Englmaier, Florian, Stefan Grimm, Dominik Grothe, David Schindler, Simeon Schudy. 2021. The value of leadership: Evidence from a large-scale field experiment .

Englmaier, Florian, Andreas Roider, Uwe Sunde. 2017. The role of communication of performance schemes: Evidence from a field experiment. *Management Science* **63**(12) 4061–4080.

Espinosa, Miguel, Christopher Stanton. 2021. Worker skills and organizational spillovers: Evidence from linked training and communications data. Tech. rep., Harvard Business School.

Friebel, Guido, Matthias Heinz, Miriam Krueger, Nikolay Zubanov. 2017. Team incentives and performance: Evidence from a retail chain. *American Economic Review* **107**(8) 2168–2203.

Friebel, Guido, Michael Kosfeld, Gerd Thielmann. 2019. Trust the police? self-selection of motivated agents into the german police force. *American Economic Journal: Microeconomics* **11**(4) 59–78.

Fudenberg, Drew, Luis Rayo. 2019. Training and effort dynamics in apprenticeship. *American Economic Review* **109**(11) 3780–3812.

Gibbons, Robert, Rebecca Henderson. 2012. *What do managers do?: Exploring persistent performance differences among seemingly similar enterprises*. Harvard Business School.

Ginther, Donna K, Janet M Currie, Francine D Blau, Rachel TA Croson. 2020. Can mentoring help female assistant professors in economics? an evaluation by randomized trial. *AEA Papers and Proceedings*, vol. 110. 205–09.

Gosnell, Greer K, John A List, Robert D Metcalfe. 2020. The impact of management practices on employee productivity: A field experiment with airline captains. *Journal of Political Economy* **128**(4) 1195–1233.

Gutner, Toddi. 2009. Finding anchors in the storm: Mentors. *The Wall Street Journal* .

Harrison, Glenn W, John A List. 2004. Field experiments. *Journal of Economic literature* **42**(4) 1009–1055.

Herbst, Daniel, Alexandre Mas. 2015. Peer effects on worker output in the laboratory generalize to the field. *Science* **350**(6260) 545–549.

Hoffman, Mitchell, Stephen V Burks. 2020. Worker overconfidence: Field evidence and implications for employee turnover and firm profits. *Quantitative Economics* **11**(1) 315–348.

Hoffman, Mitchell, Lisa B. Kahn, Danielle Li. 2017. Discretion in hiring. *The Quarterly Journal of Economics* **133**(2) 765–800.

Hoffman, Mitchell, Steven Tadelis. 2021. People management skills, employee attrition, and manager rewards: An empirical analysis. *Journal of Political Economy* **129**(1) 000–000.

Hong, Fuhai, Tanjim Hossain, John A List. 2015. Framing manipulations in contests: a natural field experiment. *Journal of Economic Behavior & Organization* **118** 372–382.

Hossain, Tanjim, John A List. 2012. The behavioralist visits the factory: Increasing productivity using simple framing manipulations. *Management Science* **58**(12) 2151–2167.

Jones, Damon, David Molitor, Julian Reif. 2019. What do workplace wellness programs do? evidence from the illinois workplace wellness study. *The Quarterly Journal of Economics* **134**(4) 1747–1791.

Kram, Kathy E. 1988. *Mentoring at work: Developmental relationships in organizational life.*. University Press of America.

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

Lee, David S. 2009. Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *The Review of Economic Studies* **76**(3) 1071–1102.

List, John A. 2020. Non est disputandum de generalizability? a glimpse into the external validity trial. Tech. rep., National Bureau of Economic Research.

Lyle, David S, John Z Smith. 2014. The effect of high-performing mentors on junior officer promotion in the us army. *Journal of Labor Economics* **32**(2) 229–258.

Mas, Alexandre, Enrico Moretti. 2009. Peers at work. *American Economic Review* **99**(1) 112–45.

Mills, Joyce F, Anna C Mullins. 2008. The california nurse mentor project: Every nurse deserves a mentor. *Nursing Economics* **26**(5) 310.

Oyer, Paul, Scott Schaefer. 2011. Personnel economics: Hiring and incentives. *Handbook of Labor Economics* **4** 1769–1823.

Payne, Stephanie C, Ann H Huffman. 2005. A longitudinal examination of the influence of mentoring on organizational commitment and turnover. *Academy of Management Journal* **48**(1) 158–168.

Porter, Catherine, Danila Serra. 2020. Gender differences in the choice of major: The importance of female role models. *American Economic Journal: Applied Economics* **12**(3) 226–54.

Rockoff, Jonah E. 2008. Does mentoring reduce turnover and improve skills of new employees? evidence from teachers in new york city. Tech. rep., National Bureau of Economic Research.

Sandvik, Jason J, Richard E Saouma, Nathan T Seegert, Christopher T Stanton. 2020. Workplace knowledge flows. *The Quarterly Journal of Economics* **135**(3) 1635–1680.

Syverson, Chad. 2011. What determines productivity? *Journal of Economic literature* **49**(2) 326–65.

Figures and Tables

Figure 1: Allocation of Cohorts and Agents to Treatment Conditions

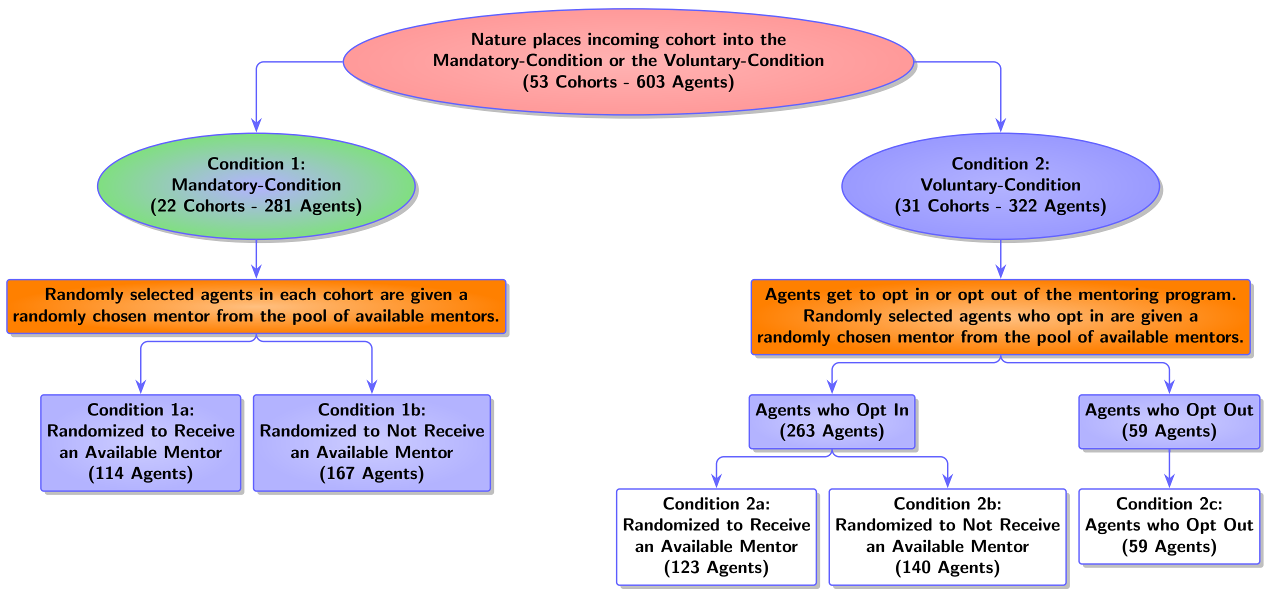

Note: This figure displays the allocation of the 53 mentor-eligible cohorts to either the Mandatory-Condition or the Voluntary-Condition, our first level of variation. It then shows the allocation of the 603 mentor-eligible agents within these cohorts into individual different treatment conditions, our second level of variation.

Figure 2: Effect of Mentoring on Productivity

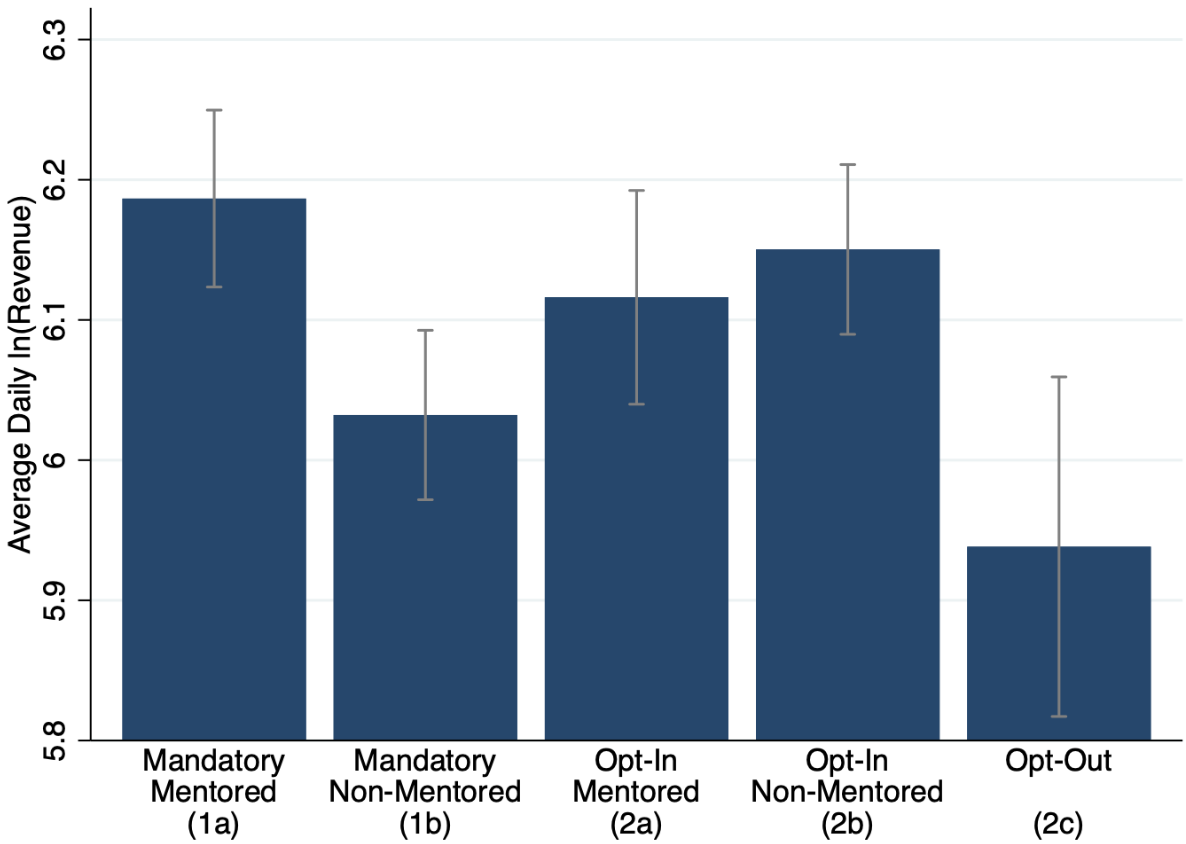

Note: This figure plots the average daily log revenue and 95% confidence intervals for agents' first two months on the sales floor, split by treatment condition. Before we aggregate revenue amounts, we net out cohort fixed effects and then we add in the average productivity level across all agents that did not receive a mentor as a baseline. The p-values from difference-in-means tests that compare various treatment conditions are as follows. (1a) = (1b) p-value $<$ 0.001; (2a) =(2b) p-value = 0.331; (2b) = (2c) p-value $<$ 0.001; (1a)-(1b) = (2a)-(2b) p-value $<$ 0.001. Similar bar charts are displayed in Figure A.2 to capture agents' productivity during months 3–6 on the sales floor.

