HARVARD | BUSINESS | SCHOOL

Employee Responses to Compensation Changes: Evidence from a Sales Firm

Jason Sandvik
Richard Saouma
Nathan Seegert
Christopher Stanton

Working Paper 18-100

Employee Responses to Compensation Changes: Evidence from a Sales Firm

Jason Sandvik
University of Utah

Richard Saouma
Michigan State University

Nathan Seegert
University of Utah

Christopher Stanton
Harvard Business School

Working Paper 18-100

Funding for this research was provided in part by Harvard Business School.

Employee Responses to Compensation Changes: Evidence from a Sales Firm

Jason Sandvik[*] Richard Saouma [†] Nathan Seegert[‡]

Christopher Stanton[§]

July 19, 2020

Abstract

Using data from an inbound sales call center, we study employee responses to compensation changes that ultimately reduced take-home pay by 7% for the average affected worker. These changes caused a significant increase in the turnover rate of the firm's most productive employees, but the response was relatively muted for less productive workers. We quantify the cost of losing highly productive employees and find that their relatively high sensitivity to changes in compensation limits managerial flexibility to adjust incentives. For the workers who remained at the firm, the compensation changes had minimal effects on performance. Our results speak to the possible sources of compensation rigidity and the difficulty managers face when setting compensation.

We thank Jen Brown, Lauren Cohen, Jeff Coles, Guido Friebel, Peter Kuhn, Bentley Macleod, Ramana Nanda, Paige Ouimet, Luke Stein, Ed Van Wesep, and seminar participants at the Arizona State Meeting of the Labor and Finance Group, Harvard Business School, SIOE Montreal, and UC Santa Barbara for helpful feedback.

[*]Department of Finance, David Eccles School of Business, University of Utah
[†]Eli Broad School of Business, Michigan State University
[‡]Department of Finance, David Eccles School of Business, University of Utah
[§]Harvard Business School and NBER; Corresponding Author: cstanton@hbs.edu

1 Introduction

How will full-time employees respond to unanticipated, adverse compensation changes? Will highly productive workers respond differently than their less productive peers? And can firm performance be substantially impacted by these employee responses? When tasked with adjusting compensation, managers must balance incentive and retention effects with expense reductions, all while limiting the damage they cause to implicit relational contracts. Given this difficulty and the uncertain responses of workers, managers largely avoid imposing adverse compensation changes.[1]

The literature has surfaced myriad negative worker responses to explain managers' reluctance to reduce compensation. The documented responses in short-term employment settings include fairness concerns (Fehr, Goette, and Zehnder, 2009) and social preferences such as warm glow and social norms (DellaVigna, List, Malmendier, and Rao, 2016). Previous studies have shown these effects can lead to output reductions (Kube, Maréchal, and Puppe, 2013) and mid-project attrition (Chen and Horton, 2016). Studies of long-term employment relationships have emphasized on-the-job responses (e.g., total output). For instance, Mas (2006) finds evidence of shirking when pay raises are deemed insufficient, and Krueger and Mas (2004) find evidence of decreased production quality among workers who perceive their compensation to be unfair. Other work documents increases in theft and antisocial behavior following unanticipated pay cuts (Greenberg, 1990; Giacalone and Greenberg, 1997).

The research most closely related to ours emphasizes the importance of management practices that retain top talent. For example, Zenger (1992) finds that employment contracts that disproportionately reward high performers contribute to their retention. In a study of lawyers, Campbell, Ganco, Franco, and Agarwal (2012b) show that high performers (as measured by their earnings) are less likely to turnover at baseline, and other research finds that firms with unequal pay structures that tilt compensation toward high performers show even less turnover at the top (Carnahan, Agarwal, and Campbell, 2012). These findings suggest that the need to retain high performers influences compensation structure but that the magnitude and type of influence

[1] Even in the face of falling demand, managers tend to favor layoffs over compensation reductions, contributing to well-documented instances of downward nominal wage rigidity (Kahn, 1997; Altonji and Devereux, 2000; Fehr and Goette, 2005; Holden and Wulfsberg, 2009; Kaur, 2014).

depends on the open question of how sensitive are the best performers to compensation *changes*. We complement prior research by linking compensation changes to employee turnover and effort across the performance distribution in a setting with long-term employment relationships—something few other studies have been able to observe.

The constraints managers face when changing compensation depend on how their most productive workers respond.[2] The underlying issue is that within-firm performance differences across workers can be significant (Lazear, 2000; Mas and Moretti, 2009; Lazear, Shaw, and Stanton, 2015; Sandvik, Saouma, Seegert, and Stanton, 2020). If the most productive workers are also the most responsive to adverse compensation changes, then managers may be more constrained than what an analysis of average responses would suggest. Survey responses of managers buttress the need to unpack responses over the productivity distribution, as Campbell III and Kamlani (1997) document that managers' concerns over losing their most productive employees frequently limit compensation reductions. This sentiment is perhaps best captured by Bewley (1998), who conducted over 300 interviews to understand why firms were so reluctant to cut pay, even in the face of falling customer demand. Bewley states: "[turnover] among the better workers is especially feared, because they are more valuable and can find new jobs more easily."

We show that the cited concerns of managers are consistent with the responses of an organization's highly productive workers; they quit in response to implicit incentive reductions. Our empirical setting is a US-based inbound sales call center. Workers belong to six different divisions, each of which answers calls from prospective customers interested in different types of digital services (e.g., television, phone, internet, etc.). Upon answering a call, workers assess a customer's needs, map these needs to available offerings, and negotiate with the customer to upsell them to high quality—more expensive—products. The presidents of one of the six divisions (henceforth Division 1) independently decided to re-balance the division's commission schedule, reducing commissions on the most commonly sold products and services and increasing them for more expensive and difficult to sell products. The changes were expected to decrease overall

[2] If firms could costlessly set individual compensation contracts for each employee (e.g., personalized commission schedules), then heterogeneous responses to compensation changes would be less of a concern.

commission pay by approximately 18% if the workers continued selling the same product mix as before. These changes to the commission schedule were unprecedented and were implemented because, several months earlier, Division 1 had expanded its call territory, causing its agents to earn more than comparable divisions. After the commission schedule changes, the workers in Division 1 could not adjust their mix of products sold, and they experienced the full expected reduction in commissions as a result. The take-home pay for the average worker fell by more than 7%. Three months after the changes in Division 1, the presidents of a second division, Division 2, enacted similar changes, which led to a decrease in average take-home pay of over 14%.

To study heterogeneous responses across the performance distribution of workers, we use worker-level output data given to us by the firm to estimate individual workers' sales productivity prior to the compensation changes. Individual productivity is widely dispersed, such that workers at the 75th percentile of the distribution sell about 50% more on a given call than those at the 25th percentile. This large dispersion motivates our investigation of the turnover and effort responses across the distribution of worker productivity.

We use three empirical approaches to estimate worker responses. We begin with a traditional difference-in-differences estimation, where we compare workers in Division 1, before and after the compensation changes, to workers in untreated control divisions. Importantly, about two months before the first compensation changes, Division 1 and the control divisions satisfied the common-trends assumption across observable outcomes. One might be concerned that our results could be biased by the territory expansion of Division 1, which motivated the compensation changes. Our second empirical strategy mitigates this concern, as it estimates heterogeneous responses for agents of different productivity levels *within* the same division. Our third approach uses survey responses to complement our main results and illuminate the possible mechanisms at play.

Our main findings point to the importance of considering both average employee responses and heterogeneous responses by agent productivity when evaluating the consequences of compensation changes. The turnover rate of highly productive agents in Division 1 increased significantly following the compensation changes. Specifically, workers with pre-treatment productivity that was one standard deviation above the mean had between a 40%–56% increase in attrition, relative to the baseline turnover rate. The average attrition rates of workers in Division 1 did not change, however, as less productive workers decreased their propensity to leave the firm. The loss of human capital from highly productive workers—who contribute significantly more to revenue than their peers—had substantial consequences for the firm. Despite initial savings on compensation expenses, the loss of highly productive agents reduced the firm's operating performance and led to a negative estimated net present value.

Turnover is rarely instantaneous, and it only happened with a lag in our setting. It took a little over four months for the cumulative loss in sales to outweigh the savings from the reduction in commission payments that resulted from the changes. Of equal importance, we observed virtually no abnormal attrition in the six weeks immediately following the compensation changes—highlighting the fact that workers did not respond to the announcement of the compensation changes by quitting *immediately*. The horizon of the turnover effects was uncertain ex-ante, and the lag in attrition is important for understanding why management initially felt that worker responses to the compensation changes were insignificant.[3]

The lack of immediate deleterious responses prompted division presidents to adjust the commission schedule in Division 2 three months after implementing the changes in Division 1.[4] This second change allows us to test whether the findings from Division 1 generalize. The average agent's commission pay in Division 2 fell by 27%, causing average take-home pay to fall by 14.5%. Unlike in Division 1, where average attrition remained relatively flat, the average attrition rate among agents in Division 2 more than doubled. The average agent in Division 2 performed—and was compensated—about as well as the highly productive agents in Division 1.[5] As such, the overall turnover effect in Division 2 is consistent with the heterogeneous turnover effect in Division 1. The observed attrition of highly productive workers from both treated divisions may

[3] A growing body of evidence suggests that experimenters and managers may opportunistically terminate experiments as soon as the desired result is attained (Berman, Pekelis, Scott, and Van den Bulte, 2018). In more ambiguous settings, where the underlying mechanisms are unknown, managers may inadvertently end interventions too early, before long-term effects are realized.

[4] Agents in Division 2 are excluded from the control sample used to estimate the effects in Division 1. No other divisions experienced compensation changes during the sample period.

[5] Agents in Division 2 sell products in large quantities to small businesses, which generates high gross profits for the firm and high commissions for the agents. Agents in Division 1 and the control divisions, on the other hand, sell products to residential customers, which generates smaller gross profits for the firm and lower commissions for the workers. The commission level for the average agent in Division 2 is within 10% of the average commission level for the top tercile of agents in Division 1.

explain the general reluctance of managers to alter their workers' compensation.

We find minimal evidence that agents responded to the compensation changes by adjusting their effort. If anything, our difference-in-differences estimates and survey response data suggest that Division 1 agents may have tried to increase their effort to offset some of the income lost due to the changes in their commission schedule. At first glance, this finding appears inconsistent with basic agency theory in static settings (Jensen and Meckling, 1976; Hölmstrom, 1979) as well as with more recent behavioral theories.[6] However, in long-term employment relationships, workers' responses are complicated by income effects, where the desire to offset a portion of lost earnings may ameliorate their desire to reduce effort in response to lower-powered incentives (Ashenfelter and Heckman, 1974; Stafford, 2015). Similarly, workers' responses may be complicated because their lower total compensation reduces their wealth. The Hicksian tradition distinguishes income (wealth) effects and price effects. Here the income effect, that workers' total compensation went down, may partially offset the decreased pay for each unit sold.

The two most likely mechanisms for the observed turnover effects are reductions in employee sentiment (Cohn, Fehr, and Goette, 2014a) and the superior outside options of highly productive workers (Bewley, 1998). Using survey responses, we find little predictive power for sentiment to explain the changes in turnover and effort. While we do find evidence of reduced sentiment after the compensation changes, workers' effort responses may have been damped by the strong incentive pay in our setting. In other settings, sentiment-based explanations may matter more. In contrast, anecdotal evidence suggests that highly productive workers did have attractive outside options. Labor shortages were common at the time of the study (Binyamin, 2017), and there were several competitors in the area vying for talented sales agents.

The increased turnover rate of highly productive agents in our setting aligns with some of the findings of Krueger and Friebel (2018), who study a reduction in incentive pay at a personnel search firm. Their study firm increased fixed wages to offset part of the reduction in incentive pay,

[6] Behavioral theories emphasize (wage) fairness concerns (Fehr et al., 2009), social comparisons (Larkin, Pierce, and Gino, 2012; Cohn, Fehr, Herrmann, and Schneider, 2014b; Obloj and Zenger, 2017), and negative reciprocity (Fehr and Falk, 1999; Dickson and Fongoni, 2018), all of which become especially salient when employees choose their effort provision following an adverse compensation change. Piece rate contracts or contracts with commissions may have different effects than adjustments to fixed wages (Esteves-Sorenson, 2017).

potentially alleviating income effects and giving rise to the effort changes they document. Several other studies have considered the ability to attract or retain different types of workers. Larkin and Leider (2012) examines how different menus of incentive schemes lead to the selection of sales workers based on their confidence, suggesting convexity in pay helps select highly motivated employees. Carnahan, Kryscynski, and Olson (2017) consider the turnover rates of NYC attorneys and find that a firm's investment in pro bono legal services attenuates the attrition rates of NYC-born attorneys after Sept. 11, 2001, relative to non-NYC-born attorneys. Given the magnitude of the heterogeneous attrition response in our setting and the significant economic implications for the firm, we conclude that understanding the link between incentives and the retention and selection of different types of workers is crucial for managerial decision-making.

2 Firm Setting and the Compensation Changes

The compensation changes that we study happened in a US-domiciled inbound sales call center, which employed over 2,000 sales agents over the course of our sample period. The agents are organized into six divisions, based on the products and services (henceforth, products) they sell. The presidents of two different divisions, Division 1 and Division 2, drastically changed the commission schedules of the agents in their divisions, which ultimately led to significant decreases in the average commission and take-home pay of their workers.[7]

2.1 Firm Setting

The firm contracts directly with national television, phone, and internet providers to market and sell their products. Prospective customers respond to the firm's marketing promotions by calling an 800-number that corresponds to a particular service provider's products. The sales agents respond to these inquiries and, when appropriate, try to upsell customers on high profit margin products (e.g., larger bundles of TV channels or faster internet speeds). Each of the six sales

[7]Division 1 and Division 2 employed 20% and 7% of the firm's sales force, respectively. Approximately 600–850 agents were employed at a given point in time during the sample period. At the time of the compensation changes, most of the firm's employees were in two separate offices, located 50 miles from each other.

divisions is managed by one or two division presidents, who have sole discretion to adjust their sales agents' commission schedules. Within each division are multiple teams of approximately 10 to 15 sales agents, who are supervised by a sales manager.

The different sales divisions are uniquely characterized based on the products their agents sell. Agents in Division 1, for instance, handle calls from residential customers, inquiring about products from a particular service provider. Agents in the control divisions respond to residential inquiries about products from different, albeit similar, service providers. Agents in Division 2 respond to inquiries from small businesses, but the products offered resemble those of the other divisions. To accommodate the lucrative opportunities of interacting with small businesses, the firm reserves space in Division 2 for its most productive and experienced agents. Inbound calls are routed to a particular division based on the phone number dialed by the prospective customer. Within a division, calls are allocated to the next available agent in the queue (i.e., whomever has been waiting the longest). This means sales opportunities are allocated to agents randomly.[8]

Agents rely on designated sales protocols and their understanding of the caller's needs to sell the products. Success depends on an agent's understanding of the products, his or her ability to master the sales protocol, and his or her ability to upsell customers onto high profit margin products. In most cases, the highest margins are earned on the most expensive products (e.g., a satellite subscription with all possible channels) or bundles of products (e.g., a service contract covering internet, telephone, and television), which are more difficult to sell. Agents generally earn more in commissions for selling high profit margin products than they do for selling low profit margin products. Agents spend about 80% of their workday either on calls or waiting for another call to arrive, and they have little scope to change the number of calls they receive.

2.2 Agent Compensation

Agent pay is made up of a fixed hourly wage, commissions, and occasional small bonuses. Agents start at an hourly wage of approximately 150% of minimum wage and receive small raises for every three months of tenure. Hourly wage rates are capped at approximately 200% of the minimum

[8] The firm is almost exclusively an *inbound* call center, meaning that agents answer calls from interested customers. Less than 3% of calls are outbound—most of which are agents following up on earlier inbound calls (e.g., returning a dropped call).

wage. Agents who stay with the company beyond a waiting period are eligible for health benefits. In addition to fixed wages, commissions are a significant part of an agent's total compensation. During the eight weeks before the compensation changes occurred, the average Division 1 agent earned $318 per week in commissions, and the average control division agent earned $201. These amounts constituted approximately 30%–40% of these agents' overall take-home pay.

An agent's commission pay is determined by his or her weekly sales revenue. The mapping of products sold by an agent to the commission pay received by the agent—i.e., the commission schedule—is determined as follows. Each product has a transfer price assigned to it by the division presidents, which the firm refers to as "revenues." These revenue amounts approximate the actual top-line revenue generated for the firm through the sale of the product. For instance, a low-end cable TV package may be assigned a revenue of $50, while a high-end package may be assigned a revenue of $200. These amounts form the basis for which agents receive commissions, and division presidents set them in a way that (1) rewards agents for each sale they make and (2) provides greater rewards for selling high profit margin products.[9] Throughout the week, agents book revenue through each sale they make. At the end of the week, these revenue amounts are summed and multiplied by the agent's commission rate, which is a function of the agent's audited call quality and selling efficiency,[10] and this determines the agent's weekly commission payment.

2.3 Changes to the Commission Schedule

In November of 2016, the presidents of Division 1 radically re-calibrated their agents' commission schedule by changing the assignment of revenue transfer prices to the products sold by their agents. Prices for customers remained unchanged, as did the per-product top-line accounting revenues realized by the firm. Similarly, agents' commission rate functions were not changed. For many of the most popular products, transfer prices were adjusted such that agents booked less revenue for selling low profit margin products. To lessen the blow to morale, the presidents

[9] Upstream service providers pay the firm for every sale in accordance with set contracts, which leads to the top-line revenue generated for the firm. All use of the term "revenue" in this paper refers to the transfer prices the firm uses to incentivize agents.

[10] Every agent has a fixed number of calls audited each week. If any conduct violations are identified, the agent's weekly commission rate is reduced. Selling efficiency is based on revenue-per-call (RPC) and revenue-per-hour (RPH). Being in higher quintiles of RPC and RPH increase an agent's commission rate.

pitched the changes as an opportunity for agents to earn more by selling high profit margin products, and some transfer price levels were shifted upward for these products. Figure 1 gives an example of changes in the commission schedule for two different types of internet packages.[11] The example suggests that an agent's revenue would drop from $375 to $315 if the agent was unable to sell a greater proportion of higher quality internet packages. Though management framed the changes as an opportunity for the workers to earn more, survey evidence in Section 2.6 indicates that the agents were aware that they were likely to take home significantly less pay due to the changes.

According to the firm's management, the changes to the commission schedule were intended to decrease the relatively high commission pay levels that Division 1 agents were earning in the months before the changes. These relatively high commissions were caused by the addition of new territories from which Division 1 agents fielded calls.[12] The inclusion of these new territories (henceforth, the territory shock) significantly increased the average commissions of Division 1 agents. Figure 2a shows the evolution of average commissions by division before and after the commission schedule changes. The pre-treatment period, Week -26 to Week 0 (with Week 0 denoting the week before the commission schedule changes), is separated into three periods around the territory shock. The weeks before Week -16 constitute the pre-territory shock period. The territory shock period runs from Week -16 to Week -8, representing the period of increasing commission levels for Division 1 agents. The period from Week -8 to Week 0 makes up the post-territory shock period. Division 1 agents' commission levels more than doubled from the pre-territory shock period to the post-territory shock period, increasing from $157 to $318, and then the effects of the territory shock stabilized in the eight weeks before the commission schedule changes.

Based on the productivity data provided by the firm, we estimated that the commission schedule changes would reduce the commission pay of the average Division 1 agent by 18%, holding fixed the pre-treatment period mix of products sold. When we compare average

[11] Table A.1 in the appendix provides a more detailed example of the commission schedule changes for a set of products that can either be sold individually, in a bundle of two products, or in a bundle of three products.

[12] The new territories were likely home to fewer competitors, and the data suggests that the sales success of Division 1 agents increased during the territory shock period.

commission levels in the pre-territory shock period ($318.39 from Weeks -8 to 0) to those in the post-treatment period ($249.71 after Week 0), the implementation of the new commission schedule led to a 21.5% decrease in commission pay for the average Division 1 agent.[13]

Because we learned of the impending commission schedule changes before they were announced, we followed the insider econometrics approach advocated by Bartel, Ichniowski, and Shaw (2004) to interview managers at the firm to assess their predictions for agent reactions. The presidents and managers in Division 1 believed agents' responses to the changes would be muted. Other leaders within the firm, however, expressed concern about increased turnover among affected agents. Few sales managers mentioned changes in effort, because, while the strength of incentives would fall, high-powered incentive pay would remain a significant component of total agent compensation. Despite managers' lack of focus on effort, agents did have discretion to meaningfully adjust their effort. For example, managers reported that, on calls involving upselling, agents often tried several different approaches before converting a sale. A plausible way for agents to reduce their effort is by trying fewer approaches for selling high profit margin products and instead simply fulfilling orders for easier to sell products. Earlier promotions suggest—and subsequent academic experiments confirm—the ability for agents in this firm to adjust their effort, as revenue temporarily increased in response to temporal sales incentives Sandvik et al. (2020). Another example of possible effort adjustment is reducing adherence to one's work schedule by taking more and longer breaks. Although schedule-adherence is tracked by the firm, agents are only penalized if their adherence level dips below a threshold of 80%, and the average pre-treatment adherence level among Division 1 agents was 83%.