Figure 3: Effect of Mentoring on Retention

(a) Agents in the Mandatory-Condition

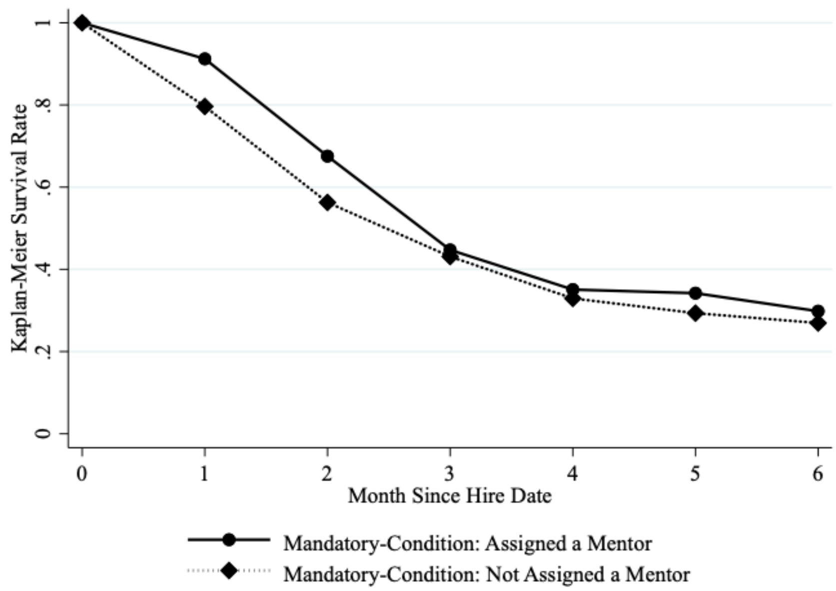

(b) Agents in the Voluntary-Condition

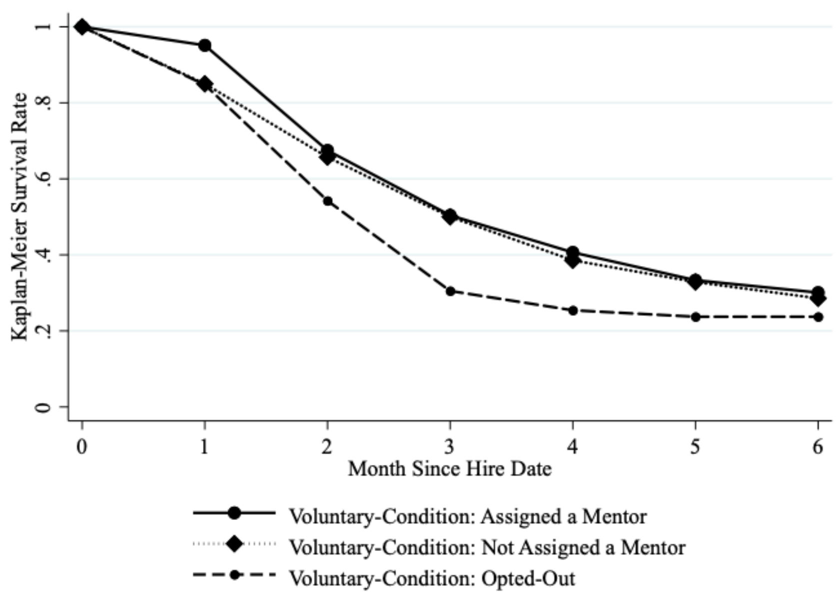

Note: Figure (a) plots Kaplan-Meier survival rates over time for agents in the Mandatory-Condition, and Figure (b) considers those in the Voluntary-Condition. The survival rate estimator considers a starting point, in our case an agent's hire date, and then, from that time, displays the fraction of agents that remain at the firm.

Figure 4: Robustness of the Treatment and Selection Effects of Mentoring on Productivity

(a) ln(Revenue)

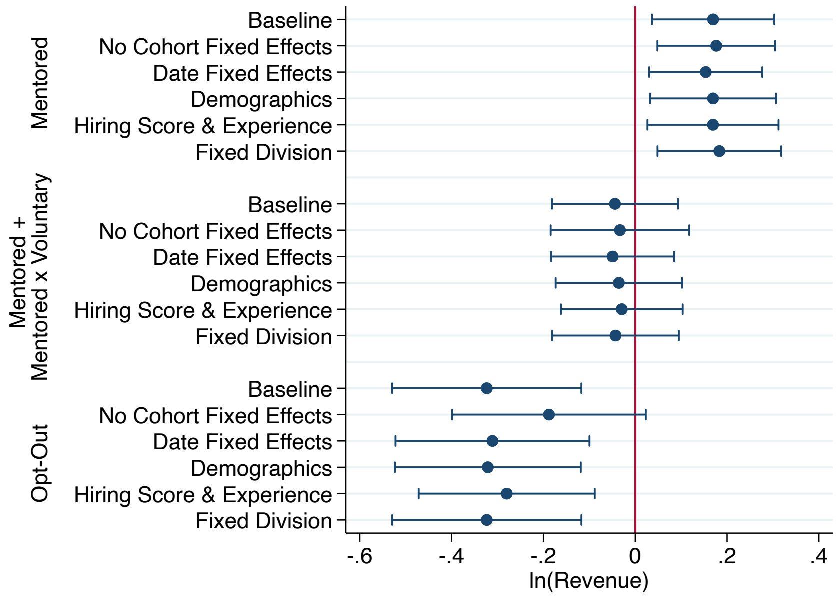

(b) ln(RPC)

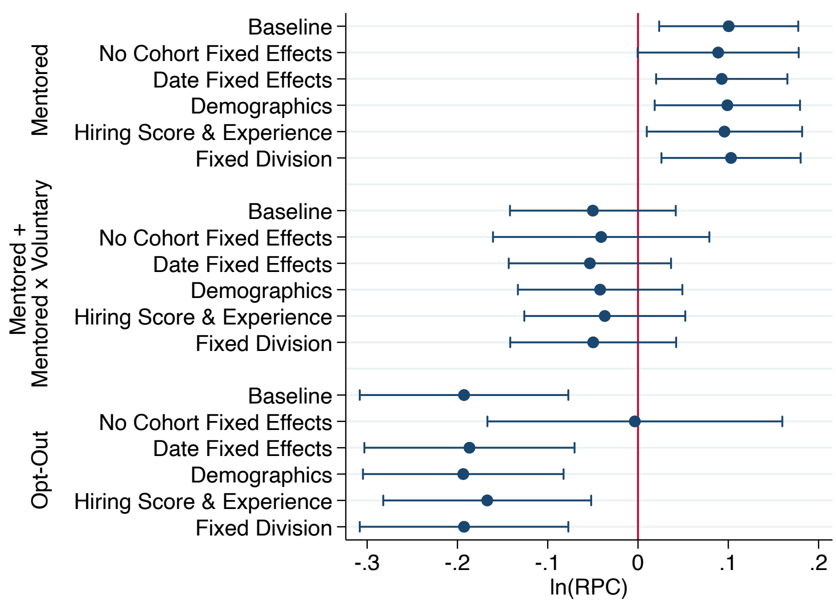

Note: These figures plot the regression coefficients (and 95% confidence intervals) on *Mentored*, the sum of *Mentored* and *Mentored* × *Voluntary*, and *Opt-Out* from Equation (1). We use log revenue as the dependent variable in Figure (a) and log RPC in Figure (b). The "Baseline" estimation in Figure (a) (Figure (b)) replicates the result from Column (1) (Column (2)) of Table 3. The second estimation excludes cohort fixed effects. The third includes date fixed effects. The fourth estimation includes controls for the agent's demographic characteristics: age, gender, marital status, and referral status. The fifth estimation includes additional controls for the agent's hiring score, previous call center experience, and previous sales experience. The sixth estimation removes observations in which agents are no longer working in the division in which they were initially hired.

Figure 5: Responses to Wrap-Up Survey

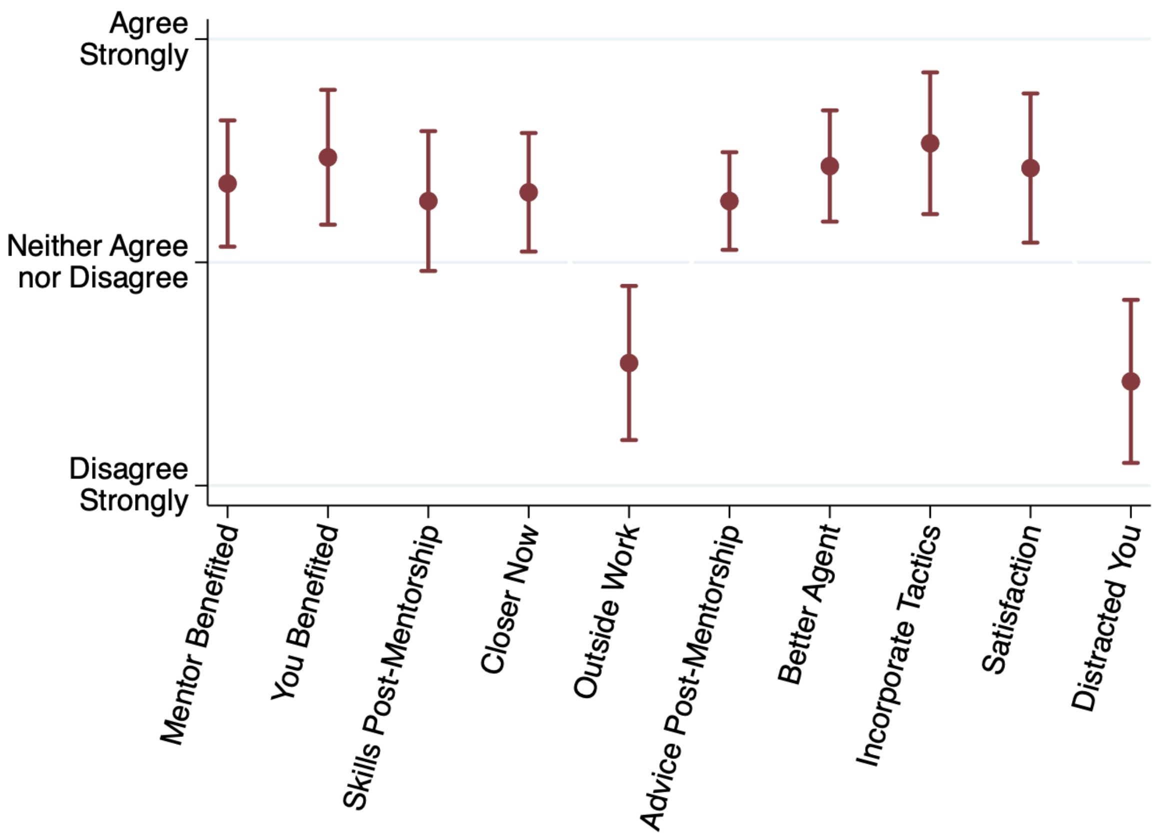

Note: This figures plots the average values and 95% confidence intervals for responses to the wrap-up survey questions. All responses were made on a scale from -3 to 3, with -3 indicating "Disagree Strongly," 0 indicating "Neither Agree nor Disagree," and 3 indicating "Agree Strongly." The statements, from left to right, are as follows: "Your mentor benefited from the mentoring relationship"; "You benefited from the mentoring relationship"; "Since your formal meetings have ended, your mentor has continued to teach you skills to help you make more sales"; "You and your mentor are closer now than you were during the mentor program"; "Since your formal meetings have ended, you and your mentor have spent time together outside of the office"; "Since your formal meetings have ended, you have reached out to your mentor for help/advice"; "You have become a better sales agent as the result of being mentored"; "Being mentored helped you incorporate important selling tactics into your day-to-day work"; "Having a mentor increased your day-to-day satisfaction at work"; "Being mentored distracted you from reaching your potential each week." Seventeen protégés completed the wrap-up survey.