Three months after the commission schedule changes occurred in Division 1, the presidents of Division 2 implemented similar changes.

2.4 Personnel and Productivity Data

We identify the consequences of the commission schedule changes using highly detailed commission, personnel, and productivity data provided by the firm. Division 1 and the control

[13]The estimated decrease is 20.5% in a regression that controls for agent demographics, division fixed effects, and time fixed effects, and we cannot reject that this point estimate equals the ex-ante estimate of an 18% reduction in commissions.

divisions all have consistent data beginning in April of 2016. The sample is organized by agent-week and runs through June of 2017. This dataset covers 2,033 unique sales agents across 61 weeks, for a total of 39,944 agent-week observations. This dataset includes proxies of worker effort—e.g., adherence, conversion rate, phone hours, revenue generated per call, total revenue generated per week—demographic details—e.g., age, race, tenure, gender, marital status—and commission pay data.[14] We refer to this as the immediate sample.

Table 1 displays pre- and post-treatment period summary statistics for the control divisions (Divisions 3–6) in Columns 1 and 4, Division 1 in Columns 2 and 5, and Division 2 in Columns 3 and 6. The pre-treatment period is restricted to the post-territory shock period (Weeks -8 to 0) to highlight the division-level characteristics immediately before the commission schedule changes occurred. We fail to reject the null at the 1% level that Division 1 and control division agents are similar in Tenure (p-value = 0.758), Age (p-value = 0.336), Race (p-value = 0.887), and Gender (p-value = 0.496). We do find that agents in the control divisions are significantly more likely to be married (p-value = 0.001). Division 1 agents have higher adherence (p-value < 0.001), but both groups are at or above the firm's mandatory level of 80%. Both groups spend a similar number of hours talking to customers each week (p-value = 0.132), though Division 1 agents realize higher commissions and greater RPC (p-value < 0.001), largely due to the territory shock experienced two months before the compensation changes occurred. Agents in Division 2 earn much more in commissions (p-value < 0.001) because they sell to small businesses, rather than residential customers. For expositional ease and because we do not have the full range of performance variables for Division 2, we discuss the effects of their commission schedule changes in Section 4.5. During the pre-treatment period, agents in the control divisions and Division 1 were predominately male, 70%–73%. The average agent was 25–26 years old and had been working at the firm for about a year.

To estimate turnover responses, we supplement our immediate sample with additional data that dates to July of 2015. This extended data has limited performance data on sales revenue, but it does contain commissions, allowing us to identify and control for seasonal (year-to-year)

[14] Adherence equals the fraction of time an agent is available to answer calls over the total time an agent is scheduled to be available; conversion rate equals the ratio of calls received where a sale is made to total calls received; and phone hours equals the number of hours in a week an agent spends talking to a customer.

patterns in compensation and attrition. We use this extended dataset, called the extended sample, for our turnover analysis.

2.5 Estimating Baseline Agent Productivity

A central test of theories around heterogeneous turnover, emphasized in the adverse selection discussions in Campbell III and Kamlani (1997) and Bewley (1998), concerns how agents of different productivity levels respond to a change in their compensation or employment contract. To identify these heterogeneous responses, we begin by estimating agents' productivity levels prior to the commission schedule changes. We estimate productivity by using a fixed effects regression analysis of log commissions, which is an omnibus measure of sales productivity that is available in both the immediate sample and the extended sample. We calculate an agent's adjusted worker fixed effect using a regression of log commissions on the worker's tenure profile, division-by-week fixed effects, and agent fixed effects.[15] To minimize the impact of measurement and sampling error, we adjust the raw worker fixed effects, according to best practices in the literature.[16] We use the resulting adjusted worker fixed effects as a measure of agents' baseline productivity, allowing us to identify high and low performers in each division before the commission schedule changes occurred in Division 1.

Because our focus is on how agents with different productivity levels respond heterogeneously to the commission schedule changes, Table 2 provides summary statistics for Division 1 in the pre-treatment period by splits of the sample into terciles based on the adjusted worker fixed

[15] We use data leading up to four weeks before the changes. Log commissions are used, rather than commissions-per-call/hour, revenue-per-call/hour, or log revenue, because we do not have data on the number of calls received, the number of hours worked, or the revenue generated in the extended sample for 2015. In addition, accounting for the tenure profile makes this measure one about underlying talent, rather than productivity improvements that may come from learning or on-the-job experience. We include agent tenure in the main specification because those with higher productivity are less likely to leave the firm and may have greater tenure.

[16] The adjusted worker fixed effects are interpretable up to a division-level average that is removed through the division-by-week fixed effects. To account for sampling variation in the estimated fixed effects, we run the main regression and collect the residuals plus the estimated worker fixed effects. We then follow the procedure of Lazear et al. (2015) and fit a restricted maximum likelihood random effects estimator and recover each worker's expected best linear unbiased predictor of their latent fixed effect. The estimator resembles an empirical Bayes procedure, where noisier sequences of data on individual workers receive less weight; less noisy data moves the estimated fixed effect away from 0 (a normalization). We call the resulting output the adjusted worker fixed effects. The adjusted worker fixed effects guard against mean reversion or classification being driven by sampling error from a short panel.

effects. As shown in Table 2, agents in the top tercile have higher tenure, in line with the firm retaining highly productive workers. Demographic characteristics also vary across the adjusted worker fixed effects terciles; namely, workers in the highest tercile are older and less likely to be single. Later specifications control for these characteristics.

The interpretation of our upcoming analysis would be muddled if the commission schedule changes affected high and low performers differently, due to their selling different product mixes before the changes. To check for this, we calculate the expected percentage change in commissions after the commission schedule changes were made, based on the sales mix in the pre-treatment period. The variable *Predicted Pct* Δ *Commission Post-Treatment* in Table 2 reports this measure. The predicted percentage change in commissions, due to the pre-treatment period sales mix, is 17%–18% across all tercile groups of workers. Although the top tercile of workers has average weekly commissions that are more than 2.5 times greater than the bottom tercile, the product mix of sales does not impact percentage changes in commissions significantly more for any one group based on pre-treatment sales behavior. Put differently, the commission schedule changes affected the expected percentage change in commissions equally across all terciles of agents in Division 1.

2.6 Surveys of Sentiment and Reactions to the Changes

We conducted a firm-wide survey before the announcement of the changes to gather information regarding agents' sentiment toward the firm. We asked sales agents from all divisions the following three questions: (1) "How likely are you to agree with the following statement, [the firm's] policies, for example on adherence, compensation, and promotion, are justified and fair?" (2) "Suppose your friend is looking for a job, how likely are you to recommend them to apply at [the firm]?" (3) "Do you think you will be promoted in the future?"[17] In addition, we conducted a follow-up survey among agents in Division 1 after the announcement of the changes and before these agents received their first paycheck reflecting the new commission schedule. In this survey,

[17] The possible answers to these questions are as follows: (1) Unit increment slider from 0 (Strongly Disagree) to 100 (Strongly Agree). (2) Unit increment slider from 0 (Not Likely) to 100 (Very Likely). (3) "Yes, within 0–3 months," "Yes, within 3–6 months," "Yes, within 6–12 months," "Yes, in over 12 months," "I don't want a promotion," and "No, promotion is not likely."

we again asked the three sentiment questions (which are discussed in detail in Section 4.4), and we also asked agents several questions related to the effects of and motivation for the commission schedule changes.[18]

Responses to the follow-up survey reveal that the average agent in Division 1 expected their commission pay to decrease by 13% (see Figure 3a), which approximates our estimate of an 18% decline in commissions.[19] The average agent reported that they *would need to* work 11% harder in response to the changes to maintain their usual commission pay (see Figure 3b).[20] Agents then reported that *they would*, on average, increase their effort by 7% in response to the changes (see Figure 3c).[21] Agents' responses to these last two questions suggest that income effects are likely at play in this setting. Several other questions on the follow-up survey asked agents about the motivation for the commission schedule changes. Over 75% of the agents felt the motivation for the changes was clearly communicated by management at the time of the announcement. When asked why the changes occurred, 42% responded with "Sales reps were overpaid,"[22] and 40% responded with "[The firm] needs to make cutbacks to stay in business." The follow-up survey also provides evidence that the changes were unanticipated by the sales agents, as only 2% of the agents say they knew the details of the changes before they were announced.

3 Identifying Assumptions and Common Trends

We are interested in how the commission schedule changes impacted the turnover and effort of the affected agents. The presence of unaffected divisions motivates the use of a difference-in-differences estimation. Difference-in-differences relies on the assumption that treated

[18] The first prompt of both surveys told respondents "This project was specifically outsourced to increase privacy, so you can answer these questions truthfully." We believe agents responded truthfully to the survey questions, though we can never be sure how agents will interpret this statement.

[19] We asked, "By how much will the changes affect your commission pay (assuming you work just as many hours and just as hard as before)?" Agents responded using a slider from -50% to 50%, indicating a 50% decrease to 50% increase.

[20] We asked, "How would your level of effort have to change in order to maintain your usual commission pay?" Agents responded using a slider from -20% to 20%, indicating a 20% decrease to 20% increase in effort exertion.

[21] We asked, "How will the revenue change affect your choice of effort relative to before the change?" Agents responded using a slider from -20% to 20%, indicating a 20% decrease to 20% increase in effort exertion.

[22] This response implied that management felt agents were overpaid, not that agents, themselves, felt they were overpaid.

and untreated individuals would follow a common trend in outcomes in the absence of the commission change. To assess the suitability of using other divisions as a control group, we consider trends across several variables. First, we provide evidence of common trends in the attrition rates of agents in Division 1 and the control divisions over many months leading up to the commission schedule changes. We then focus on trends in output measures, which proxy for effort. Using data from two months prior to the commission schedule changes in Division 1, we show that trends are similar across two different proxies of effort supply (adherence and conversion rates) and two different proxies of effort demand (call volume and phone hours). To bolster the case that outcome variables evolve in parallel, we show graphically that there is no divergence in these effort proxies between treated and control divisions after treatment. The smooth evolution of these effort measures across treated and control divisions suggests that potentially problematic trend divergence is unlikely in our setting. Finally, we discuss trends in commission pay. We show that agents in Division 1 and the control divisions follow common trends in commission pay after the territory shock and before the commission schedule changes.

In light of the territory shock, we clarify the identifying assumptions required for our setting. The common trends assumption is necessary to estimate a counterfactual—allowing researchers to understand what would have happened absent treatment. In our setting, the stability of trends between Division 1 and the control divisions after the territory shock enables the identification of counterfactuals. However, note that the territory shock itself had the potential to change agents' reference points or to make the job relatively more attractive than it had been before the territory shock. This likely makes our estimates of turnover and effort responses lower bounds of the consequences that managers may anticipate when adjusting pay. While the territory shock does not invalidate our use of a difference-in-differences estimation, we further discuss the implications for interpreting the results in Section 3.3.

3.1 Common Trends in Turnover

Our first outcome of interest is the turnover response of Division 1 agents. We graphically assess the pre-treatment trends in agent turnover to investigate the validity of comparing turnover rates between Division 1 and the control divisions. We focus on the Kaplan-Meier estimator, which plots survival rates over time, because it allows visualization of the cumulative nature of turnover. The survival rate estimator considers a starting point and then, from that time, displays the fraction of agents who remain at the firm. This allows an assessment of when retention diverges and what fraction of the total beginning workforce is affected.

Figure 4 plots the survival rates for agents in Division 1 and the control divisions. To focus on heterogeneous turnover responses, we separately plot the survival rates of high and low performers.[23] Figure 4a shows that highly productive workers in both Division 1 and the control divisions follow a similar trend in retention from Month -5 to Month 0. Similarly, low performers in Division 1 have survival rates that closely track those of low performers in the control divisions. The closeness of these survival rate trends suggests that agents in the control divisions provide an effective group for comparison to estimate the turnover responses of agents in Division 1.

3.2 Common Trends in Effort

Our second outcome of interest is the effort response of Division 1 agents. To evaluate the credibility of the assumption of common trends in effort, we investigate trends between Division 1 and the control divisions using two well-known approaches (Fowlie, Greenstone, and Wolfram, 2018; Cengiz, Dube, Lindner, Zipperer, CEP et al., 2019).[24] First, we provide graphical evidence of conditional differences in trends. Second, we report p-values from an empirical test of differential trends. Both of these tests use the following event study design:

$$y_{i,t} = \sum_t \delta_{i,t}\mathbb{I}(year = t) \times Treated_i + \beta_i Treated_i + \sum_t \lambda_t \mathbb{I}(year = t) + X_{i,t}\Gamma + \sum_j \gamma_j Div_j + \zeta_{i,t}, \quad (1)$$

where λ_t are time fixed effects that capture the common trend, $Treated_i$ equals one for agents in Division 1 and zero otherwise, and $\delta_{i,t}$ captures the possibility of different trends. $X_{i,t}$ includes controls for location and agent characteristics. Div_j is a division fixed effect, and $\zeta_{i,t}$ is an idiosyncratic error term.

[23] We split high and low performers based on whether agents have adjusted worker fixed effects that are above or below the median within their division.

[24] While a direct test of the identifying assumption is not possible, the evidence from these approaches provides some reassurance (Kahn-Lang and Lang, 2018).

Figure 5 highlights the common trends in observable outcomes that proxy for the supply of and demand for worker effort. The coefficients in these figures are estimates of $\delta_{i,t}$ from Equation (1). To improve the readability of these and other figures, we aggregate data into bi-weekly clusters. Hence Week 0 reflects data from Week 0 and Week -1, Week -2 reflects data from Week -2 and Week -3, etc. Figures 5a and 5b display trends in adherence and conversion, respectively, the two main proxies of worker effort supply. Both figures show indistinguishable differences in effort between Division 1 and the control divisions before the changes in Week 0. The p-values of tests that the Week -8 to Week 0 point estimates are jointly equal to zero are 0.44 and 0.31 for adherence and conversion, respectively.

Figures 5c and 5d display differences in brand-level call volume and phone hours, the two main proxies for demand for worker effort. Both figures again show differences in effort that are indistinguishable from the week immediately before the change. P-values from joint tests from Week -8 to Week 0 are 0.36 and 0.76 for call volume and phone hours, respectively. To show that call volumes and the amount of time spent working with customers do not change, due to the commission schedule changes, we also test that the point estimates are jointly equal to zero across the period from Week 0 to Week 8. The p-values in this case are 0.98 and 0.76. Thus any changes in worker output that we observe are not likely due to reduced call volumes or time spent talking to customers.[25] This suggests the effects of the territory shock were permanent and that they had stabilized in the two months proceeding the commission schedule changes. Across these four proxies of effort supply and demand, we fail to reject the null hypothesis of common trends between agents in Division 1 and the control divisions. In addition, the estimates from Equation (1) provide leads and lags tests to show that no other events occurred at different times around the compensation changes to shock the supply of and demand for agent effort.

3.3 Common Trends in Commissions

Figure 2a shows that commission pay in Division 1 and the control divisions, while different in levels, follows a common trend before the territory shock. The commission levels of Division 1

[25] All the agents in our sample have an employment contract requiring their presence during scheduled hours. Approximately 85% of the agents work at least 30 hours a week, and they have limited scope to adjust their hours.

agents deviate from this common trend during the territory shock period, but they appear to level off in the post-territory shock period and again track the commission levels of agents in the control divisions. The implementation of the commission schedule changes again shocks the trend of Division 1 commission levels after Week 0, but the two groups appear to follow similar trends from Week 4 through at least Week 16. In addition, the commission levels of agents in the control divisions are relatively flat in the two weeks immediately following the Division 1 commission schedule changes, suggesting that the changes had limited spillover effects into other divisions. We further discuss tests that show limited spillovers across groups in Appendix 8.1.

In our setting, we expect common trends in commission pay levels after the territory shock. We focus on this eight-week period immediately prior to the commission schedule changes because we know that trends differed during the territory shock period. Figure 2b plots the coefficients, $\delta_{i,t}$, estimated using Equation (1), just as was done to assess the common trends in effort in Section 3.2. We find evidence of common trends in commission levels when we plot the coefficients across time. The point estimates in Weeks -8 to -2 are all close to zero, and zero always exists within the 95% confidence intervals around these points. Furthermore, we fail to reject the null hypothesis that the coefficients jointly equal zero (p = 0.82). This result suggests that Division 1 and the control divisions followed common trends in commission pay levels in the two months before the commission schedule changes occurred. We overlay a plot of differences in brand-level call volume to show that the observed differences in commission levels after the commission schedule changes are not driven by brand-specific variations in call volume. Managers of the firm confirm that, in the absence of the commission schedule changes, agents in Division 1 would have continued to realize the high commission levels they enjoyed in the post-territory shock period.

In the following section, we perform our formal empirical analysis to identify the turnover and effort responses after the commission schedule changes. The figures presented in this section suggest that the differential pre-treatment trends in commission levels caused by the territory shock in Division 1 are not a major concern for our identification. Because the territory shock was demand-driven, we would be concerned if there had been a significant change in measurable demand around the commission change. Call volume and phone hours remain stable over this time, which suggests that the territory shock was permanent and that commissions would not have reverted to their prior levels, absent managements' changes to the commission schedule. While the territory shock may have caused some loss of "balance" between treatment and control groups, we detail a number of empirical strategies we use to assess this possibility in Sections 4.1 and 4.2.

4 Results and Exploration of Mechanisms

This section details Division 1 agents' turnover and effort responses to the changes in their commission schedule. We also estimate the firm-level effects of the observed worker responses. We then consider the role of sentiment in our findings and discuss the turnover effects of Division 2 agents. We begin by discussing agent attrition, which is the margin that we find has the largest impact firm performance. We then consider agent effort. Our empirical analysis is motivated by a theoretical model, which, for brevity, we discuss in Appendix 6. The model provides context for our estimates by showing that whether a compensation change is profitable depends on (1) the cost changes that affect the firm's wage bill, (2) workers' effort changes, and (3) changes in the composition of the workforce, due to asymmetric turnover based on agent productivity.

4.1 Turnover Responses

In Section 3.1, we introduced Figure 4a to show the common trends in attrition between agents in Division 1 and the control divisions in the months before the commission schedule changes. This figure also shows that the survival rates of highly productive agents in Division 1 break from those of highly productive agents in the control divisions in the post-treatment period. This figure conditions on agents who were present at the firm several months prior to the commission schedule changes, which is useful as a diagnostic tool for pre-trends. Figure 4b, on the other hand, considers survival rates relative to the sample of agents present in each group in what is labeled Month 0, the calendar month immediately before the changes occurred (October 2016). After the commission schedule changes, the survival rate of high performers in Division 1 decreases, relative to that of high performers in the control divisions, whereas the survival rate of low performers appears to increase. This is preliminary evidence of a heterogeneous turnover effect, wherein highly productive agents in Division 1 are more likely to leave the firm in response to the commission schedule changes.

Next, we present another graphical approach to assess how the composition of the workforce changed due to employee attrition after the commission schedule changes. In Figure 6, we plot the average adjusted worker fixed effects for Division 1 and the control divisions.[26] As in many sales firms, there is positive selection by worker quality over time, captured by the upward trend in average adjusted worker fixed effects in all divisions in the pre-treatment period (represented by points to the left of the vertical line). There is, however, clear evidence that average worker quality begins to deteriorate in Division 1 several weeks after the commission schedule changes. This divergence in adjusted worker fixed effects provides graphical evidence that, in response to the commission schedule changes, agents with high pre-treatment productivity exited the firm at a higher rate than did agents with low pre-treatment productivity.

With this graphical evidence in hand, we examine turnover formally by using a difference-in-differences estimator. These estimations use the extended sample, which predates the immediate sample by at least a full calendar year for each division. Our goal is to identify how turnover was affected when agents' commission schedule changed and how the turnover probability differs based on agent productivity. In an analysis of turnover, it is necessary to account for how the baseline probability of leaving the firm changes with worker tenure (Bartel and Borjas, 1981). We use a very flexible specification for how the usual probability of leaving the firm changes with tenure by including a flexible function, $g(Tenure)$, in the model. We specify this function as a fifth-order polynomial, providing enough flexibility to capture the possibility that workers with longer tenures are less likely to leave and that the relationship between attrition and tenure has several inflection points. This function is distinct from time fixed effects, which are meant to capture calendar time shocks, like seasonality, that affect all workers. We then include combinations of division and time fixed effects to capture permanent heterogeneity across

[26] These adjusted worker fixed effects are for those who are present at the firm prior to the commission schedule changes. Because the fixed effects are estimated using data that ends four weeks prior to the event date, we can assess whether the fixed effects (normalized to the event date) track one another in the pre-treatment period. Changes in the average adjusted worker fixed effects summarize how turnover differs by workers' productivity in Division 1 and the control divisions.

divisions and seasonal shocks that may be correlated with treatment. The model we estimate is:

$$\begin{aligned} Turnover_{i,t} = Div_j + \delta_1(Treated_i \times Post_t) + \delta_2(Treated_i \times Post_t \times Prod_i) + \beta_1(Prod_i) + \\ \beta_2(Treated_i \times Prod_i) + \beta_3(Post_t \times Prod_i) + TimeControls + g(Tenure) + \beta_4 X_{i,t} + \varepsilon_{i,t}. \end{aligned} \tag{2}$$

The dependent variable, $Turnover_{i,t}$, is an indicator that the week in question is worker i's last week in the firm. After the worker leaves, he or she is no longer included in the sample. The dependent variable is thus the instantaneous turnover probability, or hazard, given that the worker was at the firm in the week in question. The parameter δ_1 captures the average change in turnover probability of agents in Division 1, conditional on tenure and time controls, after the commission schedule changes occurred. This is indicated by $Post_t$, the post-treatment indicator, being interacted with $Treated_i$. We include division fixed effects, Div_j, to control for division-level differences in attrition. The matrix $X_{i,t}$ has a third-order polynomial in age, along with fixed effects for ethnicity, gender, call center location, and marital status. The separate tenure splines and age polynomials allow the effects of experience within the firm and total labor market experience to differ. We include baseline measures of worker productivity, captured by $Prod_i$, and its interaction with post-event indicators. To identify differences in productivity, we use the standardized z-score of adjusted worker fixed effects in the pre-treatment period. We use z-scores to standardize the adjusted fixed effects across Division 1 and the control divisions. This approach also facilitates the interpretation of the parameters, as a unit change in the z-score, $Prod_i$, corresponds to a standard deviation of the underlying productivity measure.

Table 3 displays the turnover responses of agents in Division 1, relative to those in the control divisions. The different columns correspond to different combinations of $TimeControls$ to account for a variety of possible seasonal differences across divisions. Across all specifications, highly productive workers in Division 1 became more likely to leave after the commission schedule changes. The point estimates on *Treated x Post x Prod* across Columns 1–4 indicate that Division 1 agents with pre-treatment productivity one standard deviation above the mean had turnover rates that increased by 1.3–2.1 percentage points in a given week, compared to Division 1 agents with average pre-treatment productivity. This turnover increase is relative to an overall sample mean of about 0.037, indicating that agents one standard deviation above the mean had between a 40%–56% increase in attrition from the sample average.

These turnover effects are precisely estimated when clustering by the identity of a worker's manager. If we instead cluster standard errors at the division level, which was the level of the intervention, our standard errors are similar. However, we do not report these standard errors, because test statistics based on them are misleading, due to the fact that we only have a few divisions and only one treated group. Instead we correct for this by conducting statistical tests using a combined randomization inference and wild bootstrap procedure designed to estimate critical regions under clustering with few treated clusters (MacKinnon and Webb, 2018).[27] This estimator has been shown to perform well in simulations and avoids problems of over-rejection that are often endemic when there are few clusters. The p-values from these tests are displayed in the bottom rows of Table 3 for δ_1 and δ_2.

Including a placebo treatment indicator offers a direct test of whether the attrition patterns we observe would have occurred at a prior time. The most natural prior time to test is the exact date in the prior calendar year, and the specification in Column 1 includes placebo indicators that are dummies for one year prior to the announcement week. This tests whether the observed turnover effect is due to the commission schedule changes or annual patterns in turnover. The zero coefficients on *Treated x Placebo x Prod* and *Treated x Placebo* indicate that the turnover patterns overall and by agent productivity level did not diverge between Division 1 and the control divisions at the same time in the past.

Columns 2 and 3 show that the estimates are robust to the inclusion of different combinations of time, week-of-year, and division fixed effects. The division by week-of-year fixed effects in Column 2 compares division-level turnover rates across calendar years, which captures possible seasonality by division and guards against the possibility that Division 1 had a similar seasonal change in turnover in the prior year. The point estimate of 0.015 suggests that workers plus or minus one standard deviation around the mean had post-treatment turnover rates of 5.2% and 2.2%, respectively. This difference provides strong evidence of a heterogeneous turnover response across the distribution of worker productivity. Column 3 includes time by division fixed effects,

[27] There is now a significant literature that addresses these issues, and applied papers have generally used some version of the wild cluster bootstrap to get valid confidence regions. See, for example, Lazear, Shaw, and Stanton (2016).

which only allows us to identify heterogeneous turnover by productivity. The advantage of these estimates, however, is that we do not need to rely on common trends by division, only common trends in turnover by different productivity groups.[28] The similarity of the overall estimates in Columns 1 and 2 and those that do not depend on common trends at the division level (Column 3) add credibility to the identifying assumptions. This heterogeneous effect also holds when we restrict the sample to begin eight weeks before the commission schedule changes occurred (Columns 4 and 5), which removes the weeks before and during the territory shock period. The results are also robust to how workers' productivity z-scores are estimated.[29] Taken together, this evidence suggests that highly productive agents in Division 1 were more likely to leave the firm in response to the commission schedule changes.

The point estimates on *Treated x Post* are not precisely estimated in any of the specifications. The inability to reject that the main effects are zero for Division 1 indicates that overall turnover did not increase among these agents. Instead, only agents who were highly productive in the pre-treatment period increased their likelihood of leaving. This is consistent with the trends in survival rates, depicted in Figure 4b. This figure shows that high performers have an increased likelihood of quitting (i.e., a decreased survival rate), whereas low performers have a *decreased* likelihood of quitting. The departure of high performers potentially increased the placement of low performing agents in the RPC quintiles used to determine commission payouts. Consequently, the implicit contract improved for low performers, increasing their incentive to stay in the firm. These offsetting effects provide intuition as to why we do not observe a significant average turnover effect among agents in Division 1.[30] This evidence confirms the concern expressed in the literature that employers fear that their best agents will quit in response to changes in compensation. This concern may explain the general rigidity of compensation arrangements and

[28]The validity of the difference-in-differences common trends assumption is not necessary for an analysis of differential effects across agent productivity levels, because we leverage within-division variation. That is, heterogeneous responses can be estimated using division-by-time fixed effects without appealing to common trends across divisions. Observing how differences between types of agents evolve—in Division 1 and control divisions—helps illustrate the sources of variation. Figures A.1a and A.1b plot the evolution of within-division differences in performance by worker pre-treatment productivity, suggesting common trends within division.

[29]The results are qualitatively similar when the adjusted worker fixed effects are estimated with the omission of the tenure polynomial.

[30]Because this firm uses relative performance evaluation in compensation, the loss of highly productive workers may have helped retain less productive workers by increasing their ranking inside the division.

managers' reluctance to cut pay.

The results from multiple additional placebo tests highlight the robustness of our turnover response estimations. Following Gubler, Larkin, and Pierce (2018), we perform fifty placebo simulations for the turnover response estimation using randomized treatment groups and treatment dates. The results of these simulations are displayed in Figure 7a. The coefficient for 47 of the 50 placebos run is smaller in size and less statistically significant than the estimated coefficient. This is approximately what one would expect from the placebo tests given the statistical significance of the estimate. In Figure 7b, we repeat our turnover estimation procedure using different control divisions as the chosen "treated" division. We also vary time periods to incorporate placebo dates. As the figure shows, our estimate is larger than alternate estimates with placebo divisions and dates.[31] These placebo tests at the division level complement the cluster wild bootstrap procedure. A now significant literature discusses the properties of these tests, and the version we use is approximately the same variation as these placebo tests.[32]

4.2 Effort Responses

Having found evidence of heterogeneity in the turnover of agents in Division 1, we next investigate whether these agents altered their effort in response to the commission schedule changes. The first specifications for estimating the effects of the commission schedule changes on worker effort are difference-in-differences regressions with the following form:

$$y_{i,t} = \alpha_i + Div_j + Trend_j + \delta_1(Treated_i \times Post_t) + \lambda_t + \beta_1 X_{i,t} + \varepsilon_{i,t}. \tag{3}$$

These specifications include time (week) fixed effects, λ_t, and division fixed effects, Div_j. Some specifications include α_i, an individual fixed effect, and some include division-specific time trends, $Trend_j$. To account for the potential that different trends across divisions bias the estimates, we check the robustness of our results by using a propensity score re-weighting estimator to match

[31] Note that we do not include placebo dates for the treated division, as we would expect a larger estimate about two months after the treatment date due to the fact that turnover accumulates over time, with a lag of about 7–8 weeks prior to an uptick.

[32] We impose the null hypothesis and resample over clusters. We then compare t-statistics after this resampling procedure to the original test statistic.

control division agents who were on similar trends as those in Division 1 before the commission schedule changes occurred. This approach aims to better balance treated and control agents, based on levels of and changes in compensation over the entire pre-treatment period.[33] In addition, we verify our results by reducing the sample to a balanced panel of agents who are present in the sample before July 2016 and after April 2017. This ensures that we are picking up variation in agents' behavior before and after the commission schedule changes and not just changes in the composition of workers.[34]

The results of the difference-in-differences estimations using Equation (3) are contained in Columns 1–5 of Table 4. We use data from the eight weeks before and the eight weeks after the commission schedule changes to estimate agents' effort responses. In Table A.2, we show that our results are robust to the inclusion of all the pre- and post-treatment data. Panel A contains results for agents' adherence and shows that, on average, agents in Division 1 do not reduce their adherence in response to the commission schedule changes. We also find negligible differences in agents' conversion rates (Panel B). The null results in Panels A and B are robust to the inclusion of agent fixed effects (Column 2), the inclusion of division-specific time trends (Column 3), the use of a re-weighting estimator (Column 4), and the use of a balanced panel (Column 5). These findings align with the graphical evidence presented in Figures 5a and 5b in Section 3.2. They suggest that agents did not avoid calls nor did they reduce their sales conversion efforts; i.e., we find little evidence of effort adjustment after the commission schedule changes.

We further consider changes in agents' effort by considering two additional proxies of worker sales effort, log revenue-per-call if (1) the commission schedule had *not* changed (Panel C) and (2) if the commission schedule had always been at the new levels (Panel D). In these specifications, we take the revenue transfer prices as given, based on the respective commission schedule regime,

[33] The details of this re-weighting procedure are provided in Appendix 7. Figures OA.1a and OA.1b in the online appendix display the weighted and unweighted measures of log commissions-per-call and commissions, respectively.

[34] Our estimation of changes in worker effort are conditional on the worker remaining at the firm. As turnover takes time to happen, however, we observe almost all treated workers with at least some sales data in the post-treatment period. The inclusion of agent fixed effects also partially addresses the concern that attrition could affect our measures of employees' effort responses. Importantly, the turnover responses that we discussed in Section 4.1 emerge several weeks (at least six) after the commission schedule changes occurred. We would expect effort responses to manifest much earlier, so it is unlikely that our estimates of effort responses are driven by abnormal attrition.

and apply these pseudo-revenues to the volume of products sold.[35] We find minimal evidence of changes in log revenue-per-call at both the old and new revenue levels.[36] That the estimates are positive in Panel D suggests agents might have increased effort after the commission schedule changes, potentially to compensate for income they stood to lose. However, this finding is not precisely estimated in any of the specifications.[37]

We find limited evidence that agents were able to substantially shift from low- to high-margin products. If workers were substituting to higher margin products under the new commission structure, we would have expected to see substantial divergence between the results using the old and new prices in Panels C and D. Instead, both sets of estimates include 0 in the confidence intervals, suggesting minimal ability to substitute to higher margin products.[38] In some specifications, however, we are able to reject the null that the coefficients using new or old prices are the same. Specifically, in Columns 1 and 2 we reject equality at the 1% level, and in Column 3 we reject equality at the 10% level. However, in Columns 4 and 5, which use a re-weighting procedure and a balanced sample, respectively, we cannot reject equality at the 10% level. Comparing these differences, their estimated magnitudes are generally small.[39] Substitution to higher margin products would have implicitly reduced the magnitude of the compensation changes experienced by the agents. As a result, the relationship between the compensation changes and turnover that we estimate is likely a lower bound for the turnover that would have

[35] The revenue transfer prices in the control divisions were constant throughout the sample period.

[36] Our results are qualitatively unchanged when we measure productivity as revenue-per-hour, rather than revenue-per-call. The call based metric is the firm's focal measure and is more salient to sales agents and their direct supervisors, which is why we focus on it. The time based metric provides an interesting complement to this measure. Panels A and B of Table OA.3 use the logarithm of revenue-per-hour as the dependent variable and report similar estimates as those in Table 4, suggesting a limited time spent per call response to the compensation changes. This evidence aligns with the fact that agents have a limited capacity to control the total number of calls received each week. Panels C and D of Table OA.3 use level RPC as the dependent variable.

[37] To account for the possibility that division-level seasonality affects the results beyond the common time fixed effects or division trends, we also check whether lags of the average division commission levels one year prior to the week in question change the results. Results are qualitatively similar with or without these controls.

[38] Similarly, we may expect that productivity would immediately decline as people learn to allocate effort and game new incentive plans (Obloj and Sengul, 2012). In our setting, however, it is difficult to disentangle learning and other time trends because agents were exposed to the new system at the same time.

[39] So the commission adjustments likely did not alter the firm's per-call unit economics, due to a substantial change in the composition of products sold. Additionally, these results are not driven by spillovers or reactions by agents in the control division, as discussed in the online appendix 8.1.

materialized absent product substitution.

A second specification identifies heterogeneous effort responses across agents, based on their pre-treatment productivity.[40] This specification interacts $Prod_i$ into the model in Equation (3), where $Prod_i$ captures differences in pre-treatment productivity across workers:

$$\begin{aligned} y_{i,t} = \alpha_i + (Div_j \times \lambda_t) + \delta_1(Treated_i \times Post_t) + \delta_2(Treated_i \times Post_t \times Prod_i) + \\ \beta_1 Prod_i + \beta_2(Post_t \times Prod_i) + \beta_3(Treated_i \times Prod_i) + \beta_4 X_{i,t} + \varepsilon_{i,t}. \end{aligned} \tag{4}$$

Column 6 of Table 4 reports both δ_1 and δ_2 from Equation (4), capturing the fact that highly productive agents may have different effort responses on some dimensions. Column 7 identifies only the parameter δ_2 by including division-by-time fixed effects in the model. We do not find a heterogeneous reduction in adherence across agents of varying productivity levels (Panel A), which suggests that neither high performers nor low performers responded to the commission schedule changes by avoiding calls or disregarding their schedules. In Panel B, we find that the conversion rates of highly productive agents decreased, relative to those of less productive agents. Similarly, the negative coefficients on *Treated x Post x Prod* in Panels C and D suggest that high performers may have reduced their revenue generation per call, relative to the average agent, but the effects are not precisely estimated.

One possible concern with the estimations of productivity changes is the possibility of mean reversion, which may be amplified in short time periods. Several empirical facts suggest a limited role for mean reversion in the productivity data we present, but we acknowledge the possibility. First, it is unlikely that mean reversion drives these results because Equation (4) accounts for this through $\beta_2(Post_t \times Prod_i)$. The parameter δ_2 on $(Treated_i \times Post_t \times Prod_i)$ thus captures any deviation from natural, agent productivity mean reversion in the post-period.[41] In addition, the

[40] The heterogeneous treatment effects are based on standardized measures, so the average worker will have an effect that is captured by "Treated x Post" because the productivity average is 0. The interpretation for other workers requires multiplying by their productivity, which has mean 0 and standard deviation 1, so the coefficients on these interactions reflect the effect of a standard deviation change around the mean.

[41] The relative reduction in the conversion of high performers is unlikely to be driven by mean reversion. Adjusted worker fixed effects—used to distinguish between high and low performers—are established using pre-treatment data up to four weeks before the changes occurred. The average conversion of agents in each of the three terciles of adjusted worker fixed effects increased from the weeks before this cutoff to the weeks after, suggesting mean reversion is not a likely cause of our findings. For example, top tercile agents in Division 1 had average conversion in September 2016 of 35%. In October, after the adjusted worker fixed

commission levels of Division 1 and the control divisions move in parallel between weeks -24 and -16, a period further removed from the compensation changes. If mean reversion had a substantial effect, we would expect larger sales reductions than we find because sales in Division 1 began at a higher level than in the control divisions. That is, substantial mean reversion would be expected to inflate the sales reductions that we attempt to estimate. Given the modest size of the estimated sales reductions, we expect that mean reversion is unlikely to be driving these results. Taken together, the results in Table 4 imply that agents, of all productivity levels, had rather muted effort responses to the commission schedule changes.

Finally, note that none of these results on effort report the wild cluster bootstrap randomization p-values. This is because of the general insignificance of the effort results, and the alternative procedure would lead to the same conclusions because it is more conservative. Figure OA.4 in the online appendix shows placebo tests where we repeat our effort estimation procedure using different control divisions as the chosen "treated" division. As the figure shows, we cannot reject the null of zero changes in effort when control divisions proxy for the treated division.

4.3 The Firm-Level Effects of the Compensation Changes

A key question for managers is to determine the correct time horizon to use when assessing the profitability of making changes to employee compensation. For the study firm, the answer was critical in deciding whether to reverse the changes or possibly to extend them to other divisions. Said differently, when would a manager be able to separate noise from a secular change in transaction profitability, due to changes in workforce composition? With the benefit of hindsight, the impact of the composition change is apparent. However, because the loss of high performers occurred with a significant lag, as shown in Figures 4b and 6, the short-term return to the firm was positive, due to the compensation savings, negligible short-run workforce changes, and minimal overall effort change. We attempt to quantify the inflection point when the unit profit change to the firm, on a per-call (or per-transaction) basis, might turn negative, due to changes in

effects had already been measured, this average increased to 37%. Across this same time horizon, bottom tercile agents maintained an average conversion of 29% and middle tercile agents increased their conversion from 32% to 33%. We cannot, however, disentangle whether changes in RPC and conversion rates are due to decreased effort or due to highly productive agents losing sales, as the result of aggressively trying to upsell.

workforce composition.

We calculate the unit profit change to the firm by combining changes in revenue-per-call and commission-per-call. We calculate these changes in revenues and costs using several estimated statistics while making several plausible assumptions. All that is required for this calculation is a measure of the agent's average RPC in the pre-treatment period, adjusted worker fixed effects, and an estimate of the change in the turnover probability as a function of adjusted worker fixed effects. We make five further simplifying assumptions. First, as supported by the lack of empirical evidence of changes in worker effort and revenue generation, we assume these effects are zero. Second, we assume that calls are re-routed to an average agent (based on the pre-treatment sales distribution) in the face of turnover.[42] Third, we assume that the per-call commission expense (CPC) for replaced agents reflects the median agent's commissions, which yields slightly greater cost savings compared to using the average. Fourth, we assume replacement agents earn hourly wages that are \$1.00 lower than departing agents; at two calls per hour, this translates into a per-call savings of about \$0.50. Fifth, to match the timing of turnover empirically, we assume that week-to-week turnover differences begin to accumulate with a five-week lag after the implementation of the commission schedule changes.

Using these assumptions, we first calculate the change in per-call revenue to the firm t weeks after the compensation adjustment as:

$$\overline{\Delta RPC_t} = \frac{1}{N} \sum_i \left[(RPC_i^{Pre} - \overline{RPC^{Pre}}) \times [(1 - \tau_i^{Pre})^t - (1 - \tau_i^{Post})^t] \right], \tag{5}$$

where τ_i^{Pre} is the per-week turnover probability for each agent prior to the changes and τ_i^{Post} is the per-week turnover probability for each agent after the changes, which is estimated from agent productivity and the parameter δ_2 in Equation (2). The expression $[(1 - \tau_i^{Pre})^t - (1 - \tau_i^{Post})^t]$ captures the change in the probability of retaining worker i through week t as a result of the compensation changes. The expression inside the summation operator captures the change in average revenue-per-call as the product of agent i's baseline RPC, relative to the mean RPC, multiplied by the change in retention probability through week t for that agent.