Table 1: Survey Data on Characteristics of Workplace Programs and Participation in Voluntary Programs

Program	Offered?	Voluntary if Offered?	Non-Participation if Voluntary?
Formal Mentorship	0.45	0.59	0.27
	(0.01)	(0.01)	(0.02)
New Hire Training	0.87	0.22	0.21
	(0.01)	(0.01)	(0.02)
Ongoing Training or Cont. Ed.	0.80	0.43	0.28
	(0.01)	(0.01)	(0.01)
N = 3,191			

Notes: This table displays means and standard deviations (in parentheses) for the prevalence and nature of a variety of workplace programs taken from responses from a nationally representative online survey conducted through the Lucid platform in June of 2022. The survey was restricted to respondents currently employed by others. Respondents were asked about whether their employer offers a program and whether it is voluntary or mandatory with the question: "Consider your current employer. Which of the following programs does your employer offer to you personally? If offered, are you required to participate (required/mandatory) or can you choose to participate or not (optional/voluntary)?" For each program, respondents chose between "Required or Mandatory," "Optional or Voluntary," or "Not offered". For the three core programs—mentorship, new hire training, and continuing education—if a respondent indicated that a program was voluntary, follow-up questions were asked about their participation and the reasons for their lack of participation, if applicable.

Table 2: Balance in Agent Demographics

Panel A: Cohort-Level Balance in Agent Characteristics

	Mandatory-Condition	Voluntary-Condition	p-value
	(1)	(2)	(2)−(1)
Age (yrs.)			
Mean	22.89	23.11	0.762
Std Dev.	2.44	2.68	
Woman			
Mean	0.44	0.40	0.338
Std Dev.	0.14	0.18	
Married			
Mean	0.14	0.15	0.787
Std Dev.	0.09	0.17	
Hiring Score			
Mean	0.84	0.85	0.200
Std Dev.	0.04	0.04	
N Cohorts	22	31	

Panel B: Cohort-Level Balance in Ex Ante Productivity

	Mandatory-Condition	Voluntary-Condition	p-value
	(1)	(2)	(2)−(1)
Revenue			
Mean	760.27	793.95	0.386
Std Dev.	112.95	153.25	
RPC			
Mean	46.80	49.27	0.361
Std Dev.	8.42	10.40	
RPH			
Mean	115.61	120.21	0.387
Std Dev.	13.69	21.83	
Calls			
Mean	17.18	17.18	0.996
Std Dev.	0.98	1.32	
Hours			
Mean	6.57	6.59	0.839
Std Dev.	0.39	0.39	
Adherence			
Mean	0.82	0.84	0.283
Std Dev.	0.04	0.04	
Conversion			
Mean	0.23	0.22	0.375
Std Dev.	0.03	0.03	
Number of Cohorts	22	31	

Panel C: Agent-Level Balance in Agent Characteristics

	Mandatory-Condition			Voluntary-Condition			
	Mentored	Non-Mentored	p-value	Mentored	Non-Mentored	p-value	Opted-Out
	(1)	(2)	(2)−(1)	(3)	(4)	(4)−(3)	(5)
Age (yrs.)							
Mean	22.41	23.82	0.128	22.56	22.79	0.773	23.30
Std Dev.	(4.39)	(9.17)		(5.79)	(6.70)		(9.08)
Woman							
Mean	0.46	0.41	0.477	0.45	0.42	0.676	0.34
Std Dev.	(0.50)	(0.49)		(0.50)	(0.50)		(0.48)
Married							
Mean	0.11	0.15	0.281	0.14	0.18	0.375	0.15
Std Dev.	(0.31)	(0.36)		(0.35)	(0.38)		(0.36)
Hiring Score							
Mean	0.82	0.84	0.144	0.85	0.85	0.997	0.83
Std Dev.	(0.09)	(0.08)		(0.08)	(0.07)		(0.09)
Referral							
Mean	0.60	0.53	0.251	0.55	0.60	0.442	0.58
Std Dev.	(0.49)	(0.50)		(0.50)	(0.49)		(0.50)
Extroversion							
Mean	4.69	4.67	0.932	5.01	5.08	0.673	4.77
Std Dev.	(1.26)	(1.35)		(1.25)	(1.27)		(1.52)
Agreeableness							
Mean	5.05	5.01	0.770	5.31	5.42	0.424	5.42
Std Dev.	(1.02)	(1.04)		(1.12)	(1.04)		(1.22)
Conscientiousness							
Mean	6.06	5.85	0.102	6.07	6.10	0.740	6.04
Std Dev.	(0.79)	(1.03)		(0.88)	(0.87)		(0.92)
Emotional Stability							
Mean	5.55	5.47	0.648	5.65	5.72	0.642	5.63
Std Dev.	(1.06)	(1.21)		(1.19)	(0.99)		(1.07)
Openness							
Mean	6.12	5.93	0.106	6.16	6.21	0.605	6.15
Std Dev.	(0.80)	(0.86)		(0.80)	(0.79)		(0.87)
Call Center Exp.							
Mean	0.56	0.43	0.226	0.25	0.39	0.044	0.50
Std Dev.	(0.50)	(0.50)		(0.44)	(0.49)		(0.51)
Sales Experience							
Mean	0.69	0.71	0.840	0.52	0.62	0.216	0.59
Std Dev.	(0.47)	(0.46)		(0.50)	(0.49)		(0.50)
Number of Agents	114	167		123	140		59

Notes. In Panel A, we average agent characteristics to the cohort-level, then take averages across cohorts. In Panel B, we take average productivity measures of agents who were not mentorship eligible within each sales division. Cohorts are assigned to a particular sales division, so the tests in Panel B estimate the balance in brand-level productivity measures between cohorts in the Mandatory-Condition versus those in the Voluntary-Condition. In Panel C, we take averages across agents within each treatment condition.

Table 3: Treatment and Selection Effects of Mentoring on Productivity and Retention

	ln(Revenue)	ln(RPC)	ln(RPH)	Adherence	$Tenure_1$	$Tenure_2$
	(1)	(2)	(3)	(4)	(5)	(6)
Mentored	0.170**	0.100**	0.104**	0.018*	0.102***	0.087
standard errors	(0.066)	(0.038)	(0.050)	(0.009)	(0.037)	(0.071)
sharpened q-value	[0.060]	[0.060]	[0.116]	[0.119]	[0.060]	[0.359]
Mentored × Voluntary	-0.214**	-0.150**	-0.146*	-0.014	-0.024	-0.096
standard errors	(0.095)	(0.060)	(0.080)	(0.011)	(0.053)	(0.089)
sharpened q-value	[0.086]	[0.060]	[0.159]	[0.345]	[0.723]	[0.403]
Voluntary Opt-Out	-0.323***	-0.193***	-0.215**	-0.008	-0.066	-0.171**
standard errors	(0.103)	(0.058)	(0.081)	(0.013)	(0.060)	(0.084)
sharpened q-value	[0.052]	[0.052]	[0.060]	[0.718]	[0.403]	[0.116]
Cohort Fixed Effects	✓	✓	✓	✓	✓	✓
Adj. R-Square	0.029	0.046	0.028	0.117	0.073	0.070
Observations	15,137	15,137	15,137	15,137	603	603
p-value: Mentored + Mentored × Voluntary	0.522	0.280	0.502	0.414	0.043	0.864
Baseline Retention					0.784	0.545

Notes. The sample used in Columns (1)–(4) is composed of agent-day productivity data for all mentor-eligible agents with post-training productivity data. The data covers agents' productivity on their first two months on the sales floor. *Mentored* equals one for agents who were randomized to received an available mentor, and zero otherwise, *Voluntary* equals one for agents in the Voluntary-Condition, and zero otherwise, and *Voluntary Opt-Out* equals one for agents who chose to opt-out of possibly receiving a mentor, and zero otherwise. Columns (5)–(6) use data with a single observation per unique hired agent to capture retention effects. $Tenure_t$ equals one for agents who achieve at least t months of tenure at the firm, and zero otherwise. We estimate ordinary least squares regressions in all columns. All specifications include cohort fixed effects. Standard errors are clustered by cohort and are reported in parentheses. Sharpened q-values are presented in brackets, following Anderson (2008). The penultimate row reports the p-values from post-estimation tests that the sum of the coefficients on *Mentored* and *Mentored* × *Voluntary* equal zero. The bottom row reports the baseline retention estimates, measured as the fraction of non-mentored agents in the Mandatory-Condition who achieve one month (two months) of tenure in Column (5) (Column (6)). *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. Long-term productivity results are displayed in Table A.2.

Table 4: Determinants of Opting Out

Dep. Variable	= 1 if Opted Out			
	(1)	(2)	(3)	(4)
Age	0.011	0.015	0.021	0.014
	(0.022)	(0.022)	(0.028)	(0.029)
Female	-0.388	-0.413	-0.038	-0.079
	(0.291)	(0.285)	(0.312)	(0.319)
Married	-0.156	-0.082	0.144	0.071
	(0.358)	(0.357)	(0.459)	(0.450)
Hiring Score		-4.060**	-4.270**	-4.740**
		(1.704)	(1.861)	(1.924)
Location			0.006	0.330
			(0.468)	(0.438)
Referral			0.119	0.213
			(0.423)	(0.461)
Extroversion			-0.145	-0.168
			(0.175)	(0.188)
Agreeableness			0.022	-0.030
			(0.206)	(0.201)
Conscientiousness			0.007	0.102
			(0.290)	(0.298)
Emotional Stability			0.047	0.029
			(0.170)	(0.180)
Openness			-0.011	0.085
			(0.254)	(0.251)
Call Center Exp.				0.727
				(0.536)
Sales Experience				-0.043
				(0.561)
Division Fixed Effects			✓	✓
Pse. R-Square	0.007	0.024	0.032	0.066
Observations	322	322	295	295

Notes. This sample is restricted to the 322 agents in the Voluntary-Condition. The dependent variable is an indicator that equals one if the agent chose to opt out, and zero otherwise. We run logistic regressions of this indicator on different potential predictors of the choice to opt out. Limited personality data reduces the sample sizes in Columns (3) and (4). Columns (3) and (4) include fixed effects for agents in the two largest divisions. In Column (4) we include indicator variables that capture previous work experience. Only agents hired after May 27th were asked about their previous call center and sales experience, so we fill in missing values as zero. We also include a dummy equal to one for agents with missing work experience data, and zero otherwise, as a control in Column (4). Standard errors are clustered by cohort and are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 5: Estimated Treatment Effect of Mentoring Among Opt-Out Agents