[42] This is a conservative assumption. New workers who join are likely to have expected sales that are below the average sales of the cross-section of agents.

Second, we calculate the change in commission-per-call and the fixed-wage bill using a similar expression, shown below. These calculations reflect the baseline cost savings, absent agent turnover, where ΔCPC_i is the change in commission-per-call as a result of the compensation changes for each agent. We take $\overline{\Delta CPC}$ as the average change in commission-per-call in the absence of turnover and then adjust the compensation savings due to agent composition as follows:

$$\overline{\Delta CostPerCall_t} = \overline{\Delta CPC} + \frac{1}{N}\sum_i \left[(\Delta CPC_i - \Delta CPC_{Med.} - 0.5) \times [(1-\tau_i^{Pre})^t - (1-\tau_i^{Post})^t]\right], \tag{6}$$

where the term in the summation is the per-call change in commissions less the change in per-call fixed wages ($0.50) weighted by the difference in retention probability.

Finally, we compute the unit profit change to the firm through week t as any revenue change plus net cost savings: $\overline{\Delta RPC_t} - \overline{\Delta CostPerCall_t}$. To get a rough estimate of the present value of these weekly changes, we use a 12.5% annual interest rate and project forward for six months. We stop the data at six months, because this firm has seasonal hiring that begins in the summer, often a period of substantial workforce changes.

Initially, the cost savings from the commission schedule changes look attractive, saving the firm about $0.68 in compensation expense per-call. Because turnover is minimal in the first few weeks after the changes, there is no offsetting reduction in revenue. However, about two months after the changes (eight weeks), the workforce composition effect reduces average revenue-per-call by $0.58, whereas the labor cost savings in Week 8 equals $0.71. Over time, the decrease in the average revenue-per-call grows more quickly than does the cost savings, and we find that Week 18 is the inflection point (a bit more than four months post-treatment), after which the net present value of the commission schedule changes is negative. Six months after the changes, the firm's gross margin per-call fell by more than 1.7 percentage points.

To put these numbers into context, we estimate the total net present value of the commission schedule change by multiplying the per-call numbers by the actual number of calls per week. At the six-month horizon, the net present value of the commission schedule changes totaled -$75,500. We emphasize that this cost estimate is a lower bound on what is likely a substantially larger cost (more negative net present value) because we do not include the costs of training new hires. Additionally, our analysis does not consider the spillover effects associated with losing high performers. Previous work has shown that high-performing employees are an important resource for raising the productivity of others (Sandvik et al., 2020), so the loss of highly productive workers likely had a deleterious effect on long-term aggregate productivity.[43]

4.4 Commission Schedule Changes and Worker Sentiment

Next we investigate three different questions regarding agents' sentiment toward the firm, as it relates to the commission schedule changes. First, we measure and report changes in agent sentiment from before and after the changes. Second, we consider whether agents' propensity to leave the firm differed based on their pre-treatment sentiment toward the company. Third, we consider whether agents' effort responses differed based on their pre-treatment sentiment.

Prior to the announcement of the commission changes in Division 1, we surveyed agents from all six divisions about their perceptions of firm fairness, their willingness to give referrals, and their future promotion prospects. The exact wording of these questions is provided in Section 2.6. Shortly after the commission schedule changes, we again surveyed agents in Division 1 to see how their sentiment towards the firm had changed.[44] These changes are documented in Table 5, split across terciles of pre-treatment sentiment in Panel A and across terciles of pre-treatment agent productivity in Panel B. The first row of Panel A reports that agents who thought the firm was relatively fair before the changes significantly reduced their perceptions of firm fairness afterward. Those who had low pre-treatment fairness perceptions, on the other hand, increased their fairness perceptions in the post-treatment period. The second row shows that, across all terciles of pre-treatment referral likelihood, agents in Division 1 reduced their reported willingness to refer others to work at the firm. The third row shows insignificant changes in perceptions of promotion prospects for agents with low prospects in the pre-treatment period. Agents with optimistic pre-treatment perceptions of promotion likelihood, however, significantly reduced their perceived promotion prospects after the commission schedule changes.

The first row of Panel B shows no significant changes in high and low performers' perceptions

[43]From a statistical perspective, managers would have needed more data, even beyond the 18-week time frame when returns to the firm became negative, to conclude reliably that the profitability impact due to the loss of highly productive workers outweighed the labor cost savings.

[44]We conducted this follow-up survey among agents in Division 1 *after* the implementation of the changes and *before* these agents received their first paycheck that reflected the new commission schedule.

of the firm's fairness. The second row reports that, across all terciles of pre-treatment productivity, agents reduced their reported willingness to refer others to work at the firm. A difference in means test shows that high performers reduced their willingness to give referrals more than did low performers (reported in the rightmost column). The third row shows relatively small changes in agents' perceptions of their own promotions prospects. In general, we estimate significant decreases in agent-level sentiment after the commission schedule changes occur, with the largest reductions coming from workers with high pre-treatment sentiment. However, sentiment changes are less related to agent productivity.

Next we investigate whether there is heterogeneity in the turnover and effort responses of treated agents, based on differences in pre-treatment responses to sentiment survey questions. Agents with positive sentiment toward the firm may have had better hopes of overcoming the discouragement resulting from the commission schedule changes. On the other hand, agents with high pre-treatment sentiment had the greatest capacity to revise their perceptions downward, which may have led to a greater distortion to their turnover and effort responses. We separately interact the treatment indicator with indicators for high firm fairness perceptions, high referral likelihood, and a belief that promotion is likely.[45] Table 6 supports our earlier finding that highly productive workers in Division 1 increased their turnover rates after the commission schedule changes, relative to the average worker in Division 1. We do not, however, find any significant heterogeneity in treatment effect across responses to the three sentiment survey questions. The point estimates on *Treated x Post x Firm Fair*, *Treated x Post x High Refer*, and *Treated x Post x Promotion* are close to zero and are not statistically significant. The findings are qualitatively similar when we include division by week-of-year fixed effects (Column 2), include time by division fixed effects (Column 3), and when we restrict the sample to include only eight weeks of pre-treatment period data (Column 4). A similar analysis is performed to investigate heterogeneous treatment effects on effort based on agents' pre-treatment sentiment levels. Table 7 reveals no evidence of statistically significant heterogeneity in treatment effect on adherence, conversion, or revenue-per-call across the distribution of agent sentiment. Taken together, these

[45] An agent's firm fairness perception is marked as high if it is above the median value. Referral likelihood is marked is high if it is above the median value. If an agent says he or she is likely to be promoted, the promotion likelihood indicator equals one.

results fail to find differential turnover and effort effects that function through various proxies of employee sentiment.

We caution, however, that our tests are likely under-powered for ruling out agent sentiment as an explanation for the results. A plausible explanation for the increased turnover among the most productive workers is that these workers lost the most money in absolute terms, which may explain their higher propensity to leave. This is consistent with both a fairness explanation and a relative change in outside options. While the survey evidence finds a limited role for the fairness explanation, there are several key limitations. First, agents may not have internalized the impact of the change at the time of the follow-up survey, as this survey occurred before the agents' first post-treatment paycheck. Second, the response rate of the follow-up survey is 30%, possibly inducing selection bias. Third, we cannot use changes in sentiment as the interactive variable of interest, as only Division 1 agents took the follow-up survey.[46] Finally, the survey questions about fairness encompass many aspects of the job, not just pay considerations. Thus, while workers' fairness concerns do not appear to be driving our results, they cannot be definitively ruled out.

4.5 Effects of the Commission Schedule Changes in Division 2

For expositional ease, we deferred the discussion of the commission schedule changes implemented in Division 2 until now. The changes in Division 2 allow us to test whether our main findings generalize. Table A.3 provides summary statistics for Division 2 in the post-territory shock period by splits of the sample into terciles based on the adjusted worker fixed effects. As was the case among agents in Division 1, agents in the top tercile of Division 2 have higher tenure, are older, and are less likely to be single. We begin our analysis of Division 2 by discussing the trends in commission pay levels for Division 2 agents, relative to control division agents, before and after their commission schedule changes. Then we report the results from our difference-in-differences analysis of turnover responses. Finally, we corroborate these findings with graphical evidence that Division 2 workers increased their propensity to quit after the commission schedule changes. We do not consider the effort responses of Division 2 workers, as data limitations prevent us from

[46] We do not have data on sentiment changes for control agents because only Division 1 agents were surveyed after the commission schedule changes occurred.

measuring changes in revenue-per-call for this division.[47]

Figure A.2a shows the evolution of average commission levels in Division 2 and the control divisions before and after the changes. As Division 2 agents sell products to small businesses, their sales and commissions are highly cyclical, which is indicated by the spikes in commissions marked by the vertical dotted lines in Weeks -40 and -20. The week before the changes were implemented, Week 0, denotes another cyclical spike in business, but the drop in commissions after Week 0 is significantly larger than those following previous productivity spikes. This large drop was caused by the implementation of the new commission schedule among agents in Division 2. In the weeks following the commission schedule changes, there is much less volatility in Division 2 commission levels, and the commission levels remain lower than they had been in over seven months. This stability in post-treatment commission levels can be seen in Figure A.2b. The point estimates in Weeks 2–8 are very similar and are significantly less than the point estimates in Weeks 0, -2, and -4. We overlay a plot of differences in brand-level call volume to show that the observed drop in commission levels is not driven by a decrease in the number of calls field by Division 2 agents. If anything, it appears to be the case that call volume in Division 2 increased in the post-treatment period, relative to the call volume in the control divisions.

Table A.4 provides evidence that agents in Division 2 did, in general, increase their rate of attrition after their commission schedule changed. The point estimate on *Treated x Post* of 0.013 in Column 1 can be interpreted as doubling weekly attrition, relative to a baseline turnover rate of 0.0083 per week in Division 2: (0.013 + 0.0083) / 0.0083. We do not detect within-division heterogeneous turnover responses in Division 2, as the coefficients on *Treated x Post x Prod* are not precisely estimated. We do not estimate a fourth specification with a shortened pre-treatment sample, as Division 2 did not experience a territory shock like that experienced by Division 1.[48] Agents in Division 2 are highly productive and highly compensated, so they are already in the right tail of the firm-wide productivity distribution and resemble the best agents in Division 1. Highly productive agents in Division 1 and all agents in Division 2 would have had similar outside

[47] Each division within the firm is responsible for gathering and archiving its own productivity data, and Division 2's processes could not provide the data necessary to estimate counterfactual revenue amounts.

[48] Due to data limitations associated with Division 2, we cannot include interactions with *Placebo*, as we did for Division 1.

options and similar incentives to search for other employment after their commission schedules changed, effectively reducing their take-home pay. This may explain why we observe a heterogeneous turnover effect in Division 1 and an overall increase in turnover in Division 2. Our relatively limited data access for Division 2 prevents us from performing the same net present value calculation for Division 2 as we performed for Division 1 in Section 4.3. The negative NPV resulting from the Division 1 compensation changes and the large increase in attrition of Division 2 workers, however, suggests that the Division 2 compensation changes were also likely NPV-negative in the long-term.

Figure A.3 corroborates the evidence in Table A.4 by showing the Kaplan-Meier survival rates for agents in Division 2, relative to agents in the control divisions.[49] Figure A.3a shows that from Month -5 to Month 0, agents in Division 2 had very high survival rates, relative to agents in the control divisions. Agents in the control divisions maintained their steady survival rate trend after the commission schedule changes were implemented in Division 2. The survival rates of Division 2 agents, however, quickly decreased in the wake of the commission schedule changes.[50] Figure A.3b considers survival rates relative to the sample of agents present in each group in what is labeled Month 0, the calendar month before the changes occurred. After the commission schedule changes, the survival rate of agents in Division 2 decreases at a much faster rate than does that of agents in the control divisions. This graphical evidence supports the findings in Table A.4 that Division 2 agents increased their propensity to quit in response to the commission schedule changes.

[49] We split high and low performers based on whether workers are above or below the median level of adjusted worker fixed effects within their division.

[50] The turnover effect in Division 2 appears to manifest more quickly than the turnover effect in Division 1. This could be due to (1) the relatively larger impact of the changes on Division 2 agents' commission pay, (2) anticipation or information leakage of the changes in Division 2 or (3) the fact that February (the month of the Division 2 changes) is potentially a more natural time to switch jobs than is November (the month of the Division 1 changes). It is likely that all of these factors contributed to the quicker turnover response of Division 2 workers.

5 Concluding Discussion

An important question for managers considering compensation changes is how worker responses may vary across the employee productivity spectrum. In this paper, we document heterogeneous responses to an adverse compensation change affecting call-center salespeople of varying productivity levels. We find that workers with pre-treatment productivity one standard deviation above the mean were 40%–56% more likely to leave the firm in response to compensation changes that effectively reduced overall pay by 7% for the average worker. This increase in the attrition of the most productive employees took several weeks to manifest, and we find very little evidence to suggest that worker sentiment drove the observed turnover responses. These results point strongly to the possibility that highly productive agents left the firm upon securing more attractive outside options, whereas their less productive peers were unable—or unwilling—to do so. While the compensation changes reduced costs, the foregone margins associated with the loss of highly productive workers caused net firm performance to decrease. In comparison to the literature on employee compensation and their on-the-job performance (Lazear, 2000), we find surprisingly little evidence that workers reduced their effort following the compensation changes. We conjecture that the lack of a negative effort response is due to a combination of income effects, income targeting, or reference points (Mas, 2006). Disentangling these explanations is beyond the variation available in this study but merits further research.

While our study firm provides an ideal setting to estimate heterogeneous worker responses to an adverse compensation change, our analysis faces a few limitations. First, we only have data from one firm, making it impossible for us to track employees following their departure from the study firm. As such, we do not know if employees who leave the study firm are making optimal decisions.[51] Second, one of the compensation changes that we study was motivated by an earlier territory shock that raised workers' commission levels. The shock might limit the generalizability of our results if it changed agents' reference points. However, we find similar results when we investigate another compensation change in a second division. Third, we do not observe variation in non-performance pay (e.g., hourly wages), which has been the general focus of the existing

[51] In fact, the existing research suggests that they may not because high performers who switch firms may lose some of their firm-specific comparative advantage (Groysberg, Lee, and Nanda, 2008).

literature (Fehr and Falk, 1999; Dickson and Fongoni, 2018). While this limits our comparability to these studies, our results complement these findings by estimating the consequences of changes in performance pay. Importantly, a growing fraction of the workforce has performance pay (Lemieux, MacLeod, and Parent, 2009), suggesting that our results may be increasingly important for understanding real-world compensation contracts.

Our work contributes to studies on the micro-foundations for downward nominal wage rigidity by linking compensation changes with heterogeneous worker responses. Heterogeneous performance has been a focus in the strategy and management literature (Zenger, 1992; Campbell et al., 2012b; Carnahan et al., 2012), but less is known about heterogeneous responses to *changes* in compensation. Our ability to link responses to underlying productivity validates early evidence from surveys and interviews (Campbell III and Kamlani, 1997; Bewley, 1998), which suggests that fears of losing top talent constrain managers from adjusting compensation (Kahn, 1997).

These findings also have implications for our understanding of monopsony power. The fact that the average worker often does not respond to compensation adjustments is thought to indicate the presence of monopsony in the labor market (Manning, 2003; Dube, Giuliano, and Leonard, 2018). While we find only a negligible average turnover response, the fact that we find a positive relationship between pre-treatment productivity and post-treatment attrition suggests that the empirical evidence upon which our understanding of labor market power is built may be incomplete. Specifically, studies that only examine an average turnover elasticity to infer market power—or firms' bargaining position with employees—may not capture the relatively high mobility and bargaining power of highly productive employees. Understanding the importance of heterogeneity across workers also links to work on strategic human capital as a source of competitive advantage (Barney, 1991; Crook, Todd, Combs, Woehr, and Ketchen Jr, 2011). Our results speak directly to the limits on human capital mobility (Campbell, Coff, and Kryscynski, 2012a) and firms' capacity to appropriate value from employees (Molloy and Barney, 2015).[52]

Our findings have two direct implications for managers. First, high performing employees are

[52]We recognize that worker productivity is influenced largely by both general- and firm-specific human capital. However, to the extent that an agent's firm-specific human capital signals their capacity and willingness to make such investments in the future (Morris, Alvarez, Barney, and Molloy, 2017), overall agent productivity is a reasonable proxy for the labor market's demand for said agent.

the most sensitive subgroup to adverse compensation changes. Therefore while compensation changes may reduce payroll costs across all impacted employees, retention risks—and the subsequent costs—are greatest among top performers. This finding suggests that pay changes that insulate the most productive employees may be beneficial, but more research is required to understand the potential adverse effects of workplace inequality. Second, the presence of performance-pay may limit negative on-the-job responses to adverse compensation changes via income effects.

Finally, earlier work has found that context matters for how individuals respond to changes in their environment, through framing and communication effects (Kahneman, Knetsch, and Thaler, 1986; Chen and Horton, 2016; Englmaier, Roider, and Sunde, 2017). An area for future study is to consider how workers' responses to externally motivated compensation changes might vary compared to the within-firm motivations examined here (e.g., business cycle shocks or pandemics, as examined in Mascarenhas and Aaker (1989)). Taken together, our findings suggest that managers must be prepared for unanticipated consequences when using incentives to change employee behavior.

References

Altonji, Joseph G, Paul J Devereux. 2000. The extent and consequences of downward nominal wage rigidity. *Research in Labor Economics*. Emerald Group Publishing Limited, 383–431.

Ashenfelter, Orley, James Heckman. 1974. The estimation of income and substitution effects in a model of family labor supply. *Econometrica* **42**(1) 73–85.

Barney, Jay. 1991. Firm resources and sustained competitive advantage. *Journal of management* **17**(1) 99–120.

Bartel, Ann, Casey Ichniowski, Kathryn Shaw. 2004. Using "insider econometrics" to study productivity. *The American Economic Review* **94**(2) 217–223.

Bartel, Ann P, George J Borjas. 1981. Wage growth and job turnover: An empirical analysis. *Studies in Labor Markets*. University of Chicago Press, 65–90.

Berman, Ron, Leonid Pekelis, Aisling Scott, Christophe Van den Bulte. 2018. P-hacking and false discovery in a/b testing. Tech. rep.

Bewley, Truman F. 1998. Why not cut pay? *European Economic Review* **42**(3-5) 459–490.

Binyamin, Appelbaum. 2017. Lack of workers, not work, weighs on the nation's economy. *New York Times* .

Burdett, Kenneth, Dale Mortensen. 1998. Wage differentials, employer size, and unemployment. *International Economic Review* **39**(2) 257–73.

Campbell, Benjamin A, Russell Coff, David Kryscynski. 2012a. Rethinking sustained competitive advantage from human capital. *Academy of Management Review* **37**(3) 376–395.

Campbell, Benjamin A, Martin Ganco, April M Franco, Rajshree Agarwal. 2012b. Who leaves, where to, and why worry? Employee mobility, entrepreneurship and effects on source firm performance. *Strategic Management Journal* **33**(1) 65–87.

Campbell III, Carl M, Kunal S Kamlani. 1997. The reasons for wage rigidity: Evidence from a survey of firms. *The Quarterly Journal of Economics* **112**(3) 759–789.

Carnahan, Seth, Rajshree Agarwal, Benjamin A Campbell. 2012. Heterogeneity in turnover: The effect of relative compensation dispersion of firms on the mobility and entrepreneurship of extreme performers. *Strategic Management Journal* **33**(12) 1411–1430.

Carnahan, Seth, David Kryscynski, Daniel Olson. 2017. When does corporate social responsibility reduce employee turnover? Evidence from attorneys before and after 9/11. *Academy of Management Journal* **60**(5) 1932–1962.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, Ben Zipperer, IZA CEP, et al. 2019. The effect of minimum wages on low-wage jobs. *Forthcoming, Quarterly Journal of Economics* .

Chen, Daniel L, John J Horton. 2016. Research note—are online labor markets spot markets for tasks? A field experiment on the behavioral response to wage cuts. *Information Systems Research* **27**(2) 403–423.

Cohn, Alain, Ernst Fehr, Lorenz Goette. 2014a. Fair wages and effort provision: Combining evidence from a choice experiment and a field experiment. *Management Science* **61**(8) 1777–1794.

Cohn, Alain, Ernst Fehr, Benedikt Herrmann, Frédéric Schneider. 2014b. Social comparison and effort provision: Evidence from a field experiment. *Journal of the European Economic Association* **12**(4) 877–898.

Crook, T Russell, Samuel Y Todd, James G Combs, David J Woehr, David J Ketchen Jr. 2011. Does human capital matter? A meta-analysis of the relationship between human capital and firm performance. *Journal of applied psychology* **96**(3) 443.

DellaVigna, Stefano, John A List, Ulrike Malmendier, Gautam Rao. 2016. Estimating social preferences and gift exchange at work. Tech. rep., National Bureau of Economic Research.

Dickson, Alex, Marco Fongoni. 2018. Asymmetric reference-dependent reciprocity, downward wage rigidity, and the employment contract .

Dube, Arindrajit, Laura Giuliano, Jonathan Leonard. 2018. Fairness and frictions: The impact of unequal raises on quit behavior. Tech. rep., National Bureau of Economic Research.

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.

Esteves-Sorenson, Constanca. 2017. Gift exchange in the workplace: Addressing the conflicting evidence with a careful test. *Management Science* .

Fehr, Ernst, Armin Falk. 1999. Wage rigidity in a competitive incomplete contract market. *Journal of Political Economy* **107**(1) 106–134.

Fehr, Ernst, Lorenz Goette. 2005. Robustness and real consequences of nominal wage rigidity. *Journal of Monetary Economics* **52**(4) 779–804.

Fehr, Ernst, Lorenz Goette, Christian Zehnder. 2009. A behavioral account of the labor market: The role of fairness concerns. *Annu. Rev. Econ.* **1**(1) 355–384.

Fowlie, Meredith, Michael Greenstone, Catherine Wolfram. 2018. Do energy efficiency investments deliver? Evidence from the weatherization assistance program. *The Quarterly Journal of Economics* **133**(3) 1597–1644.

Giacalone, Robert A, Jerald Greenberg. 1997. *Antisocial behavior in organizations*. Sage.

Greenberg, Jerald. 1990. Employee theft as a reaction to underpayment inequity: The hidden cost of pay cuts. *Journal of applied psychology* **75**(5) 561.

Groysberg, Boris, Linda-Eling Lee, Ashish Nanda. 2008. Can they take it with them? The portability of star knowledge workers' performance. *Management Science* **54**(7) 1213–1230.

Gubler, Timothy, Ian Larkin, Lamar Pierce. 2018. Doing well by making well: The impact of corporate wellness programs on employee productivity. *Management Science* **64**(11) 4967–4987.

Holden, Steinar, Fredrik Wulfsberg. 2009. How strong is the macroeconomic case for downward real wage rigidity? *Journal of Monetary Economics* **56**(4) 605–615.

Hölmstrom, Bengt. 1979. Moral hazard and observability. *The Bell Journal of Economics* **10**(1) 74–91.

Jensen, Michael C, William H Meckling. 1976. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics* **3**(4) 305–360.

Kahn, Shulamit. 1997. Evidence of nominal wage stickiness from microdata. *The American Economic Review* **87**(5) 993–1008.

Kahn-Lang, Ariella, Kevin Lang. 2018. The promise and pitfalls of differences-in-differences: Reflections on '16 and pregnant' and other applications. *NBER Working Paper Number 24857* .

Kahneman, Daniel, Jack L Knetsch, Richard Thaler. 1986. Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review* 728–741.

Kaur, Supreet. 2014. Nominal wage rigidity in village labor markets. Tech. rep., National Bureau of Economic Research.

Krueger, Alan B, Alexandre Mas. 2004. Strikes, scabs, and tread separations: Labor strife and the production of defective bridgestone/firestone tires. *Journal of political Economy* **112**(2) 253–289.

Krueger, Miriam, Guido Friebel. 2018. A pay change and its long-term consequences. Tech. rep.

Kube, Sebastian, Michel André Maréchal, Clemens Puppe. 2013. Do wage cuts damage work morale? Evidence from a natural field experiment. *Journal of the European Economic Association* **11**(4) 853–870.

Larkin, Ian, Stephen Leider. 2012. Incentive schemes, sorting, and behavioral biases of employees: Experimental evidence. *American Economic Journal: Microeconomics* **4**(2) 184–214.

Larkin, Ian, Lamar Pierce, Francesca Gino. 2012. The psychological costs of pay-for-performance: Implications for the strategic compensation of employees. *Strategic Management Journal* **33**(10) 1194–1214.

Lazear, Edward P. 2000. Performance pay and productivity. *American Economic Review* **90**(5) 1346–1361.

Lazear, Edward P, Kathryn L Shaw, Christopher Stanton. 2016. Making do with less: Working harder during recessions. *Journal of Labor Economics* **34**(S1) S333–S360.

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

Lemieux, Thomas, W Bentley MacLeod, Daniel Parent. 2009. Performance pay and wage inequality. *The Quarterly Journal of Economics* **124**(1) 1–49.

MacKinnon, James G, Matthew D Webb. 2018. Wild bootstrap randomization inference for few treated clusters. Tech. rep., Queens Economics Department.

Manning, Alan. 2003. *Monopsony in motion: Imperfect competition in labor markets*. Princeton University Press.

Mas, Alexandre. 2006. Pay, reference points, and police performance. *The Quarterly Journal of Economics* **121**(3) 783–821.

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

Mascarenhas, Briance, David A Aaker. 1989. Strategy over the business cycle. *Strategic Management Journal* **10**(3) 199–210.

Molloy, J.C., J.B. Barney. 2015. Who captures the value created with human capital? A market-based view. *Academy of Management Perspectives* **29**.

Morris, Shad S, Sharon A Alvarez, Jay B Barney, Janice C Molloy. 2017. Firm-specific human capital investments as a signal of general value: Revisiting assumptions about human capital and how it is managed. *Strategic Management Journal* **38**(4) 912–919.

Obloj, Tomasz, Metin Sengul. 2012. Incentive life-cycles: Learning and the division of value in firms. *Administrative Science Quarterly* **57**(2) 305–347.

Obloj, Tomasz, Todd Zenger. 2017. Organization design, proximity, and productivity responses to upward social comparison. *Organization Science* **28**(1) 1–18.

Sandvik, Jason, Richard Saouma, Nathan Seegert, Christopher Stanton. 2020. Workplace knowledge flows .

Stafford, Tess M. 2015. What do fishermen tell us that taxi drivers do not? An empirical investigation of labor supply. *Journal of Labor Economics* **33**(3) 683–710.

Zenger, Todd R. 1992. Why do employers only reward extreme performance? Examining the relationships among performance, pay, and turnover. *Administrative Science Quarterly* 198–219.

Figures and Tables

Figure 1: Revenue Transfer Price Changes Within the Commission Schedule

.

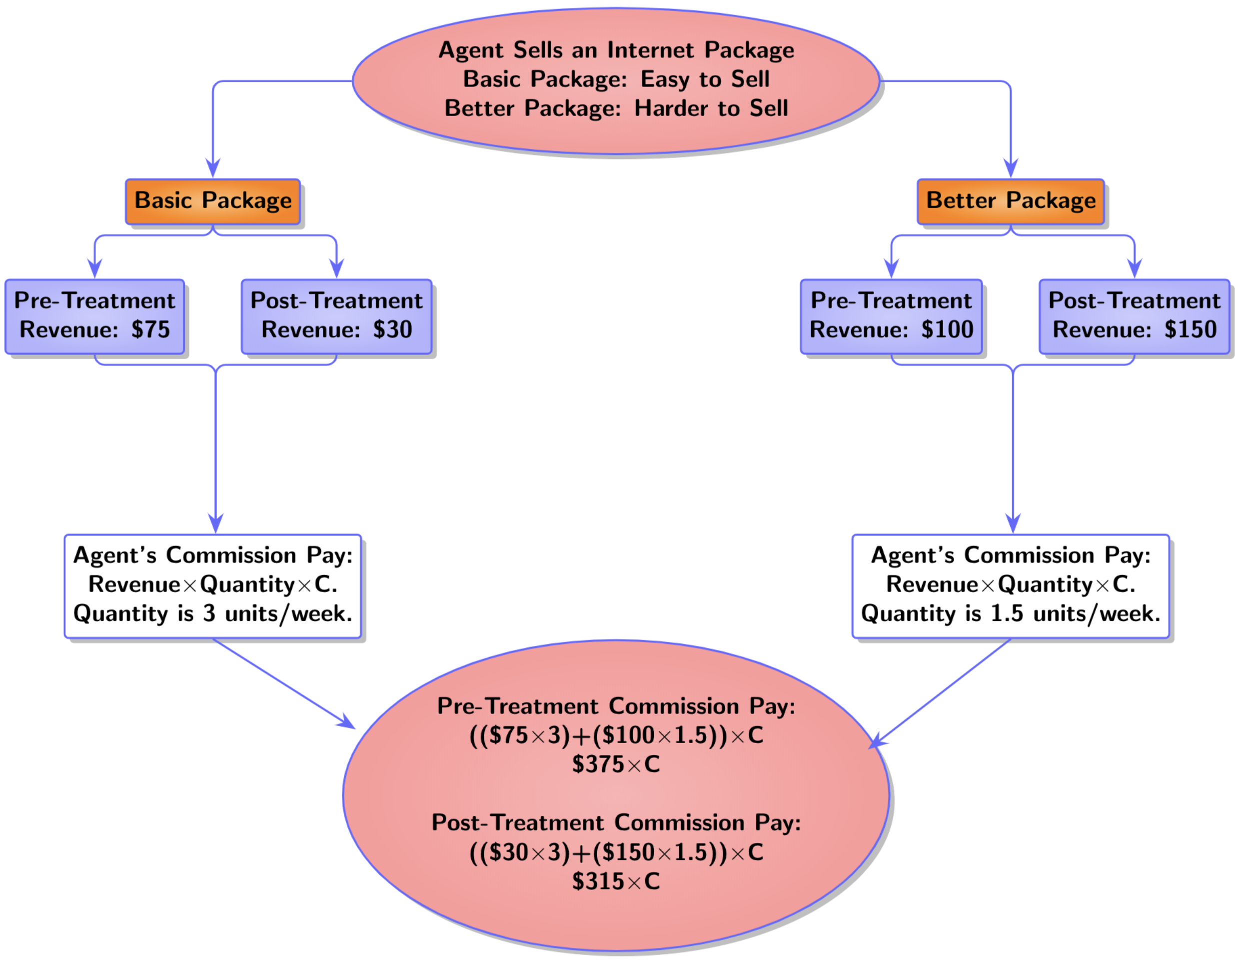

Notes: This figure provides an example of changes in the commission schedule for two different types of internet packages. The pre- and post-treatment revenue transfer prices for the basic package are displayed in the left branch. The pre- and post-treatment revenue transfer prices for the better package are displayed in the right branch. A basic package is easier to sell than a better package, captured by the higher quantity of sales per agent-week, 3 vs. 1.5. The agent's commission rate, C, is multiplied by the product of the revenue transfer price and quantity sold to determine the amount of commission pay the agent receives for selling a particular package.

Figure 2: Commissions in Division 1 and the Control Divisions

(a) **Commission Level Time Series in Different Periods**

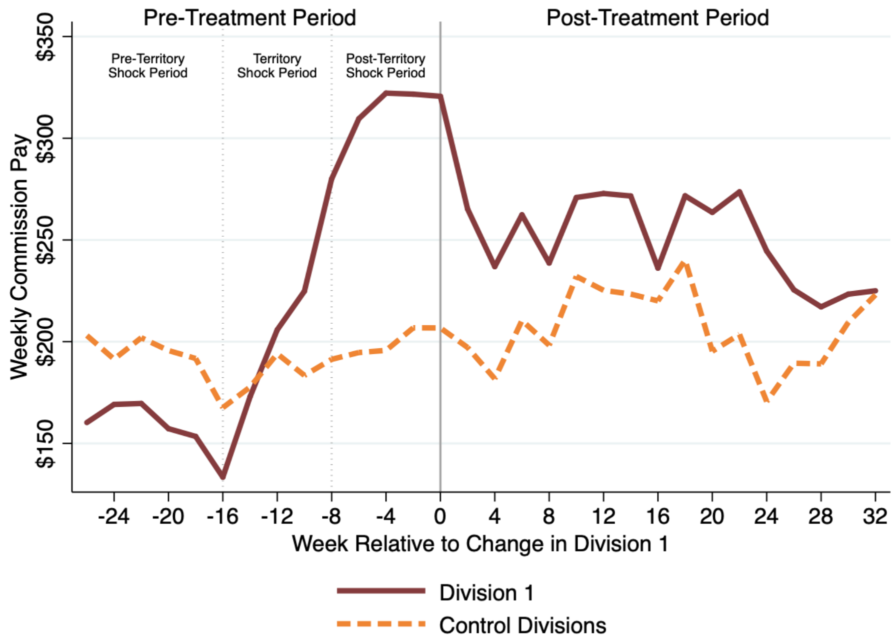

(b) **Trends in Commission Levels and Total Call Volumes**

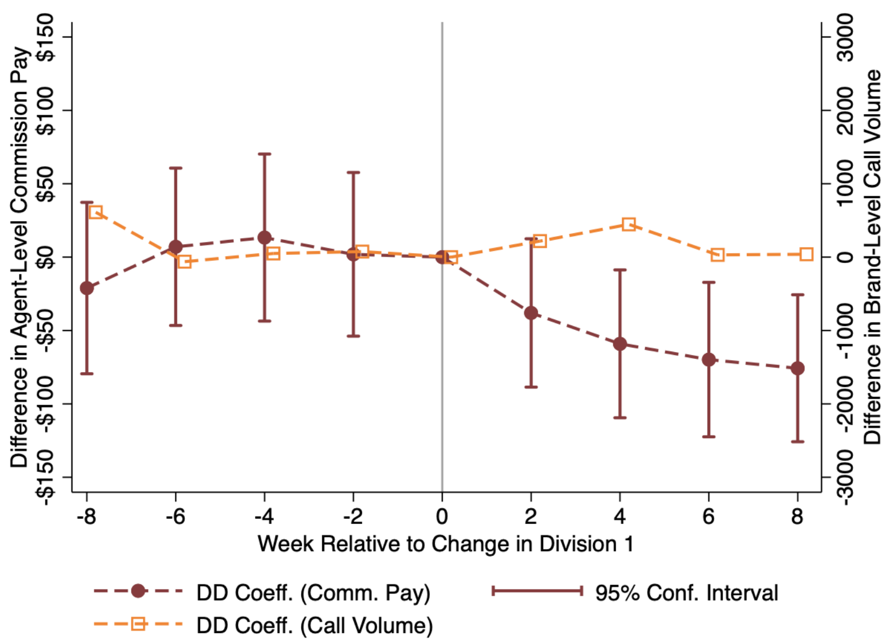

Notes: Figure (a) plots the average weekly commission pay levels for agents in Division 1 and the control divisions. The solid vertical line corresponds to the two weeks immediately before the week of the commission schedule changes in Division 1. Figure (b) plots the difference-in-differences coefficients that capture differential trends in commission pay levels and total call volume between Division 1 and the control divisions. We fail to reject the null hypothesis that the coefficients in Weeks -8 to 0 jointly equal zero (p = 0.82).

Figure 3: Reported Changes in Commissions and Effort

(a) **Expected Change in Commissions**

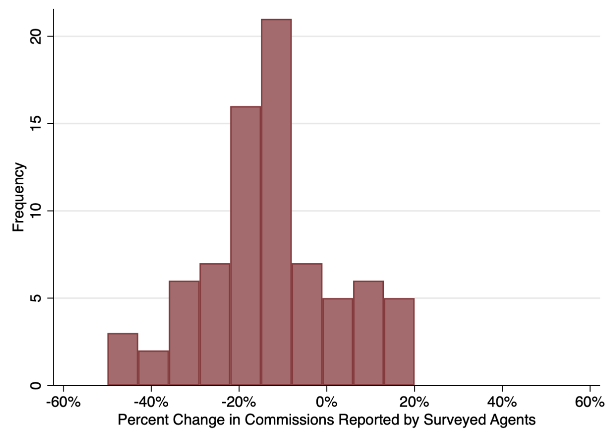

(b) **Effort Necessary to Maintain Earnings**

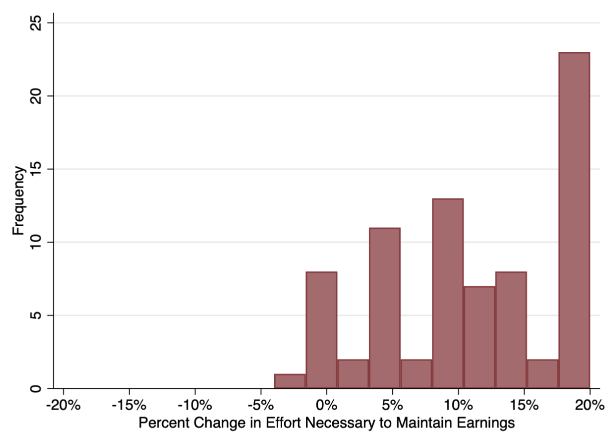

(c) **Actual Change in Effort**

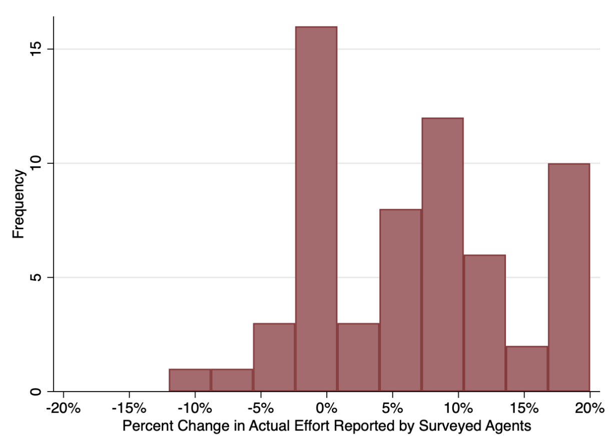

Notes: Histogram (a) displays survey responses to a question asking agents what their expected change in commissions would be due to the commission schedule changes. Histogram (b) displays responses to a question asking how much agents' effort would need to change to maintain their normal level of earnings. Histogram (c) displays responses to a question of what changes in effort workers actually planned to make. See Section 2.6 for more details.

Figure 4: Survival Rates By Productivity in Division 1 and the Control Divisions

(a) **Survival Rates Relative to Month -5**

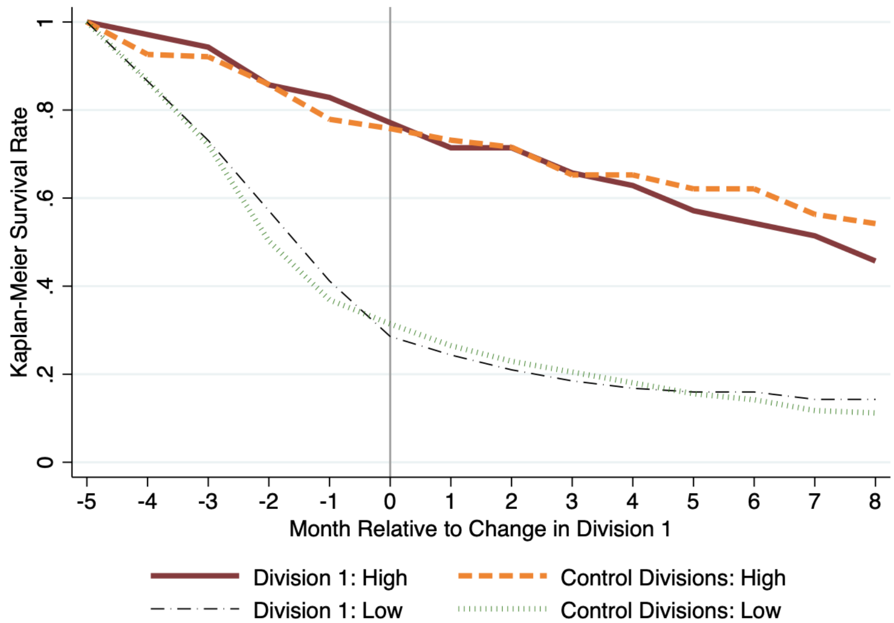

(b) **Survival Rates Relative to Month 0**

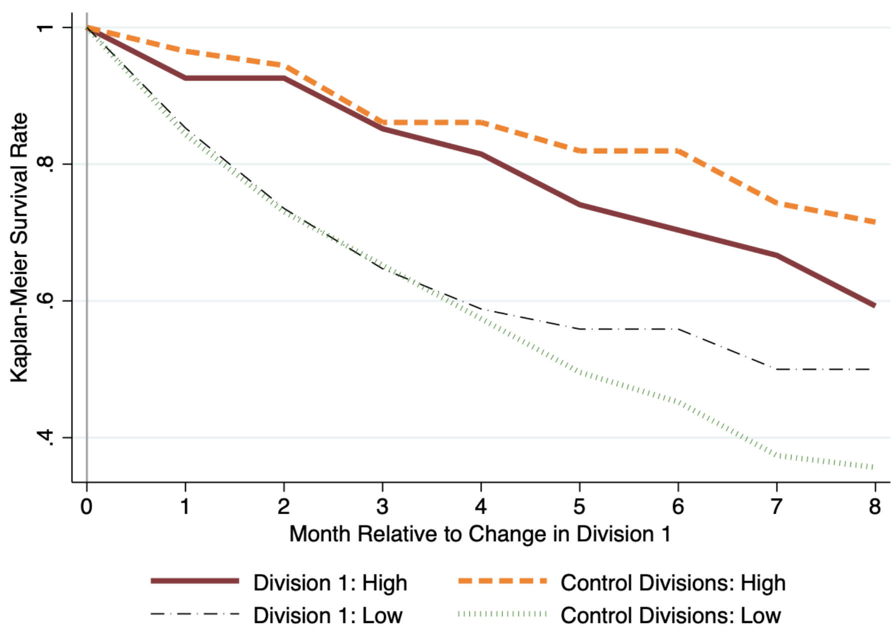

Notes: These figures plot Kaplan-Meier survival rates over time. The survival rate estimator considers a starting point and then, from that time, displays the fraction of agents that remain at the firm. Because turnover can be lumpy, with multiple exits in some weeks and no exits in others, we aggregate survival rates to the monthly level. The sample is split by high and low performers based on whether agents' adjusted worker fixed effects are above or below the median within their division.