Panel A: Pre-Registered Estimates of Opt-Out Treatment Effect (Months 1–2)

	ln(Revenue)	ln(RPC)	ln(RPH)	Adherence
	(1)	(2)	(3)	(4)
Opt-Out Mentored Effect	1.212**	0.633*	0.757*	0.130**
	(0.553)	(0.345)	(0.439)	(0.063)

Panel B: GMM Estimation With Cohort FE (Months 1–2)

	ln(Revenue)	ln(RPC)	ln(RPH)	Adherence
	(1)	(2)	(3)	(4)
Opt-In Baseline Effect	0.331***	0.200***	0.217**	0.007
	(0.116)	(0.069)	(0.093)	(0.011)
Opt-In Mentored Effect	-0.082	-0.073	-0.076	0.005
	(0.072)	(0.048)	(0.066)	(0.006)
Opt-Out Mentored Effect	0.794**	0.543**	0.561*	0.057
	(0.359)	(0.234)	(0.294)	(0.039)
Opt-In Likelihood		0.729		

Panel C: GMM Estimation Without Cohort FE (Months 1–2)

	ln(Revenue)	ln(RPC)	ln(RPH)	Adherence
	(1)	(2)	(3)	(4)
Opt-In Baseline Effect	0.436**	0.181	0.376**	0.090***
	(0.177)	(0.142)	(0.150)	(0.027)
Opt-In Mentored Effect	-0.066	-0.057	-0.073	0.003
	(0.074)	(0.057)	(0.065)	(0.007)
Opt-Out Mentored Effect	0.531*	0.281	0.232	-0.019
	(0.290)	(0.221)	(0.228)	(0.036)
Opt-In Likelihood		0.729		

Notes. The results in Panel A show estimates of the treatment effect among agents who opt out of mentorship. To estimate standard errors, we block-bootstrap by cohort (N = 53) over the whole procedure, with 500 bootstrap replications for each column. Panels B and C report the GMM estimates as described in the text. The Opt-In Baseline Effect is the difference in productivity among untreated agents who participate in the program and those who opt out. The Opt-In Mentored Effect is the effect of mentoring for agents who opt in relative to their baseline effect. The Opt-Out Mentored Effect is the effect of mentoring in the mandatory program for those who would not participate in the voluntary program. We include (exclude) cohort fixed effects in the Panel B (Panel C) estimations of equation (4). Standard errors for the GMM estimation are clustered by cohort. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 6: Testing for Discouragement, Leakage, and Crowd-Out

Panel A: Comparing New Hires and Veterans in Mentor-Eligible and Non-Eligible Cohorts

	Ln(Revenue): Months 1–2			Ln(RPC): Months 1–2		
	(1)	(2)	(3)	(4)	(5)	(6)
New Hire	-0.422***			-0.483***		
	(0.103)			(0.068)		
New Hire × Mandatory	0.027	0.076	0.082	-0.010	0.026	0.031
	(0.094)	(0.101)	(0.103)	(0.055)	(0.059)	(0.060)
New Hire × Voluntary	0.079	0.098	0.099	0.083*	0.082	0.081
	(0.083)	(0.088)	(0.088)	(0.049)	(0.052)	(0.052)
Mentored × Mandatory	0.208***	0.213***	0.202***	0.106**	0.115**	0.107**
	(0.071)	(0.071)	(0.076)	(0.046)	(0.046)	(0.048)
Mentored × Voluntary	-0.040	-0.042	-0.035	-0.042	-0.048	-0.044
	(0.073)	(0.074)	(0.072)	(0.049)	(0.048)	(0.047)
Division-Location-Date FE	✓	✓	✓	✓	✓	✓
Location-New Hire FE	✓			✓		
Division-Location-New Hire FE		✓	✓		✓	✓
New Hire-Date FE		✓	✓		✓	✓
Demographic Controls			✓			✓
Adj. R-Square	0.084	0.089	0.100	0.112	0.117	0.124
Observations	41,867	41,858	41,858	41,867	41,858	41,858
$New_B = 0$, $New_S = 0$	0.622	0.527	0.513	0.139	0.273	0.279
New_B - $New_S = 0$	0.569	0.809	0.854	0.090	0.316	0.372
New_B - New_S + $Men_B = 0$	0.071	0.034	0.041	0.803	0.286	0.308

Panel B: Comparing New Hires and Veterans in Mentor-Eligible and Pre-Experimental Cohorts

	Ln(Revenue): Months 1–2			Ln(RPC): Months 1–2		
	(1)	(2)	(3)	(4)	(5)	(6)
New Hire	-0.503***			-0.366***		
	(0.054)			(0.038)		
New Hire × Mandatory	0.211**	0.217**	0.083	-0.037	-0.023	-0.094
	(0.094)	(0.095)	(0.096)	(0.058)	(0.059)	(0.060)
New Hire × Voluntary	0.218**	0.246***	0.105	0.021	0.037	-0.040
	(0.087)	(0.091)	(0.089)	(0.054)	(0.057)	(0.057)
Mentored × Mandatory	0.237***	0.228***	0.206***	0.130***	0.128***	0.114**
	(0.076)	(0.075)	(0.077)	(0.048)	(0.048)	(0.049)
Mentored × Voluntary	-0.030	-0.029	-0.028	-0.034	-0.039	-0.038
	(0.073)	(0.073)	(0.071)	(0.048)	(0.048)	(0.046)
Division-Location-Date FE	✓	✓	✓	✓	✓	✓
Location-New Hire FE	✓			✓		
Division-Location-New Hire FE		✓	✓		✓	✓
Demographic Controls			✓			✓
Adj. R-Square	0.160	0.161	0.166	0.165	0.165	0.168
Observations	73,359	73,359	73,359	73,359	73,359	73,359
$New_B = 0$, $New_S = 0$	0.022	0.016	0.481	0.582	0.550	0.286
New_B - $New_S = 0$	0.935	0.756	0.811	0.300	0.281	0.337
New_B - New_S + $Men_B = 0$	0.007	0.020	0.030	0.180	0.204	0.254

Notes. Panel A reports tests of the net effect of discouragement and leakage by comparing the performance of three groups of agents: (1) new hires who were in mentor-eligible hiring cohorts; (2) new hires who were not in mentor-eligible hiring cohorts during the time of the experiment; and (3) seasoned veterans who began working at the firm before the onset of the mentorship program. The specifications are described in Section 6.1. The dependent variable is ln(Revenue) in Columns (1)–(3) and ln(RPC) in Column (4)–(6). In the bottom three rows, New_B stands for new hire in the Mandatory-Condition, New_S stands for new hire in the Voluntary-Condition, and Men_B stands for those who were mentored in the Mandatory-Condition. The estimations in Panel B are analogous to those in Panel A, with the exception that instead of using new hires from non-mentor-eligible cohorts as a control group, we use new hires who began working at the firm before the mentorship program began. This allows us to test for crowd-out effects. Standard errors are clustered by agent and are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 7: Balanced Panel with Productivity Replacement

	ln(Revenue)	ln(RPC)	ln(RPH)
	(1)	(2)	(3)
Mentored	0.209**	0.118**	0.133*
	(0.101)	(0.058)	(0.072)
Mentored × Voluntary	-0.185	-0.124	-0.124
	(0.135)	(0.078)	(0.101)
Voluntary Opt-Out	-0.280***	-0.152***	-0.180***
	(0.085)	(0.045)	(0.064)
Cohort Fixed Effects	✓	✓	✓
Adj. R-Square	0.023	0.024	0.022
Observations	23,392	23,392	23,392

Notes. The results in this table show estimates of the treatment effects of mentorship when accounting for the replacement productivity of newly hired agents who replace terminated agents. For mentor-eligible agents who leave the firm before the two-month mark, we extend the time series of their productivity provision to two months and replace their post-termination productivity values with the average productivity of a newly hired replacement agent. We do this by computing the average productivity of newly hired non-eligible agents in the same location-division-year-quarter as the departed agent. We then re-estimate our main regression models using Equation (1). Standard errors are clustered by cohort and are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