Figure 5: Trends in Proxies for Effort Supply and Effort Demand

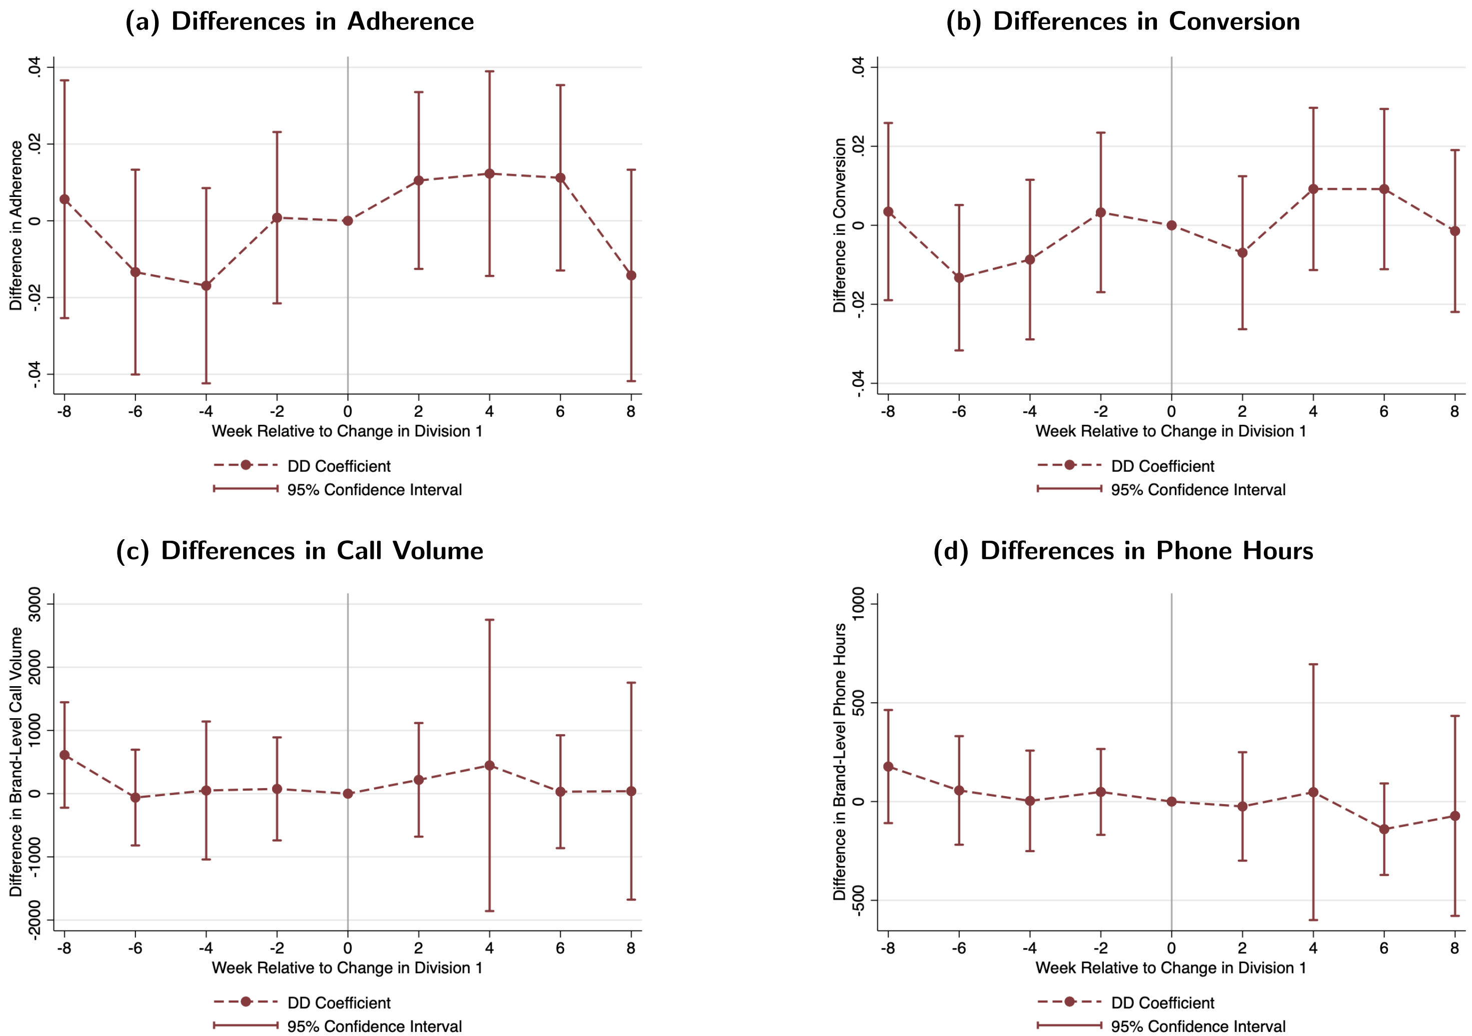

Notes: The coefficients in these figures are estimates of $\delta_{i,t}$ from Equation (1), using different outcome variables of interest. Adherence and conversion are the two proxies for an agent's supply of effort. Call volume and phone hours are the proxies for customers' demand for worker effort. To improve the readability of these figures, we aggregate data into bi-weekly clusters. The p-values of tests that the Week -8 to Week 0 point estimates are jointly equal to zero are 0.44, 0.31, 0.36, and 0.76 for adherence, conversion, call volume, and phone hours, respectively.

Figure 6: Adjusted Worker Fixed Effects Before and After the Compensation Changes

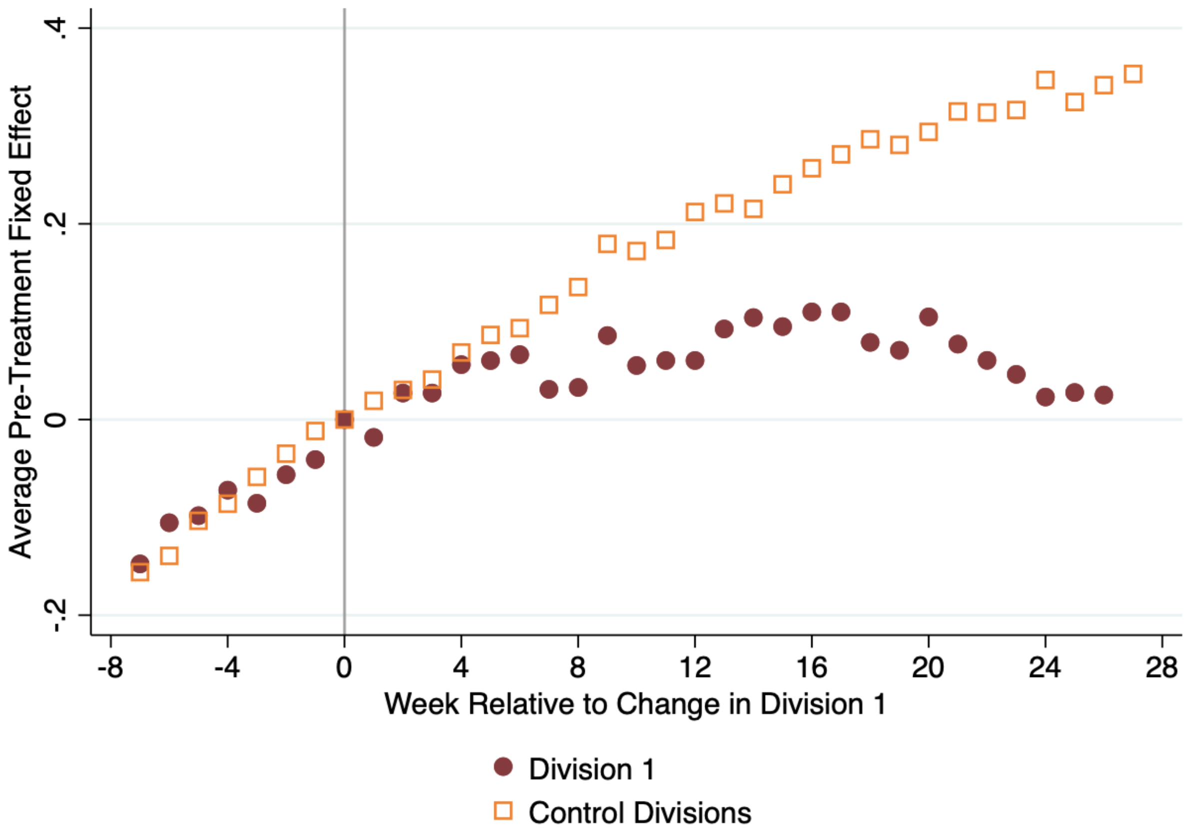

Notes: This figure plots the average adjusted worker fixed effects for Division 1 and the control divisions. These adjusted worker fixed effects are for those who are present at the firm prior to the commission schedule changes. Adjusted worker fixed effects are calculated from a regression of log commissions on worker dummy variables, division-by-week dummy variables, and a cubic spline in tenure. We then correct for sampling variation using the method in Lazear et al. (2015). The series are normalized to correspond at the announcement date, which is depicted by the vertical line.

Figure 7: Placebo Treatment Tests

(a) **Agents and Treatment Date Randomized**

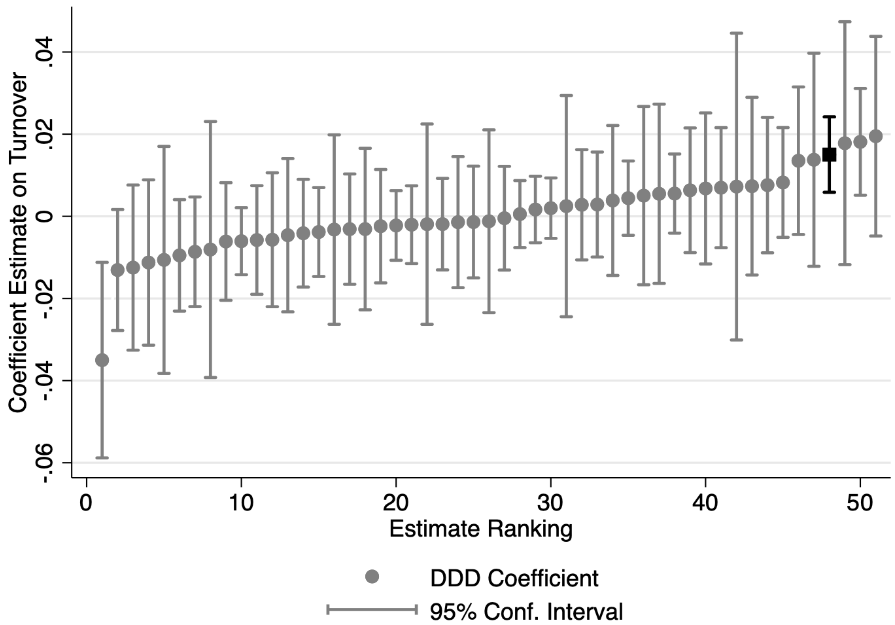

(b) **Divisions and Treatment Date Randomized**

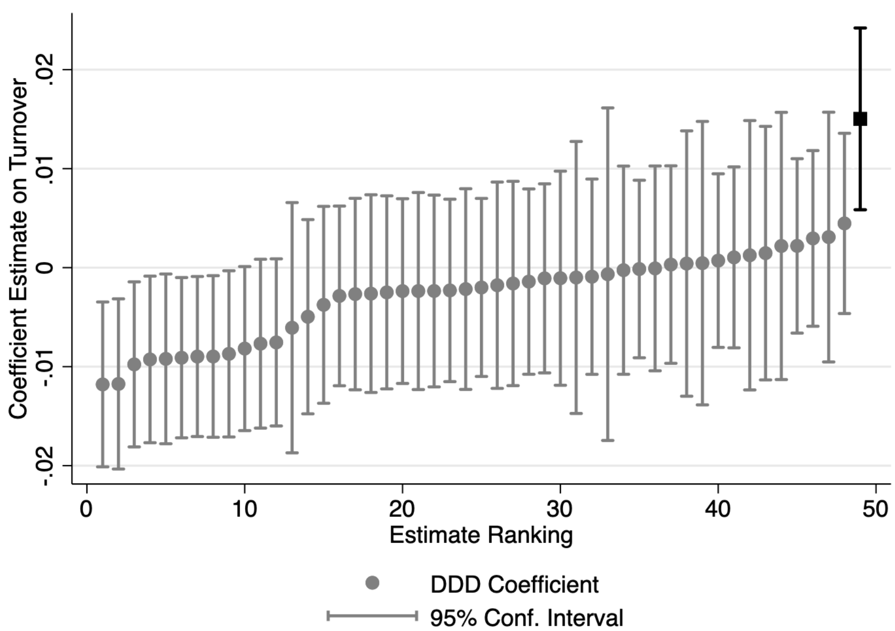

Notes: Figure (a) plots fifty placebo simulations (gray dots) for the turnover response estimation using randomized treatment groups and treatment dates. The black square shows the actual result from Column 2 of Table 3. For each placebo simulation, we randomly select 180 agents to constitute the treated division, with the other agents making up the control group. We then randomly choose an intervention week between September 1st, 2016 and January 31st, 2017. This process is similar to that used in Gubler et al. (2018). Figure (b) plots similar placebo simulations for the turnover response estimation using different divisions as the "treated" division, while the other divisions (including the actual treated division) make up the control group.

Figure 8: Changes in Reported Sentiment Towards the Firm for Division 1

(a) **Perceptions of Firm Fairness**

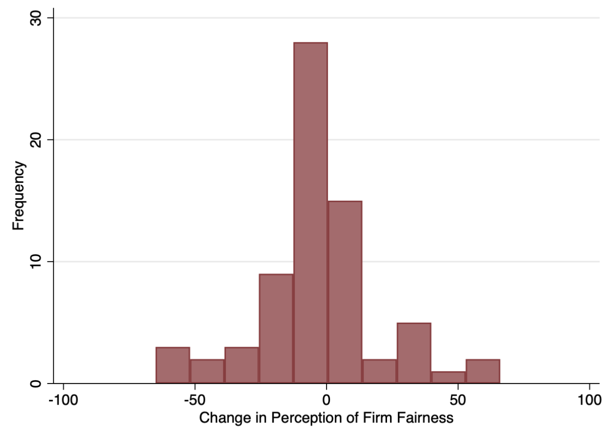

(b) **Propensity to Refer Friends to the Firm**

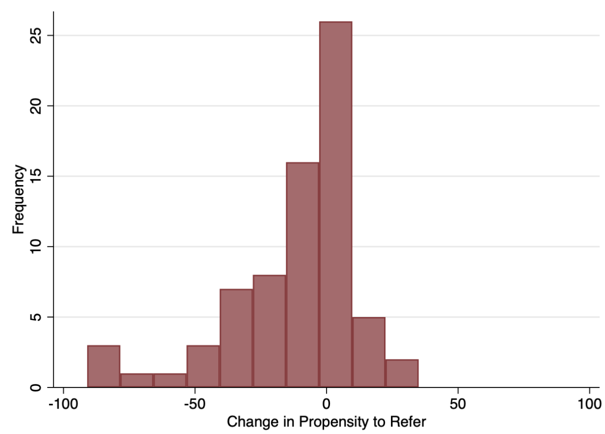

(c) **Perceptions of Future Promotion**

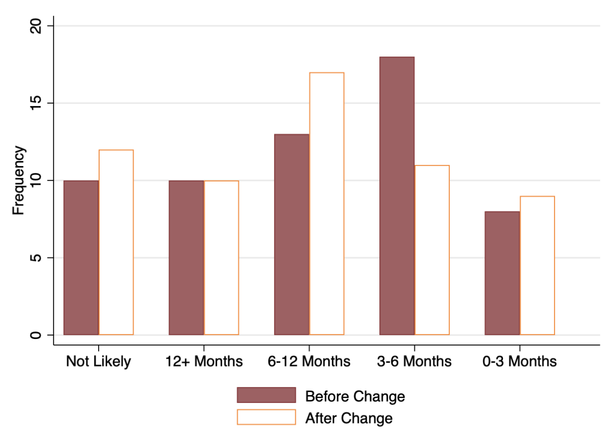

Notes: These figures plot the changes in Division 1 agents' sentiment before and after the commission schedule changes. The unit of analysis is a sales agent who is present in both survey waves. See Section 2.6 for details about question wording.

Table 1: Summary Statistics Pre- and Post-Treatment

	Eight Weeks Pre-Treatment			All Weeks Post-Treatment		
	Control Divisions	Treated Division 1	Treated Division 2	Control Divisions	Treated Division 1	Treated Division 2
	(1)	(2)	(3)	(4)	(5)	(6)
Commission	200.91	318.39	502.68	206.92	249.71	308.56
	(184.79)	(283.59)	(333.35)	(189.57)	(230.30)	(226.25)
Adherence	0.80	0.83	0.79	0.81	0.82	0.80
	(0.12)	(0.12)	(0.15)	(0.12)	(0.12)	(0.14)
Conversion	0.26	0.33	0.29	0.25	0.32	0.29
	(0.10)	(0.09)	(0.12)	(0.09)	(0.10)	(0.13)
Log RPC_{Old}	4.11	4.19		4.14	4.09	
	(0.47)	(0.51)		(0.49)	(0.52)	
Log RPC_{New}	4.11	3.96		4.14	3.90	
	(0.47)	(0.50)		(0.49)	(0.50)	
Phone Hours	19.89	20.71	17.41	20.41	20.18	15.92
	(7.59)	(7.32)	(6.25)	(8.38)	(7.88)	(7.26)
Total Calls	62.35	71.13	49.19	64.58	69.88	45.30
	(26.79)	(27.76)	(19.52)	(28.99)	(29.50)	(21.31)
Tenure (days)	356.56	369.02	672.98	450.36	399.81	608.06
	(419.52)	(389.51)	(558.98)	(505.66)	(411.39)	(594.50)
Age	25.84	25.18	29.71	26.20	25.99	28.33
	(7.15)	(6.50)	(8.64)	(7.32)	(7.75)	(7.98)
Single	0.52	0.68	0.44	0.38	0.44	0.33
	(0.50)	(0.47)	(0.50)	(0.49)	(0.50)	(0.47)
White	0.71	0.72	0.60	0.68	0.66	0.62
	(0.45)	(0.45)	(0.49)	(0.47)	(0.48)	(0.49)
Male	0.70	0.73	0.70	0.73	0.73	0.64
	(0.46)	(0.44)	(0.46)	(0.44)	(0.44)	(0.48)
Agent-Weeks	4,024	867	357	13,817	3,474	950
Agents	632	138	51	874	234	89

Notes: This table presents summary statistics for the control divisions and Division 1 that come from data eight weeks prior to the Division 1 commission schedule changes (Columns 1 and 2) and all weeks after the Division 1 changes (Columns 4 and 5). The summary statistics for Division 2 are measured using data from the eight weeks prior to (Column 3) and all weeks after the Division 2 commission schedule changes (Column 6). The *Commission* measure is average weekly commissions; *Adherence* is a measure of schedule adherence, which captures the amount of time an agent is available to take calls; *Conversion* is the probability of having positive sales revenue on a given call; *Log* RPC_{Old} measures an agent's revenue-per-call (RPC) if the commission schedule had *not* changed; *Log* RPC_{New} measures an agent's revenue-per-call (RPC) if the commission schedule had always been at the new levels. *Phone Hours* capture the amount a time an agent spends talking with customers; and *Total Calls* is the number of calls fielded by an agent each week. Data limitations prevent us from measuring *Log* RPC_{Old} and *Log* RPC_{New} for Division 2. Standard deviations are reported in parentheses.