A Appendix Figures and Tables

Figure A.1: Comparisons of Agent Characteristics

(a) Gender

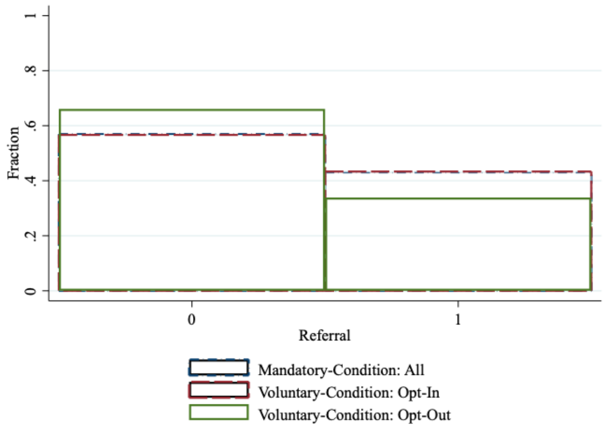

(b) Marital Status

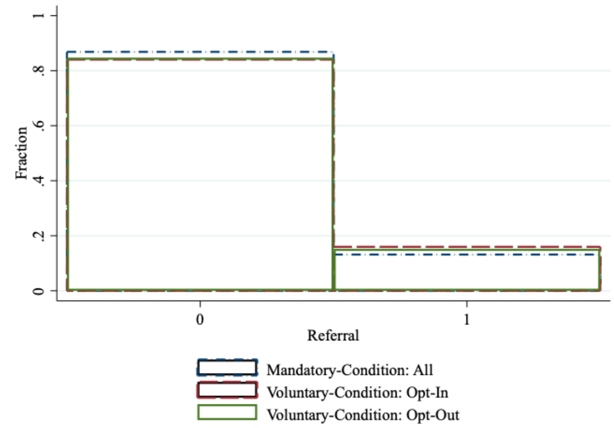

(c) Referral Status

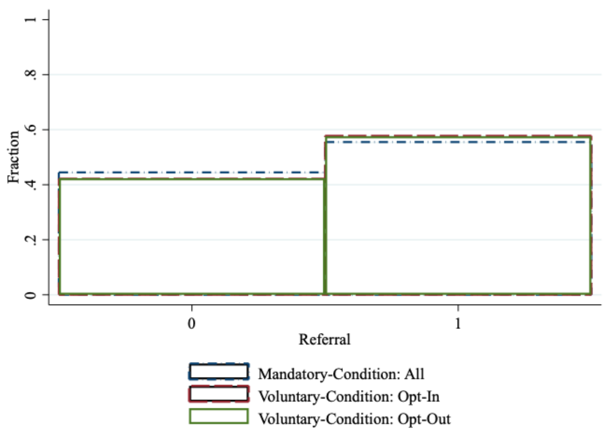

(d) Call Center Experience

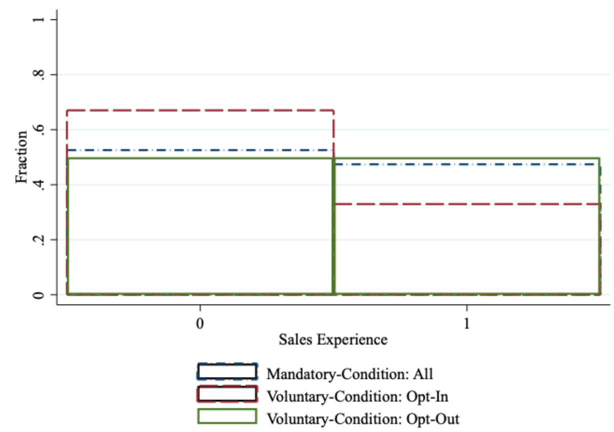

(e) Sales Experience

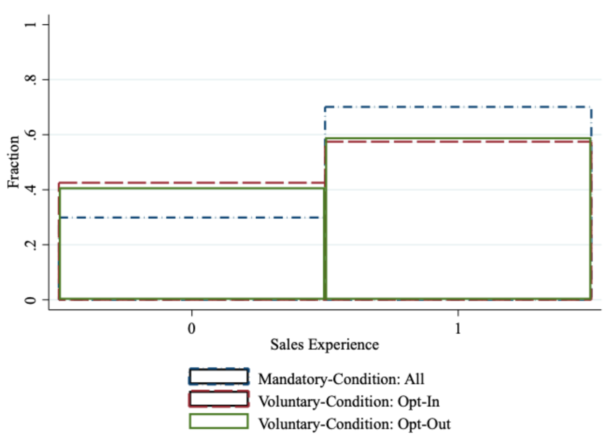

(f) Conversion

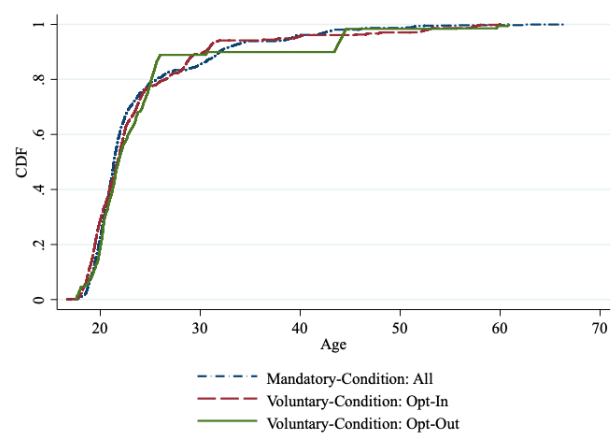

(g) Gender

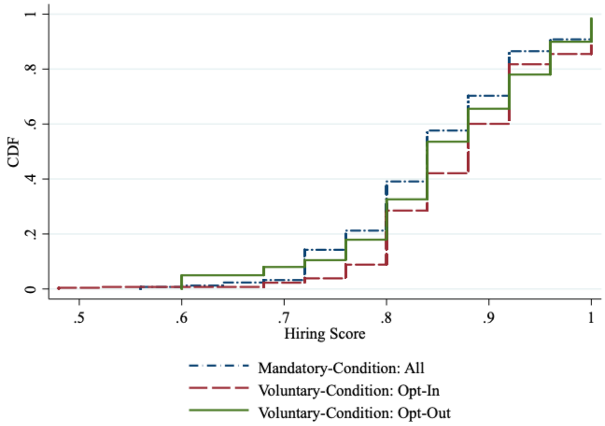

(h) Marital Status

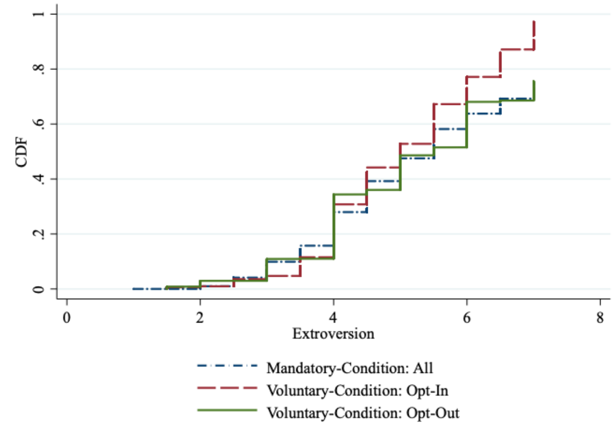

(i) Referral Status

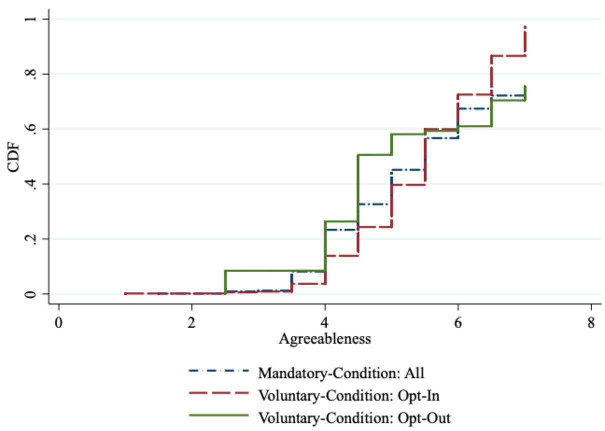

(j) Call Center Experience

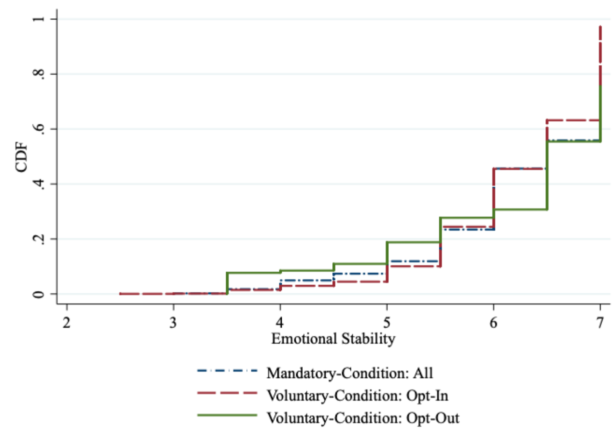

(k) Sales Experience

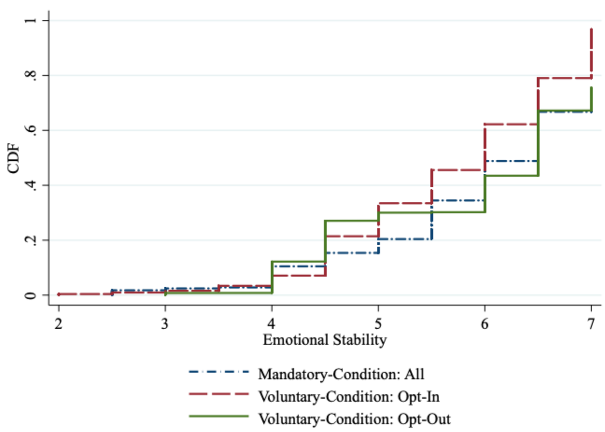

(l) Conversion

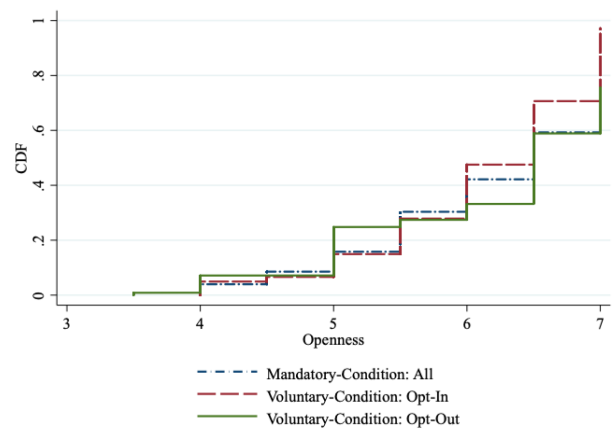

Notes: These figures plot comparisons of agent characteristics across three different groups: (1) all agents in the Mandatory-Condition; (2) agents in the Voluntary-Condition who opted into mentorship; (3) agents in the Voluntary-Condition who opted out of mentorship. For binary characteristics, we plot the fraction of agents in each group using histograms. For continuous characteristics, we plot the cumulative distribution function for each group.

Figure A.2: Effect of Mentoring on Productivity (Months 3–6 on Sales Floor)

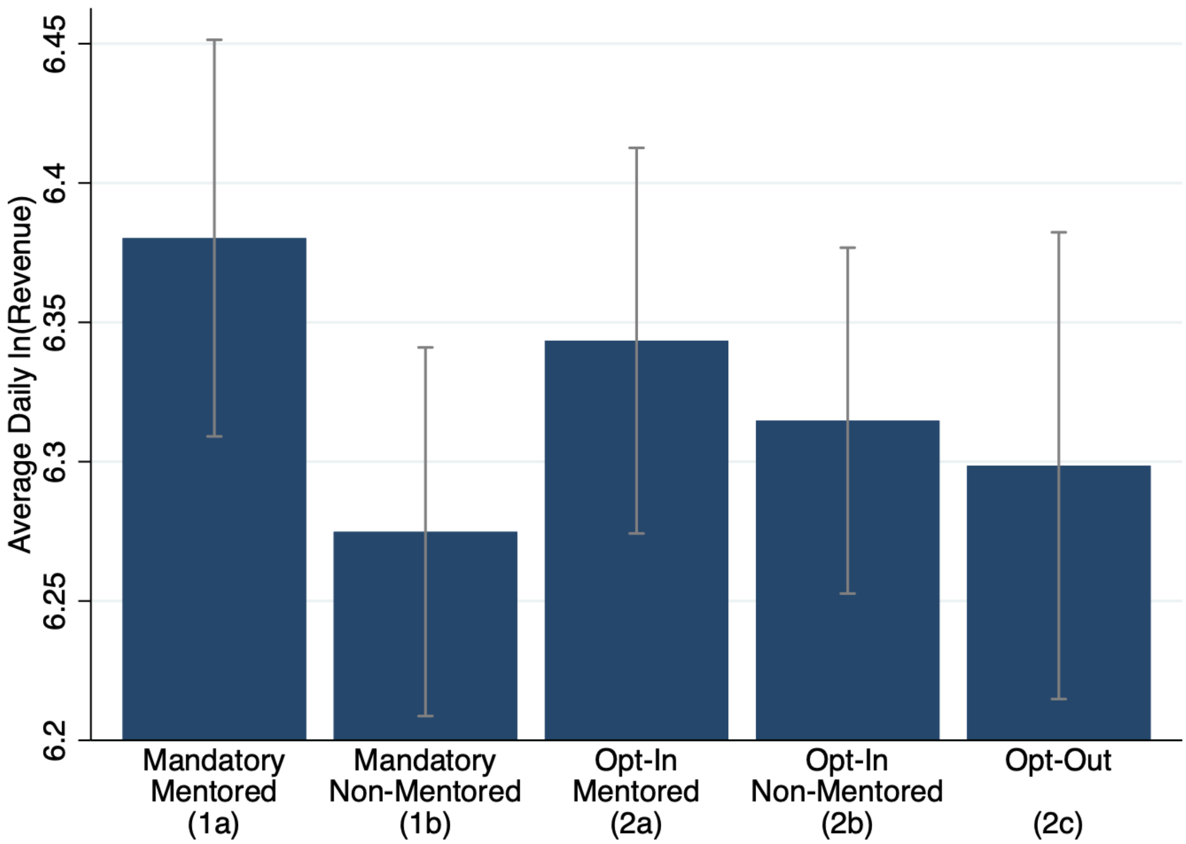

Note: This figure plots the average daily log revenue and 95% confidence intervals for agents' third to sixth months on the sales floor, split by treatment condition. Before we aggregate revenue amounts, we net out cohort fixed effects and then we add in the average productivity level across all agents that did not receive a mentor as a baseline.

Table A.1: Balance in Mentor Demographics

	Mandatory-Condition	Voluntary-Condition	p-value
	(1)	(2)	(2)−(1)
Mentor Age (yrs.)			
Mean	24.06	23.68	0.516
Std Dev.	4.66	4.44	
Mentor Woman			
Mean	0.28	0.22	0.278
Std Dev.	0.45	0.42	
Mentor Married			
Mean	0.08	0.20	0.010
Std Dev.	0.27	0.40	
Mentor Tenure			
Mean	1.32	1.21	0.390
Std Dev.	0.83	1.07	
Number of Protégés	114	123	

Notes. In this table we report average characteristics of the agents who mentored protégés in the Mandatory-Condition in Column (1) and of the agents who mentored protégés in the Voluntary-Condition in Column (2). The p-values from the difference-in-means tests are reported in the rightmost column. Mentors were not designated exclusively to either of the mentoring conditions. In other words, a mentor's first protégé could have been assigned to the Mandatory-Condition, whereas their second protégé could have been assigned to the Voluntary-Condition. Mentors were never informed as to whether their protégés were in the Mandatory-Condition or the Voluntary-Condition.

Table A.2: Long-Term Treatment and Selection Effects of Mentoring

	ln(Revenue)	ln(RPC)	ln(RPH)	Adherence	Tenure_3	Tenure_4
	(1)	(2)	(3)	(4)	(5)	(6)
Mentored	0.162	0.083	0.107	0.005	0.009	0.031
standard errors	(0.106)	(0.064)	(0.094)	(0.007)	(0.070)	(0.071)
sharpened q-value	[0.240]	[0.345]	[0.403]	[0.718]	[0.761]	[0.723]
Mentored × Voluntary	-0.126	-0.103	-0.084	0.004	-0.013	-0.019
standard errors	(0.135)	(0.084)	(0.114)	(0.010)	(0.095)	(0.097)
sharpened q-value	[0.463]	[0.359]	[0.638]	[0.723]	[0.761]	[0.761]
Voluntary Opt-Out	-0.043	0.013	0.023	-0.012	-0.252***	-0.212**
standard errors	(0.101)	(0.042)	(0.080)	(0.013)	(0.090)	(0.081)
sharpened q-value	[0.723]	[0.761]	[0.761]	[0.463]	[0.060]	[0.060]
Cohort FE	✓	✓	✓	✓	✓	✓
Adj. R-Square	0.049	0.058	0.058	0.083	0.007	0.003
Observations	13,075	13,075	13,075	13,075	603	603

Notes. The sample used in Columns (1)–(4) is composed of agent-day productivity data for all mentor-eligible agents with post-training productivity data. The data covers agents' productivity on their third to sixth months on the sales floor. *Mentored* equals one for agents who were randomized to received an available mentor, and zero otherwise, *Voluntary* equals one for agents in the Voluntary-Condition, and zero otherwise, and *Voluntary Opt-Out* equals one for agents who chose to opt out of possibly receiving a mentor, and zero otherwise. Columns (5)–(6) use data with a single observation per unique hired agent to capture retention effects. $Tenure_t$ equals one for agents who achieve at least t months of tenure at the firm, and zero otherwise. We estimate ordinary least squares regressions in all columns. All specifications include cohort fixed effects. Standard errors are clustered by cohort and are reported in parentheses. Sharpened q-values are presented in brackets, following Anderson (2008). *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A.3: Treatment and Selection Effects of Mentoring (Pre-May vs. Post-May)

Panel A: Productivity in Months 1–2

	ln(Revenue)	ln(RPC)	ln(RPH)	Adherence	$Tenure_1$	$Tenure_2$
	(1)	(2)	(3)	(4)	(5)	(6)
Mentored	0.197***	0.107***	0.121**	0.024**	0.121***	0.150*
	(0.069)	(0.034)	(0.055)	(0.012)	(0.029)	(0.086)
Mentored × Post	-0.096	-0.023	-0.059	-0.019	-0.063	-0.206
	(0.172)	(0.109)	(0.124)	(0.018)	(0.105)	(0.129)
Mentored × Voluntary	-0.256*	-0.172*	-0.174	-0.017	-0.061	-0.218**
	(0.144)	(0.100)	(0.128)	(0.017)	(0.076)	(0.106)
Mentored × Voluntary × Post	0.121	0.046	0.076	0.015	0.093	0.289*
	(0.228)	(0.152)	(0.184)	(0.023)	(0.134)	(0.161)
Voluntary Opt-Out	-0.270**	-0.206***	-0.228***	-0.011	0.013	-0.328***
	(0.109)	(0.061)	(0.077)	(0.025)	(0.029)	(0.075)
Voluntary Opt-Out × Post	-0.126	0.023	0.023	0.008	-0.170	0.303**
	(0.229)	(0.131)	(0.178)	(0.028)	(0.113)	(0.143)
Cohort FE	✓	✓	✓	✓	✓	✓
Adj. R-Square	0.029	0.046	0.028	0.118	0.073	0.074
Observations	15,137	15,137	15,137	15,137	603	603
No Differential Post Effects	0.864	0.993	0.969	0.569	0.455	0.074

Panel B: Productivity in Months 3–6

	ln(Revenue)	ln(RPC)	ln(RPH)	Adherence	$Tenure_3$	$Tenure_4$
	(1)	(2)	(3)	(4)	(5)	(6)
Mentored	0.129	0.037	0.064	0.011*	0.042	0.092
	(0.126)	(0.071)	(0.117)	(0.006)	(0.090)	(0.090)
Mentored × Post	0.129	0.177	0.165	-0.026	-0.110	-0.198
	(0.213)	(0.125)	(0.153)	(0.020)	(0.133)	(0.125)
Mentored × Voluntary	0.123	0.092	0.025	0.016	-0.009	0.047
	(0.150)	(0.088)	(0.136)	(0.012)	(0.125)	(0.109)
Mentored × Voluntary × Post	-0.462*	-0.405**	-0.270	-0.001	0.050	0.009
	(0.255)	(0.153)	(0.190)	(0.024)	(0.181)	(0.166)
Voluntary Opt-Out	0.052	0.101***	-0.009	0.007	-0.309**	-0.207**
	(0.111)	(0.034)	(0.056)	(0.012)	(0.133)	(0.085)
Voluntary Opt-Out × Post	-0.103	-0.124	0.111	-0.034	0.131	0.037
	(0.178)	(0.075)	(0.132)	(0.027)	(0.185)	(0.158)
Cohort FE	✓	✓	✓	✓	✓	✓
Adj. R-Square	0.050	0.059	0.059	0.086	0.005	0.007
Observations	13,075	13,075	13,075	13,075	603	603
No Differential Post Effects	0.121	0.042	0.325	0.091	0.590	0.115

Notes. The sample used in Columns (1)–(4) is composed of agent-day productivity data for all mentor-eligible agents with post-training productivity data. *Mentored* equals one for agents who received mentorship, and zero otherwise, *Voluntary* equals one for agents in the Voluntary-Condition, and zero otherwise, *Voluntary Opt-Out* equals one for agents who chose to opt out of possibly receiving a mentor, and zero otherwise, and *Post* equals one for cohorts that entered the firm after May 19th, and zero otherwise. We estimate ordinary least squares regressions in all columns. Columns (5)–(6) use data with a single observation per unique hired agent to capture retention effects. $Tenure_t$ equals one for agents who achieve at least t months of tenure at the firm, and zero otherwise. The bottom row reports p-values from post-estimation tests that the coefficients on *Mentored* × *Post*, *Mentored* × *Voluntary* × *Post*, and *Voluntary Opt-Out* × *Post* are jointly equal to zero. All specifications include cohort fixed effects. Standard errors are clustered by cohort and are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A.4: Treatment and Selection Effects of Mentoring on Calls and Hours

Panel A: Calls and Hours in Months 1–2

	ln(Calls)	ln(Hours)	ln(Calls/Hour)
	(1)	(2)	(3)
Mentored	0.051	0.050***	0.001
	(0.033)	(0.018)	(0.025)
Mentored × Voluntary	-0.037	-0.050**	0.013
	(0.043)	(0.022)	(0.037)
Voluntary Opt-Out	-0.074	-0.075**	0.001
	(0.047)	(0.036)	(0.029)
Cohort FE	✓	✓	✓
Adj. R-Square	0.124	0.082	0.177
Observations	15,137	15,137	15,137

Panel B: Calls and Hours in Months 3–6

	ln(Calls)	ln(Hours)	ln(Calls/Hour)
	(1)	(2)	(3)
Mentored	0.057	0.040*	0.018
	(0.046)	(0.023)	(0.037)
Mentored × Voluntary	-0.015	-0.041	0.026
	(0.062)	(0.043)	(0.052)
Voluntary Opt-Out	-0.050	-0.068	0.017
	(0.072)	(0.059)	(0.055)
Cohort FE	✓	✓	✓
Adj. R-Square	0.108	0.063	0.200
Observations	13,075	13,075	13,075

Notes. The sample used in Columns (1)–(3) is composed of agent-day productivity data for all mentor-eligible agents with post-training productivity data. *Mentored* equals one for agents who received mentorship, and zero otherwise, *Voluntary* equals one for agents in the Voluntary-Condition, and zero otherwise, *Voluntary Opt-Out* equals one for agents who chose to opt out of possibly receiving a mentor, and zero otherwise, We estimate ordinary least squares regressions in all columns. All specifications include cohort fixed effects. Standard errors are clustered by cohort and are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A.5: Lee Bounds Estimates of the Effect of Mentoring on Productivity

Panel A: Mandatory-Condition, Productivity in Months 1–2

	ln(Revenue)	ln(RPC)	ln(RPH)	Adherence
	(1)	(2)	(3)	(4)
Mentored_{lower}	0.084**	-0.002	-0.002	0.006*
	(0.040)	(0.028)	(0.031)	(0.003)
Mentored_{upper}	0.464***	0.273***	0.299***	0.037***
	(0.035)	(0.028)	(0.029)	(0.003)
Observations	7,819	7,819	7,819	7,819

Panel B: Mandatory-Condition, Productivity in Months 3–6

	ln(Revenue)	ln(RPC)	ln(RPH)	Adherence
	(1)	(2)	(3)	(4)
Mentored_{lower}	0.129***	0.040	0.048*	-0.006**
	(0.036)	(0.026)	(0.028)	(0.003)
Mentored_{upper}	0.355***	0.214***	0.229***	0.010***
	(0.034)	(0.025)	(0.027)	(0.002)
Observations	6,592	6,592	6,592	6,592

Panel C: Voluntary-Condition, Productivity in Months 1–2

	ln(Revenue)	ln(RPC)	ln(RPH)	Adherence
	(1)	(2)	(3)	(4)
Mentored_{lower}	-0.045	-0.048	-0.053	0.002
	(0.058)	(0.035)	(0.042)	(0.006)
Mentored_{upper}	-0.030	-0.036	-0.042	0.004
	(0.038)	(0.028)	(0.030)	(0.003)
Observations	7,756	7,756	7,756	7,756