Table 2: Summary Statistics for Division 1 By Productivity Level

	Tercile of Adjusted Worker Fixed Effects		
	Bottom Third	Middle Third	Top Third
	(1)	(2)	(3)
Commission	170.88	298.02	476.63
	(167.92)	(210.59)	(342.00)
Predicted Pct Δ Commission Post-Treatment	-0.18	-0.18	-0.17
	(0.04)	(0.03)	(0.03)
Adherence	0.81	0.84	0.84
	(0.16)	(0.11)	(0.09)
Conversion	0.29	0.34	0.37
	(0.09)	(0.08)	(0.09)
Log RPC_{Old}	3.93	4.24	4.37
	(0.59)	(0.41)	(0.41)
Log RPC_{New}	3.69	4.00	4.15
	(0.59)	(0.40)	(0.40)
Phone Hours	18.41	21.26	22.99
	(7.32)	(7.02)	(6.26)
Total Calls	64.95	72.42	77.55
	(25.86)	(25.75)	(25.00)
Tenure (days)	149.69	215.07	691.07
	(67.70)	(110.89)	(485.51)
Age	23.12	23.67	28.14
	(4.78)	(3.77)	(8.48)
Single	0.79	0.76	0.58
	(0.41)	(0.43)	(0.50)
White	0.73	0.70	0.73
	(0.44)	(0.46)	(0.44)
Male	0.68	0.78	0.73
	(0.47)	(0.41)	(0.45)
Survey Response to Firm Fairness	0.57	0.48	0.35
	(0.50)	(0.50)	(0.48)
Survey Response to Referral Likelihood	0.73	0.58	0.52
	(0.45)	(0.49)	(0.50)
Survey Response to Promotion Likelihood	0.55	0.83	0.59
	(0.50)	(0.38)	(0.49)
Agent-Weeks	249	292	297
Agents	40	40	40

Notes: This table presents summary statistics for Division 1 using data eight weeks prior to the commission schedule changes. Each column represents an approximate tercile of the distribution of adjusted worker fixed effects in the pre-treatment period. Adjusted worker fixed effects are calculated from a regression of log commissions on worker dummy variables, division-by-week dummy variables, and a cubic spline in tenure. We then correct for sampling variation using the method in Lazear et al. (2015). We are not able to estimate adjusted worker fixed effects for every agent, resulting in slightly smaller agent and agent-week counts compared to those in Table 1. The *Predicted Percentage* Δ *Commission Post-Treatment* is a calculation of how total commissions would decline for each agent due to the commission schedule changes

Table 3: Linear Probability Model Estimates of Turnover Responses

	Last Week in Firm				
	(1)	(2)	(3)	(4)	(5)
Treated x Post x Prod	0.021**	0.015**	0.016*	0.013**	0.012*
	(0.007)	(0.005)	(0.007)	(0.005)	(0.005)
Treated x Post	-0.006	-0.006		-0.002	
	(0.004)	(0.007)		(0.010)	
Treated x Placebo x Prod	-0.006		-0.002		
	(0.004)		(0.004)		
Treated x Placebo	0.000				
	(0.004)				
Week Fixed Effects	✓	✓		✓	
Division x Week-of-Year Fixed Effects		✓			
Week x Division Fixed Effects			✓		✓
Post-Territory Shock Period				✓	✓
Observations	51,497	51,497	51,497	19,689	19,689
Mean Turnover Probability in Division 1			0.037		
p-value on Treated x Post x Prod	0.018	0.084	0.100	0.044	0.043
p-value on Treated x Post	0.480	0.315		0.244	

Notes: The dependent variable is an indicator that equals one if it is the worker's last week at the firm. The sample includes all current employees in Division 1 and the control divisions with non-missing data. Estimates come from a linear probability model that captures changes in the turnover probability for the existing workforce. Each model includes a 5th order polynomial for workers' tenure to account for a potentially arbitrary baseline relationship between tenure and turnover. *Prod* refers an agent's sales z-score, which is the standardized measure of an agent's pre-treatment productivity estimated as their adjusted worker fixed effect according to the procedure in Lazear et al. (2015). For additional details, see Section 2.5. The specification in Column 2 includes division by week-of-year fixed effects to account for seasonality. The specification in Column 3 includes week by division fixed effects. Columns 4 and 5 use a shortened pre-treatment period that only includes the weeks of data after the territory shock period. *Placebo* is an indicator for the date 52 weeks prior to the treatment date. Two forms of inference are presented, one using standard errors clustered by manager (in parentheses) and the second using p-values with division-level clusters (see the final two lines) computed using the wild cluster bootstrap randomization inference procedure in MacKinnon and Webb (2018). We use the t-statistic version of the procedure that imposes the null hypothesis.

$*p < 0.1$; $**p < 0.05$; $***p < 0.01$

Table 4: Estimates of Effort Responses

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Adherence to Schedule							
Treated x Post	0.006	0.012	0.047**	0.032	0.032	0.048**	
	(0.009)	(0.010)	(0.015)	(0.022)	(0.022)	(0.015)	
Treated x Post x Prod						-0.006	-0.006
						(0.005)	(0.005)
Observations	8,119	8,119	8,119	7,064	3,706	8,119	8,119
Panel B: Conversion Rate							
Treated x Post	0.004	0.002	0.010	0.006	0.008	0.016*	
	(0.007)	(0.005)	(0.007)	(0.008)	(0.007)	(0.007)	
Treated x Post x Prod						-0.020***	-0.020***
						(0.005)	(0.005)
Observations	8,283	8,283	8,283	6,903	3,743	8,283	8,283
Panel C: Log RPC at Old Prices							
Treated x Post	-0.025	-0.039	0.005	-0.016	0.005	0.022	
	(0.031)	(0.025)	(0.033)	(0.036)	(0.048)	(0.032)	
Treated x Post x Prod						-0.041	-0.042
						(0.024)	(0.024)
Observations	9,229	9,229	9,229	7,840	4,126	9,229	9,229
Panel D: Log RPC at New Prices							
Treated x Post	0.027	0.007	0.027	0.002	0.016	0.046	
	(0.032)	(0.025)	(0.036)	(0.039)	(0.054)	(0.035)	
Treated x Post x Prod						-0.048	-0.049
						(0.024)	(0.025)
Observations	9,229	9,229	9,229	7,840	4,126	9,229	9,229
Week Fixed Effects	✓	✓	✓	✓	✓	✓	
Agent Fixed Effects		✓	✓	✓	✓	✓	✓
Division Trend Controls			✓	✓	✓	✓	
Week x Division Fixed Effects							✓
Re-Weighted				✓			
Balanced Sample					✓		

Notes: The sample includes all current employees in Division 1 and the control divisions with non-missing data. The models in Columns 1–6 include fixed effects for week, division, and office location. All models include cubic splines for tenure and a cubic polynomial in age. The OLS regression in Column 1 includes dummies for ethnicity, gender, and marital status. The specifications in Columns 2–7 include individual fixed effects. Columns 3–6 include division-specific trend controls. The specification in Column 4 uses a re-weighting estimator based on the propensity score for being in Division 1 (see Appendix 7). The balanced panel in Column 5 restricts to workers who are present prior to July, 2016 and after April, 2017. Columns 6 and 7 consider heterogeneous responses based on worker productivity, and Column 7 omits week fixed effects and division-specific trend controls and instead includes week by division fixed effects. Differing numbers of observations across panels reflect differences in data availability. The sample used restricts to eight weeks of pre-treatment data and eight weeks of post-treatment data. The results are similar when all available pre- and post-treatment data is used (See Table A.2). Reported standard errors are clustered by manager.

$*p < 0.1$; $**p < 0.05$; $***p < 0.01$

Table 5: Sentiment Descriptive Statistics

Panel A: Changes in Sentiment by Pre-Treatment Sentiment

	Pre-Treatment Sentiment				
	All	0%–33%	33%–66%	66%–100%	Diff.
	(1)	(2)	(3)	(4)	(4)–(2)
ΔFairness Perceptions	-1.43	14.17**	-12.25**	-5.74**	-19.91***
	(2.74)	(5.14)	(4.51)	(2.37)	(5.66)
ΔReferral Likelihood	-12.51***	-9.09*	-11.14**	-16.36***	-7.27
	(2.90)	(4.87)	(4.44)	(5.38)	(7.40)
ΔPromotion Prospects	-0.17**	0.09	0.04	-0.57***	-0.66***
	(0.07)	(0.06)	(0.11)	(0.14)	(0.17)
Agents	70	23	24	23	

Panel B: Changes in Sentiment by Pre-Treatment Productivity

	Pre-Treatment Productivity (Z-Score)				
	All	0%–33%	33%–66%	66%–100%	Diff.
	(1)	(2)	(3)	(4)	(4)–(2)
ΔFairness Perceptions	-1.43	-2.74	-3.92	2.48	5.22
	(2.74)	(4.35)	(5.12)	(4.80)	(6.48)
ΔReferral Likelihood	-12.51***	-5.30	-11.52**	-19.77***	-14.46**
	(2.90)	(3.32)	(4.85)	(5.93)	(7.04)
ΔPromotion Prospects	-0.17**	-0.04	-0.33**	-0.14	-0.10
	(0.07)	(0.10)	(0.16)	(0.12)	(0.16)
Agents	70	23	24	23	

Notes: This table documents the changes in the self-reported sentiment levels of Division 1 agents from before to after the commission schedule changes. In Panel A, the data is split across terciles of pre-treatment sentiment where Column 2 contains agents with the lowest sentiment and Column 4 contains agents with the highest sentiment. Panel B splits the data based on terciles of pre-treatment agent productivity. The results of difference-in-means tests between Columns 4 and 2 are reported in the far right column. Standard errors of means are reported in parentheses.

*$p < 0.1$; **$p < 0.05$; ***$p < 0.01$

Table 6: Heterogeneous Turnover Responses Based on Worker Sentiment

	Last Week in Firm				
	(1)	(2)	(3)	(4)	(5)
Treated x Post x Prod	0.012	0.014**	0.010	0.013*	0.013*
	(0.006)	(0.005)	(0.007)	(0.005)	(0.005)
Treated x Post x Firm Fair	0.009	0.009	0.009	0.013	0.012
	(0.009)	(0.009)	(0.009)	(0.008)	(0.008)
Treated x Post x High Refer	0.005	0.004	0.004	0.005	0.004
	(0.007)	(0.006)	(0.007)	(0.006)	(0.007)
Treated x Post x Promotion	-0.011	-0.011	-0.011	-0.011	-0.010
	(0.011)	(0.011)	(0.011)	(0.012)	(0.012)
Week Fixed Effects	✓	✓		✓	
Division x Week-of-Year Fixed Effects		✓			
Week x Division Fixed Effects			✓		✓
Shortened Pre-Period				✓	✓
Observations	51,497	51,497	51,497	19,689	19,689
Mean Turnover Probability in Division 1			0.037		
p-value on Treated x Post x Prod	0.039	0.037	0.157	0.069	0.056

Notes: The dependent variable is an indicator that equals one if it is the worker's last week at the firm. The sample includes all current employees in Division 1 and the control divisions with non-missing data. Estimates come from a linear probability model that captures changes in the turnover probability for the existing workforce. Each model includes a 5th order polynomial for workers' tenure to account for a potentially arbitrary baseline relationship between tenure and turnover. *Prod* refers an agent's sales z-score, which is the standardized measure of an agent's pre-treatment productivity estimated as their adjusted worker fixed effect according to the procedure in Lazear et al. (2015). For additional details, see Section 2.5. We separately interact the treatment indicator with indicators for high firm fairness perceptions, high referral likelihood, and a belief that promotion is likely. An agent's firm fairness perception is marked as high if it is above the median value. Referral likelihood is marked is high if it is above the median value. If an agent says they are likely to be promoted in the future, their promotion likelihood indicator equals one. The specification in Column 2 includes division by week-of-year fixed effects to account for seasonality. The specification in Column 3 includes week by division fixed effects. Columns 4 and 5 use a shortened pre-treatment period that only includes the weeks of data after the territory shock period. Two forms of inference are presented, one using standard errors clustered by manager (see parentheses) and the second using p-values with division-level clusters (see the final two lines) computed using the wild cluster bootstrap randomization inference procedure in MacKinnon and Webb (2018). We use the t-statistic version of the procedure that imposes the null hypothesis.

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Table 7: Heterogeneous Effort Responses Based on Worker Sentiment

	Adherence to Schedule	Conversion Rate	Log RPC at Old Prices	Log RPC at New Prices
	(1)	(2)	(3)	(4)
Treated x Post x Prod	-0.007	-0.021***	-0.046*	-0.046*
	(0.006)	(0.005)	(0.019)	(0.018)
Treated x Post x Firm Fair	-0.008	0.004	-0.001	0.005
	(0.016)	(0.010)	(0.064)	(0.060)
Treated x Post x High Refer	-0.003	-0.013	-0.067	-0.062
	(0.012)	(0.008)	(0.049)	(0.044)
Treated x Post x Promotion	-0.001	0.005	0.070	0.085
	(0.013)	(0.011)	(0.062)	(0.051)
Agent Fixed Effects	✓	✓	✓	✓
Week x Division Fixed Effects	✓	✓	✓	✓
Observations	8,119	8,283	9,229	9,229

Notes: The sample includes all current employees in Division 1 and the control divisions with non-missing data. All models include agent fixed effects and fixed effects division and office location. To account for experience effects, all models include cubic splines for tenure with the firm and a cubic polynomials in age. Each specification also includes week by division fixed effects. We separately interact the treatment indicator with indicators for high firm fairness perceptions, high referral likelihood, and a belief that promotion is likely. An agent's firm fairness perception is marked as high if it is above the median value. Referral likelihood is marked is high if it is above the median value. If an agent says they are likely to be promoted in the future, their promotion likelihood indicator equals one. Differing numbers of observations across columns reflect differences in data availability. The sample used restricts to eight weeks of pre-treatment data and eight weeks of post-treatment data. The results are similar when all available pre- and post-treatment data is used. Reported standard errors are clustered by manager.

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

6 Motivating Framework

We motivate the analysis with a simple model of heterogeneous agent responses to commission changes as a function of productivity differences. Heterogeneous responses are difficult to sign without assumptions, making them empirical objects of analysis. We then consider how turnover changes across the productivity distribution affect profitability.

Let e_i denote agent i's sales effort and assume further that his sales revenue, y_i is given by $y_i = \theta_i e_i + \epsilon$ where $\theta_i > 0$ is the agent's productivity or type, and ϵ is mean-zero noise. To simplify the exposition, all agents are assumed to be risk-neutral and collect a linear share of their revenues, R, in addition to a common fixed wage, α, such that we can represent agent i's expected utility by $U(\alpha, R, \theta_i, e_i) = \alpha + R\theta_i e_i - c(e_i)$.

The cost of effort function $c(\cdot)$ is strictly increasing and convex, with $c(0) = c'(0) = 0$. Let e^* denote the unique solution to the agent's problem:

$$e^* = \underset{e}{\operatorname{argmax}} R\theta e - c\left(e\right)$$

such that agent i's value function evaluated at e^* can be expressed as $V(\alpha, R; \theta_i)$.

The optimal effort, e_i^*, is strictly positive, as $c'(0) = c(0) = 0 < R$. Accordingly, the function U has strictly increasing differences in e_i and R, as well as in e_i and θ_i. By application of Topkis's Theorem, both $\frac{\partial e_i^*}{\partial R}$ and $\frac{\partial e_i^*}{\partial \theta_i}$ are themselves strictly positive. However the heterogeneous effort responses across agents of different pre-treatment productivity are captured by $\frac{\partial^2 e_i^*}{\partial R \partial \theta_i}$, which we cannot sign without additional assumptions.

Proposition 1. *An agent's change in effort with respect to commissions is increasing in agent productivity, θ, as long as c''' is sufficiently small.*

Proof. Proof of Proposition 1. The goal is to show that the marginal effect of productivity, θ, on agent i's effort response to a change in commissions is directly proportional to the curvature of the agents' cost function. Specifically:

$$\frac{\partial^2 e_i^*}{\partial R \partial \theta_i} \propto \left(c''\left(e_i^*\right)\right)^2 - c'''\left(e_i^*\right) R\theta_i. \tag{7}$$

To prove Equation (7), we begin with the first order condition $R\theta_i = c'\left(e_i^*\right)$. Differentiating both sides with respect to R yields $\theta_i = c''\left(e_i^*\right) \frac{\partial e_i^*}{\partial R}$. Differentiating again by θ_i yields: $1 = c'''\left(e_i^*\right) \frac{\partial e_i^*}{\partial R} \frac{\partial e_i^*}{\partial \theta_i} + c''\left(e_i^*\right) \frac{\partial^2 e_i^*}{\partial R \partial \theta_i}$, substituting the earlier terms and rearranging yields:

$$\frac{\partial^2 e_i^*}{\partial R \partial \theta_i} = \frac{\left(c''\left(e_i^*\right)\right)^2 - c'''\left(e_i^*\right) R\theta_i}{\left(c''\left(e_i^*\right)\right)^3},$$

which completes the proof as $c'' > 0$ by assumption. □

When the agent's costs follow a standard power function, e.g. $c(e) = e^n/n$, the expression characterizing $\frac{\partial^2 e_i^*}{\partial R \partial \theta_i}$ is strictly positive. We conclude that in most standard settings, agents have weakly larger effort responses to commission changes as their type increases. Accordingly, we treat effort changes by agent type as an empirical question, and instead turn our attention to turnover effects.

Beginning with the seminal work of Burdett and Mortensen (1998), the job ladder model has been used extensively to capture worker mobility. The standard model maintains an attrition (quit) rate of $Q(w) = \delta + \lambda[1 - F(w)]$, where $\delta > 0$ captures exogenous job destruction, $\lambda \in [0, 1]$ captures search frictions via an arrival rate of outside job opportunities, and w is a random variable with density $f(\cdot)$ and associated CDF $F(\cdot)$ capturing the distribution of *fixed* wage offers to the agent from outside firms. We define the agent's reservation wage, $w(\theta_i)^*$, as the lowest fixed-wage yielding an expected utility of $V(\alpha, R; \theta_i)$.[53] To simplify the ensuing analysis, we assume that the agent's type, θ_i, does not influence his expected utility outside of the firm—that is, we assume that agent productivity is entirely firm-specific. As the following proposition shows, however, the agent's type will influence his reservation fixed-wage.

Proposition 2. *First, low-productivity agents are more likely to leave the firm than high-productivity agents. Second, the marginal attrition associated with a commission reduction is greatest for high-productivity agents. Third, the distribution of incoming offers ultimately determines if the change in turnover rate is increasing in agent productivity.*

Proof. Proof of Proposition 2. The optimal effort e_i^* is increasing in type (see proof to Proposition 1), therefore revealed preference implies that the agents' expected utility $V(\alpha, R; \theta_i)$ is itself increasing in θ_i. Because the agents have a (strictly) positive utility for wages, the unique fixed-wage, $w(\theta_i)^*$, which makes an agent indifferent between the outside offer and his internal utility, $V(\alpha, R; \theta_i)$, must itself be increasing in θ_i. Consider two agents, with productivity levels $\theta_j > \theta_i > 0$. Since $w(\theta_j)^* > w(\theta_i)^*$, all offers $\overline{w} \geq w(\theta_j)^*$ are sufficient to lure both types of agents away from the firm. Offers $\underline{w} \in [w(\theta_i)^*, w(\theta_j)^*)$, on the other hand, will lure the agent with type θ_i but are insufficient to lure the agent with type θ_j. Accordingly, an agent with productivity θ_i will leave the firm while an agent with productivity $w(\theta_j)^*$ will remain with probability $F(w(\theta_j)^*) - F(w(\theta_i)^*) > 0$.