Panel D: Voluntary-Condition, Productivity in Months 3–6

	ln(Revenue)	ln(RPC)	ln(RPH)	Adherence
	(1)	(2)	(3)	(4)
Mentored_{lower}	0.023	-0.037	-0.004	-0.005*
	(0.060)	(0.036)	(0.044)	(0.003)
Mentored_{upper}	0.190***	0.073***	0.123***	0.006***
	(0.037)	(0.026)	(0.029)	(0.002)
Observations	6,872	6,872	6,872	6,872

Notes. Panels A and B use agent-day productivity data for agents in the Mandatory-Condition. Panels C and D use agent-day productivity data for agents in the Voluntary-Condition, excluding those agents who opt out of the program. We estimate treatment effect bounds that account for non-random attrition as proposed by Lee (2009). *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A.6: Meeting Completion Rates

	Mandatory-Condition	Voluntary-Condition	p-value
	(1)	(2)	
Number of Agents	114	123	
At Least One Recorded Meeting	110	114	
No Recorded Meeting	4	9	
Number Recorded Meetings (avg.)	2.58	2.27	0.073
	(1.43)	(1.22)	
Meeting Completion Ratio (avg.)	0.83	0.72	0.009
	(0.29)	(0.34)	

Notes. In this table we report the mentor meeting completion details of protégés in the Mandatory-Condition and the Voluntary-Condition. "No Recorded Meeting" indicates that there is no record that the mentor-protégé pair ever met with one another. The "Meeting Completion Ratio" measure is based on the number of possible meetings the mentor-protégé pair could have had. While the protocol called for four weekly meetings, if either the mentor or protégé missed a week of work then that mentor-protégé pair would have only had three possible meetings.

Table A.7: Instrumental Variables Estimates of Meetings with Mentors Across Different Conditions

	Mandatory-Condition				Voluntary-Condition			
	First Stage	Second Stage	First Stage	Second Stage	First Stage	Second Stage	First Stage	Second Stage
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.836***		2.842***		0.746***		2.377***	
	(0.046)		(0.206)		(0.061)		(0.172)	
Meeting Completion Ratio		0.203**				-0.073		
		(0.080)				(0.091)		
Number Recorded Meetings				0.060**				-0.023
				(0.023)				(0.029)
Cohort Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
Cragg-Donald Wald F	50,058		19,338		32,493		17,991	
Centered R-Square	0.843	0.002	0.718	0.005	0.734	0.000	0.666	-0.000
Observations	6,744	6,744	6,744	6,744	7,035	7,035	7,035	7,035

Notes. In this table we estimate two-stage least squares regressions using *Mentored* as an instrumental variable for mentor-protégé meeting completion. Columns (1)–(4) consider only agents in the Mandatory-Condition, and Columns (5)–(8) consider only agents in the Voluntary-Condition who did not choose to opt out of the mentorship program. The dependent variable in Columns (1) and (5) is *Meeting Completion Ratio*, the fraction of possible (i.e., scheduled) mentor-protégé meetings that the protégé completed. The dependent variable in Columns (3) and (7) is *Number Recorded Meetings*, the number of mentor-protégé meetings that the protégé completed. Then in Columns (2), (4), (6), and (8), we regress log revenue on the predicted values of these meeting completion numbers. All specifications include cohort fixed effects. Standard errors are clustered by cohort and are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A.8: Differences in Worksheet Content

	Total Words per Worksheet	Skill Words per Worksheet	Support Words per Worksheet	Other Words per Worksheet
	(1)	(2)	(3)	(4)
Mandatory-Condition	-0.023	0.253	0.145*	-0.422
	(2.686)	(0.311)	(0.079)	(2.584)
Adj. R-Square	-0.006	-0.002	0.013	-0.006
Observations	159	159	159	159
Mean DV	47.36	4.27	0.448	42.65

Notes. This table considers differences in worksheet content between protégés in the Mandatory-Condition and those in the Voluntary-Condition. For each worksheet, we identify the fraction of words in the responses that relate to job-specific skills or knowledge (*Skill*), those that relate to receiving support, encouragement, and friendship (*Support*), and those that are neither related to skill nor support (*Other*), which include stop words. These become the dependent variables in our regression specifications of worksheet content on mentorship type and opt-out likelihood. Robust standard errors are reported in parentheses. The mean of the dependent variable is listed below the observation count line in each panel. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

B Documentation

Mentor Instructions

What is a Mentor?

In *The Odyssey*, Odysseus prepared to fight in the Trojan War. Before leaving home to fight in the war, he asked his trustworthy friend, named Mentor, to train and educate his son, Telemachus. Similarly, mentors today are meant to train and educate their protégés. Management at has chosen you to be a mentor---a source of further skill development---for newly hired sales agents. You have been selected specifically because you've demonstrated a willingness to teach other sales agents and help them become a successful and productive sales agent.

The responsibility to mentor a newly hired sales agent should not be taken lightly. Management strongly believes new agents will benefit from the additional training and the insider knowledge received as a result of being mentored by a talented, more seasoned agent. Because of this, has devoted significant resources to give mentors and protégés the best opportunity to spend productive time together, so please take your mentorship responsibilities seriously.

What will You Do as a Mentor?

As a mentor, you will do the following:

1. You will meet with your protégé at **least once a week**.
    a. Before meeting, your protégé will complete the Protégé Worksheet.
        i. If he/she has not completed it, you will kindly help him/her do so.
    b. During your meeting, you and your protégé will discuss his/her responses. You should also take this time to do the following:
        i. Impart knowledge and skill by explaining, giving useful examples, and demonstrating processes, and asking thought-provoking questions.
        ii. Discuss actions you've taken to become a successful sales agent.
        iii. Provide him/her with any tips and sales tactics that help you overcome customer concerns and that help you up-sell to better services.
        iv. Practice the designated sales protocol with them and help them gain a strong understand of the products, services, and bundles available.

2. After meeting with your protégé, you will deliver the finished worksheet to .
    a. will initial and timestamp the worksheet and make a record that you completed your weekly meeting responsibility.

3. Every two weeks, you will be asked to complete an on-line survey.
    a. These questions are meant to gauge the progress of your protégé and the overall benefit of the mentoring relationship.
    b. Please answer these questions honestly, as they are not meant to punish but, instead, to help assess the effectiveness of the mentorship program.

Protégé Worksheet (Week 1)

Protégé: ____
Mentor: ________________________________ Number of times mentor has reached out: _____
Date: ____

Weekly Self-Reflection:

What are **your expectations** regarding your sales ability? Does your mentor know this?

What may **prevent you from** having a successful first week? Does your mentor know this?

Think of the **MOST** successful call you had recently. What made it **successful**?

Think of the **LEAST** successful call you had recently. What made it **unsuccessful**?

Weekly Goal:
What **ONE goal** are you setting for yourself for this coming week?

What will you do to **reach this goal**? Have you told your mentor about this goal? __________
____

For Mentors to Respond:
How have you, as a mentor, been a source of **skill development** for your protégé? What have you done so far to help him/her succeed on the sales floor here at ?

Protégé's Initials	Mentor's Initials	Intern's Initials & Timestamp

Protégé Worksheet (Week 2)

Protégé: ____
Mentor: _______________________________________ Number of times mentor has reached out: __
Date: ____

Weekly Self-Reflection:
What was your **big win** from Week 1? Does your mentor know this yet?

Do you feel **comfortable reaching out** to your mentor? coach? manager? Why or why not?

Do you feel you have been **provided the tools** you need to be successful? Why or why not?

Think of the **MOST** successful call you had recently. What made it **successful**?

Think of the **LEAST** successful call you had recently. What made it **unsuccessful**?

Weekly Goal:
What **ONE goal** are you setting for yourself for this coming week?

What will you do to **reach this goal**? Have you told your mentor about this goal? ___________
____

For Mentors to Respond:
How have you, as a mentor, been a source of **skill development** for your protégé? What have you done so far to help him/her succeed on the sales floor here at ?

Protégé's Initials	Mentor's Initials	Intern's Initials & Timestamp

Protégé Worksheet (Week 3)

Protégé: ____

Mentor: _________________________________ Number of times mentor has reached out: ____

Date: ____

Weekly Self-Reflection:

What was your **big win** from Week 2? Does your mentor know this yet?

How are you **adjusting to your work environment**? What does your team **expect of you**?

Do you understand your **department's compliance rules**? Why or why not?

Think of the **MOST** successful call you had recently. What made it **successful**?

Think of the **LEAST** successful call you had recently. What made it **unsuccessful**?

Weekly Goal:

What **ONE goal** are you setting for yourself for this coming week?

What will you do to **reach this goal**? Have you told your mentor about this goal? __________

____

For Mentors to Respond:

How have you, as a mentor, been a source of **skill development** for your protégé? What have you done so far to help him/her succeed on the sales floor here at ?

Protégé's Initials	Mentor's Initials	Intern's Initials & Timestamp

Protégé Worksheet (Week 4)

Protégé: ____

Mentor: ____________________________________ Number of times mentor has reached out: _____

Date: ____

Weekly Self-Reflection:

What was your **big win** from Week 3? Does your mentor know this yet?

What is **something new you have learned** since being on the phones?

Do you understand **'s attendance policy**? What **motivates you** to work hard?

Think of the **MOST** successful call you had recently. What made it **successful**?

Think of the **LEAST** successful call you had recently. What made it **unsuccessful**?

Weekly Goal:

What **ONE goal** are you setting for yourself for this coming week?

What will you do to **reach this goal**? Have you told your mentor about this goal? ___________

____

For Mentors to Respond:

How have you, as a mentor, been a source of **skill development** for your protégé? What have you done so far to help him/her succeed on the sales floor here at ?

Protégé's Initials	Mentor's Initials	Intern's Initials & Timestamp

C Worksheet Response Examples

Panel A: Think of the most successful call you had recently. What made it successful?

Skill	· I pitched TV really well · Having different examples of pitches from my coach to fall back on
Support	· I was confident and tried to connect · The person I spoke with was very nice

Panel B: Think of the least successful call you had recently. What made it unsuccessful?

Skill	· Customer didn't want to pay the deposit, [I] didn't rebuttal · Not doing call flow, not caring, not enough discover
Support	· Not being confident in my ability to rebuttal · The person was rude and wanted me fired

Panel C: What will you do to reach this goal? Have you told your mentor about this goal?

Skill	· I'll follow the call flow · I will create better pitches · Be better with the triple play, use what [the] mentor told [me] · My mentor is going to help me pitch DTV by giving me her tips on what helped her · Practice on every unserviceable call · Try upsell technique
Support	· [Goal to achieve] 1500 a day, build confidence in it · Be more positive · Stay positive · Stay in communication with [my coach] · Motivations - self discipline · Check in with my coach and be confident

Panel A: Words Associated with Sales Skills and Knowledge

Adherence, Conversion, Customer, Direct, Dish, Double, DPI, DTV, Internet, Knowledge, Phone, Pitch, Price, Pricing, Process, Revenue, RPC, RPH, Sale, Security, Sell, Skill, Sold, System, Television, Triple, TV

Panel B: Words Associated with Receiving Support

Annoy, Breath, Confidence, Confident, Cool, Encourage, Encouraging, Friend, Introduce, Kind, Laugh, Mean, Motivate, Motivation, Nice, Patience, Patient, Positive, Rude, Social, Support, Welcome, Welcoming

D AEA Pre-Registration Text

Here we replicate the AEA pre-registration text. Differences between the AEA pre-registration and our actual implementation are denoted in footnotes.

D.1 Abstract

Mentoring is increasingly encouraged in workplaces, and a number of firms have implemented formal programs. While a growing body of research suggests that mentoring relationships benefit those being mentored (protégés), there is scant evidence to delineate whether these favorable outcomes are driven by the mentoring experience on average, by the self-selection of protégés into mentoring who anticipate having the largest gains (selection based on gains), or by the self-selection of protégés who would have performed well in the absence of mentoring (selection based on levels). We use a field experiment to evaluate a workplace mentoring program inside a large sales organization.

Experienced employees opt-in as mentors, and new hires are slated as potential protégés. The project objective is to study the mentoring consequences across protégés who actively elect to be formally mentored relative to those who are randomly allocated a mentor. We estimate treatment effects on sales productivity and turnover for those who select into mentoring and for those who opt out.

D.2 Intervention(s)

We analyze the effectiveness of a workplace mentoring program where employees opt-into mentoring or are randomly assigned a mentor. More details are provided in the design field.

D.2.1 Intervention Start Date

2019-05-27

D.2.2 Intervention End Date

2019-12-20

D.3 Primary Outcomes (end points)

Log revenue-per-call (RPC), an indicator for worker turnover, log completed tenure, the firm's internal adherence to schedule measure (e.g. time spent working whilst at work), and the firm's internal engagement metrics (online surveys asking for willingness to recommend employment at firm, comfort with leadership, etc.).

D.3.1 Primary Outcomes (explanation)

Agent's weekly RPC is a measure of sales productivity that removes demand variation outside of the worker's control. RPC is the primary productivity measure used by the firm, combining both agent's firm-specific knowledge and their individual effort.[29] Worker turnover measures whether

[29] RPC was the primary endpoint based on our experience analyzing the productivity of veteran agents within the firm (Sandvik et al., 2020), but total revenue picks up different margins of adjustment for new agents, which is why we report both metrics, along with revenue-per-hour (RPH).

the interventions changed the agents' propensity to leave the firm. Log of completed tenure is a different measure of retention that has been used in the prior literature and the attendance measure provides an adjacent measure of agent effort. Finally, engagement measures are hypothesized to be forward looking measures of productivity.

D.4 Experimental Design Details

Seasoned sales agents are invited to apply as internal mentors to incoming recruits (the firm "qualifies" mentors as having sufficient sales experience). New mentorship opportunities are periodically announced, and prior mentors are permitted to re-enter the mentor pool. The firm communicates that serving as a mentor is a useful first step to being considered for a managerial position. New mentors complete a survey asking them about their personality, interests, work preferences, and values. Mentors are randomly assigned with probability 50% to receive a set of instructions emphasizing that mentoring is about teaching protégés how to do the job. The remaining mentors receive instructions emphasizing that mentoring is about providing protégés support. Sales agents are hired in batches (cohorts). Newly hired sales agents complete two weeks of training, primarily in a classroom or listening in on other agents' sales calls. New agents then complete the same personality and preference survey that mentors take. At the end of their two-week training, each cohort of agents is eligible for randomization into a mentoring treatment arm. Any mentoring relationship commences as soon as the agent completes their training.

The randomization procedure is as follows:

D.4.1 Cohort Level Randomization

The initial level of randomization is cohorts of new hires (potential protégés). Each cohort (a group of new hires who are joining the firm at the same time, are in the same training group, and will be working in the same sales division and office location) will be randomized into one of two conditions: Mandatory-Condition or Voluntary-Condition. 40% of the cohorts will be in the Mandatory-Condition group and 60% of the cohorts will be in the Voluntary-Condition group.

D.4.2 Within Cohort Randomization

For cohorts in Mandatory-Condition, new hires will receive a mentor with probability 50%. This will be communicated privately between sales floor staff and the individual workers. Agents in the Mandatory-Condition who do not receive a mentor will not receive communication regarding the program. For cohorts in Voluntary-Condition, sales floor staff verbally explain the firm's mentorship program, answer questions, and provide each agent a confidential ballot where they can decide whether or not to enter a lottery which randomly determines whether the agent is allocated a randomly assigned mentor, or no mentor at all. Of the agents who enter the lottery, approximately 50% will be assigned a mentor. Agents who choose not to be mentored will never be assigned a mentor.

D.5 Compliance Tracking

The firm's training staff will track whether mentors and protégés meet. This tracking will be aided by worksheets. Upon completion of the worksheets, the firm will reward "kudos" points that can be accumulated to purchase items from the company store. As mentioned earlier, mentors may participate more than once, however they will never have more than one protégé at a time.[30] Eligible protégés and mentors will each take an electronic survey at the end of the formal program. The survey for protégés will ask about the protégé's initial excitement when told about the mentoring program, their perceived engagement with their mentor, and an estimate of the effectiveness of mentoring. This question will be phrased as: "What was your average RPC last week? What do you think your average RPC would have been had you not been working with a mentor?" The survey for mentors will ask about the protégé's enthusiasm for the mentorship program and an estimate of the mentor's perceived treatment effect on the protégé. This question will be phrased as: "If your protégé had not received mentoring, his/her RPC would have been [40% lower — slider — 40% higher].[31] Note that numbers greater than zero mean that mentoring was not effective for improving protégé performance. Please be candid, as your responses will not be shared with management."

D.6 Edit June 4, 2019

To assess the potential for spillovers, we have revisited the design in consultation with the company such that there will be "hold out" cohorts for one division-office who never receive mentoring. Any cohorts/individuals who are switching brands also will be held-out. Work-from-Home cohorts will also present a possible "hold out" group for comparison and all cohorts in a smaller third office (which no longer exist, but for whom historical data is available) were "hold out" cohorts who knew nothing about mentoring. A "sentiment survey" will be administered to all agents in their 5th week on the sales floor.[32] This will be one week after mentored agents finish hiring. We will gather information on their feelings towards the onbaording process and ask questions, common in the literature, to solicit their sentiment towards the firm, their perceptions of their ability, their enthusiasm about the job, etc. We will use this survey to test for spillovers based on survey responses.

D.7 Randomization Method

Randomization done by computer. Participants will be informed if randomized in.

D.8 Randomization Unit

Clustered randomization of cohorts in a first level, with individual randomization within the cohort. See design details.

[30] As the program progressed, the internal mentoring staff felt that many of the mentors could effectively mentor multiple protégés as once. As a result, we adjusted the protocol such that it was possible for a single individual to mentor multiple new hires concurrently, but mentor-protégé pairs always met individually, meaning the protocol was the exact same from the point of view of the protégé.

[31] The wrap-up survey completion rates of mentors and protégés were very poor, so we do not have meaningful data for this question. Anecdotally, the average responses of both sets of individuals suggests that protégés' RPC would have been lower in the absence of mentorship, but the inference is not precise.

[32] We were not able to administer this survey. The firm had several of its own survey initiatives occurring simultaneously, so additional surveys connected to the mentorship program were not conducted due to the concern of "survey fatigue" among the sales agents.

D.9 Was the treatment clustered?

Yes

D.10 Sample size: planned number of clusters

The exact sample size is stochastic and depends on the firm's actual hiring. We have 46 planned clusters.

D.11 Sample size: planned number of observations

In one office, the firm has projected 269 new hires in 22 cohorts. There are 350 new hires in 24 cohorts projected in the second office.

D.12 Sample size (or number of clusters) by treatment arms

Please see design field.

D.13 Minimum detectable effect size for main outcomes

Using pre-intervention data to estimate the intra-class correlation coefficient and residual variation, the minimum detectable effect size for log RPC between those randomized into and out of mentoring is 0.07 (accounting for sample design and clustering).

D.14 Analysis Plan

The Treatment Effect of Mentoring on those who opt in is:

$$\beta_{OptInMentor} = mean(Y_{OptInMentor}) - mean(Y_{OptInNoMentor}).$$

We will estimate this mean difference using a regression of Y on an indicator for receiving a mentor along with cohort fixed effects and indicators for the type of instructions mentors receive.[33] The sample will be the workers in the voluntary treatment cohorts who opt into mentoring.

The Treatment Effect of Mentoring on those who opt out can then be derived by writing the average gain from mentoring in the population as:

$$mean(Y_{RandomMentor}) - mean(Y_{NoMentor}) = \beta_{OptInMentor}\pi_{OptIn} + \beta_{OptOutMentor}\pi_{OptOut}.$$

The β parameters are the heterogeneous treatment effects and the π are the population fraction who opt in and opt out. This yields:

$$\beta_{OptOutMentor} = [mean(Y_{RandomMentor}) - mean(Y_{NoMentor}) - \beta_{OptInMentor}\pi_{OptIn}]/\pi_{OptOut},$$

where the difference in means is net of cohort fixed effects and indicators for mentoring instruction type. The population average treatment effect (ATE) of mentoring can be estimated from a

[33] As described in Section 7.1, differences in mentor instructions did not have a meaningful impact on the treatment effect of mentorship. Because of this and for brevity, we omit this indicator from the models in our heterogeneous treatment effects tests.

regression of Y on a dummy for receiving a mentor and cohort fixed effects in cohorts that have (entirely) randomly assigned mentoring. This yields:

$$\beta_{OptOutMentor} = [ATE - \beta_{OptInMentor}\pi_{OptIn}]/\pi_{OptOut}.$$

Inference for $\beta_{OptOutMentor}$ will come from block bootstrapping the statistic. Selection bias will be measured among voluntary treatment cohorts as the regression analogue of:

$$mean(Y_{OptInNoMentor} - mean(Y_{OptOut}),$$

where the means are net of cohort fixed effects. This procedure allows us to estimate sales productivity differences among proteges who opt into mentoring and those who do not. We use the sample of agents in the voluntary cohorts who did not receive a mentor. We regress Y on an indicator that the agent opted into mentoring along with cohort fixed effects and their mentor instruction-type fixed effects. Other regressions will look at opt-in as a function of early sales and demographic characteristics (gender, age, office location) and past experience (prior sales or call center experience).

We plan to validate these estimates using the electronic survey responses collected after the protégé graduates from the formal mentoring program, approximately 4 weeks following the initial onboarding instruction (e.g. how to use the systems, enroll for benefits, etc.).[34] We will compare average perceived gains from mentors and protégés to the actual estimated treatment effects across different assignment conditions. We will then assess whether the effectiveness of the mentoring pair differs based on characteristics of the mentor and protégé. We will regress protégé sales on fully saturated interactions of demographic characteristics for the mentor-protégé pair (old/young based on coarse buckets; gender) as well as similarity in survey responses on the intake survey.

Finally, to assess whether mentoring detracts from–or improves sales–for the mentor, we will regress mentor log RPC and other sales measures on indicators demarking whether the mentor is eligible to mentor but has not yet done so, whether they have previously mentored in the program, or whether they are actively mentoring a protege. This regression will include mentor fixed effects and mentor tenure.[35]

[34] As mentioned earlier, we were not able to administer this survey.

[35] Tests that compare the characteristics of mentors and proteges, and those that look at the impact of mentorship on mentor productivity, are likely to be discussed in a separate article.