To prove the second statement, we must establish that $\frac{\partial^2 w(\theta_i)^*}{\partial R \partial \theta_i} > 0$, which suffices as the distribution of outside offers is independent of internal compensation contracts. By definition, $w(\theta_i)^*$ is the lowest external, fixed-wage offer that yields utility $V(\alpha, R; \theta_i)$ to agent i. Revealed preference guarantees that an agent's expected utility $V(\alpha, R; \theta_i)$ is strictly increasing in R. Accordingly, the minimum external wage $w(\theta_i)^*$ increases (decreases) for all types as the commission rate R increases (decreases). To see this formally, note that the envelope theorem yields $U'(e_i^*) = 0$, hence:

$$\frac{dV}{dR} = \frac{\partial U}{\partial R} + U'(e_i^*)\frac{\partial e_i^*}{\partial R} = \frac{\partial U}{\partial R} = \theta_i e_i^* > 0,$$

where the final inequality holds by the strict convexity of $c(\cdot)$ and the fact that both $c(0)$ and $c'(0)$ are equal to zero.

We must next prove that the marginal effect increases concomitantly with agent productivity:

$$\begin{aligned}\frac{d^2 V(\alpha, R, \theta_i)}{dR d\theta_i} &= \frac{\partial^2 U}{\partial R \partial \theta_i} + \frac{\partial U}{\partial R} U'(e_i^*)\frac{\partial e_i^*}{\partial \theta_i} \\ &= \frac{\partial^2 U}{\partial R \partial \theta_i} = 2e_i^* > 0.\end{aligned}$$

We have thus established that: (1) decreasing the commission rate R makes all agents more vulnerable to poaching, and (2) following a reduction of R, a highly productive agent, say an agent with productivity θ_j, decreases their external reservation rate, $w(\theta_j)^*$ by more than a less

[53] Without loss of generality, we assume that the fixed-wage offers require the agent to exert a fixed level of (un-modeled) effort with known dis-utility equal to 0.

productive agent reduces their own external reservation wage $w(\theta_i)^*$, where $\theta_i < \theta_j$. This does not, however, establish that high-productivity agents are more likely to leave the firm following a wage reduction, because separation nonetheless requires an external offer. To see this, consider a discrete change in R from $\overline{R}$ to $\underline{R}$ with $\overline{R} > \underline{R}$. Abusing notation, let $\underline{w}(\theta_i)^* = V(\alpha, \underline{R}, \theta_i)$ and $\overline{w}(\theta_i)^* = V(\alpha, \overline{R}, \theta_i)$. Accordingly, we can define $W(\theta_i)^* = [\underline{w}(\theta_i)^*, \overline{w}(\theta_i)^*)$ as the set of external wages which would suffice to lure an agent with productivity θ_i under the commission rate $\underline{R}$ but not under the commission rate $\overline{R}$. For $\theta_i < \theta_j$, we have shown that $\|W(\theta_i)^*\| < \|W(\theta_j)^*\|$, however $\int_{W(\theta_i)^*} f(w)dw$ may exceed $\int_{W(\theta_j)^*} f(w)dw$. In other words, the distribution of incoming external offers ultimately determines whether high- or low-productivity workers are more likely to separate from the firm following a reduction in the commission rate, R. □

The intuition behind the first statement in Proposition 2 is relatively straight-forward: because all agents face the same distribution of outside offers, those with the lowest reservation utility are more likely to accept a relatively low outside offer, and hence are the most likely to leave. The second finding is slightly more nuanced; while all agents are more likely to accept an outside offer once their (internal) commission rate, R, decreases, a commission reduction will decrease a high-productivity agent's reservation wage $w(\theta_j)^*$ by more than the same commission rate change affects a low-productivity agent's reservation wage $w(\theta_i)^*$—the difference in reservation wage adjustment is determined by the agents' (common) effort cost function, $c(e)$. Despite the larger reservation wage adjustment, the theory is unable to predict how a change in the commission rate, R, will effect relative attrition rates because we have not imposed restrictions on the distribution, $f(\cdot)$ of external offers w. If, however, internal productivity did influence external offers; e.g. if agents can project their productivity to external employers, then highly-productive agents will be that much more likely to separate from the firm. Even had we modeled such a mechanism, without very strict assumptions now on the conditional distribution of external opportunities, whether or not highly-productive agents are more likely to separate from the firm following an adverse commission change, would remain an empirical question. The answer to this question influences how compensation changes map into firm profits.

Proposition 3. *The sensitivity of changes in profits with respect to sales commissions depends on the turnover propensity of high-productivity agents relative to low-productivity agents. The turnover of highly productive agents mitigates any cost savings from reducing R.*

Proof. Proof of Proposition 3. We consider a representative sales opportunity allocated to a random agent. Let $g(\theta|R)$ denote the density of agent types at the firm under the commission structure R. The expected profits from the sales opportunity are

$$(1-R)\int \theta e^*(\theta, R)\, dG(\theta|R).$$

Differentiation with respect to R yields

$$\frac{\partial \pi}{\partial R} = -\int \theta e^*(\theta, R)\, dG(\theta|R) + (1-R)\int \left\{\theta \frac{\partial e^*}{\partial R} g(\theta|R) + \theta e^* \frac{\partial g(\theta|R)}{\partial R}\right\} d\theta.$$

The first term, $-\int \theta e^*(\theta, R)\, dG(\theta|R)$, is negative, as raising commissions while holding sales fixed provides the agent with a transfer. When $\frac{\partial g(\theta|R)}{\partial R} = 0$, such that there is no sorting, the sign of the second term is positive, meaning the agent's positive effort response may offset the firm's decreased profits from the transfer made to the agent. When $\frac{\partial g(\theta|R)}{\partial R} > 0$, the average quality of the workforce increases with R, further offsetting the firm's decreased profits stemming from marginal transfers to the agent. $\square$

A reduction in commissions has two different effects: profits increase because of cost savings, while effort reductions offset some of these savings. When the change in the composition of the workforce is greatest for highly productive workers, that is $\frac{\partial g(\theta|R)}{\partial R}$ is increasing in θ, the loss of highly productive workers further offsets the cost savings from the commission changes. The magnitude of the composition and effort changes is the empirical question that we examine.

7 Re-weighting Estimators

This section provides details about the implementation of the re-weighting estimators that attempt to match individuals in control divisions with individuals in Division 1. The purpose is to match individuals based on their sales trajectories. The first step is to estimate the probability of being in Division 1. We use the data from the pre-treatment period for this purpose but hold out the data one month prior to the commission schedule changes. The second step is to use the propensity score from this estimation procedure to form weights which will be used in later regressions. The third step is to assess how well the re-weighting estimates fit, using a "hold out" sample of data one month prior to the commission schedule changes.

In the first step, we estimate logit models where the dependent variable is being in Division 1. Each worker present in the pre-treatment period for Division 1 and the control divisions enters the sample once. The first month of available data includes the X variables and demographic characteristics in levels. The regressors in X are an indicator for male, the agent's age, and the agent's monthly averages of log commissions, log commission per call, log revenue, log total calls, tenure, and adherence. For each of the regressors on productivity, we also include one and two month differences over future months to capture trends in these measures. We then estimate the logit model and form $\hat{P}$, the predicted probability of being in Division 1.

The weights in the second step are formed as $W_i = Treated_i + (1 - Treated_i)\frac{\hat{P}}{1-\hat{P}}$ where $\hat{P}$ is the treatment probability estimated from the logistic regression on pre-treatment data and $Treated_i$ indicates the worker is in Division 1. Figures OA.1a and OA.1b assess fit, making it clear that per-call fit works reasonably well. Fit for overall revenue is not as good, suggesting that the territory shock yielded an up-tick in sales success among Division 1 agents. As a result, we prefer specifications at the per-call level to remove potential demand confounders when interpreting changes in effort supply. These per-call measures of productivity allow us to measure output while controlling for demand.[54]

[54]Given that the divergence between the re-weighted control group trend and the trend for Division 1 occurs before the commission schedule changes, we suspect demand changes are responsible for divergence in the levels measures.

Appendix Figures and Tables

Figure A.1: Common Trends by Worker Type within Division

(a) **Log Commissions by Median Ability**

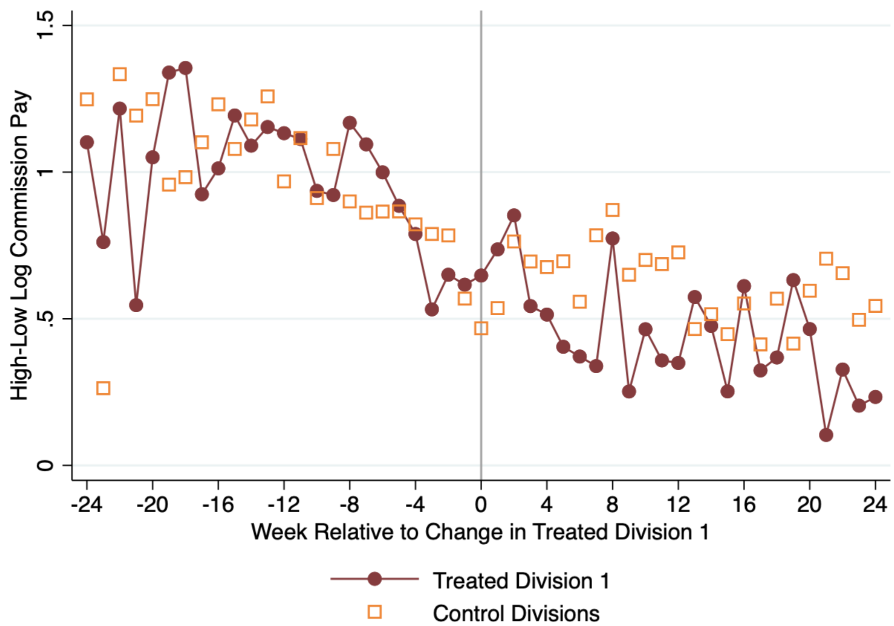

(b) **Log Revenue by Median Ability**

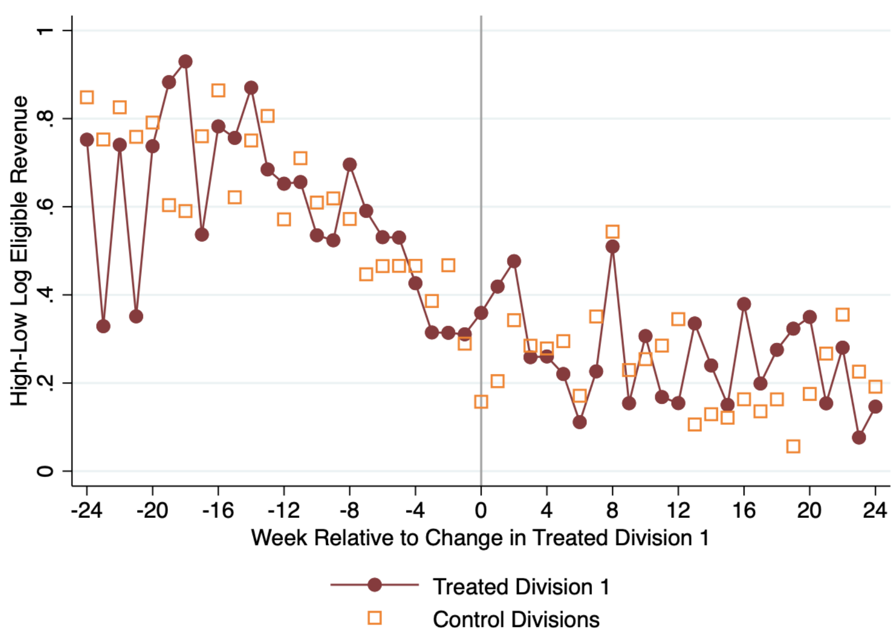

Notes: These figures plot the evolution of within-division differences in performance by worker pre-treatment productivity. Figure (a) considers trends in log commissions, whereas Figure (b) considers trends in log revenue. Week 0 on the x-axis denotes the two weeks immediately before the commission schedule changes occurred. The y-axis in each figure captures the differences in output between high and low performers.

Figure A.2: Commission Trends in Division 2 and the Control Divisions

(a) **Trends in Commission Levels**

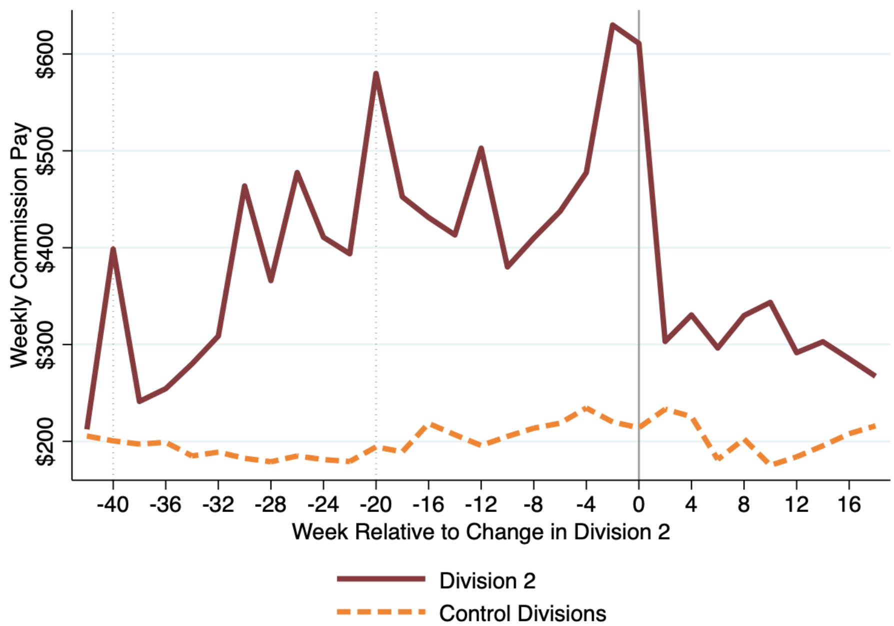

(b) **Differences in Commission Levels and Total Call Volume**

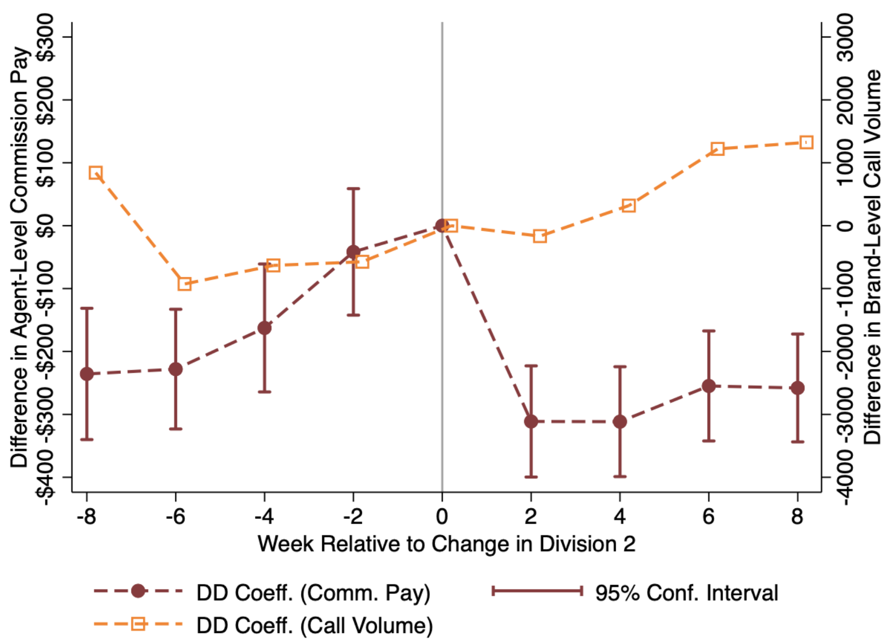

Notes: Figure (a) plots the average weekly commission pay levels for agents in Division 2 and the control divisions. The solid vertical line corresponds to the two weeks immediately before the week of the commission schedule changes in Division 2. Figure (b) plots the difference-in-differences coefficients that capture differential trends in commission pay levels and total call volume between Division 2 and the control divisions.

Figure A.3: Survival Rates in Division 2 and the Control Divisions

(a) **Survival Rates Relative to Month -5**

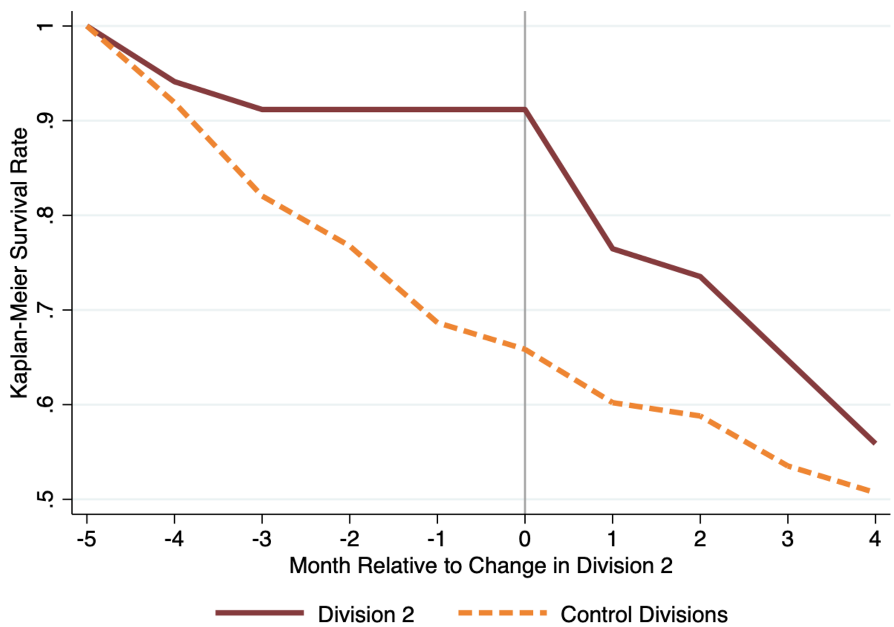

(b) **Survival Rates Relative to Month 0**

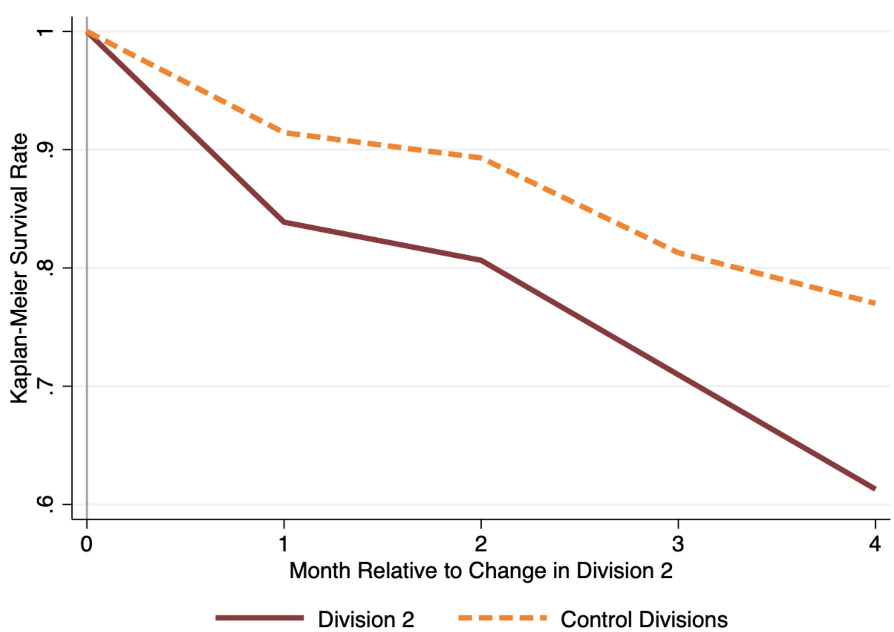

Notes: These figures plot Kaplan-Meier survival rates over time. The survival rate estimator considers a starting point and then, from that time, displays the fraction of agents that remain at the firm. The graphical properties of the cumulative survival rate allow an assessment of when retention diverges over time and what fraction of the total beginning workforce is affected. Because turnover can be lumpy, with multiple exits in some weeks and no exits in others, we aggregate survival rates to the monthly level.

Table A.1: Illustration of Commission Changes

	Pre-Change	Post-Change	Difference (2)−(1)	Commissions (3) × 10%
	(1)	(2)	(3)	(4)
One of Three Products				
Transfer Price per Sale	\$15	\$10	-\$5	
Avg. Sales per Agent-Week	39.9	38.5	-1.40	
Avg. Revenue per Agent-Week	\$598.50	\$385.00	-\$213.50	-\$21.35
Bundle of Two Products				
Transfer Price per Sale	\$50	\$25	-\$25	
Avg. Sales per Agent-Week	7.56	4.97	-2.59	
Avg. Revenue per Agent-Week	\$378.00	\$124.25	-\$253.75	-\$25.38
Bundle of Three Products				
Transfer Price per Sale	\$100	\$125	\$25	
Avg. Sales per Agent-Week	5.33	5.99	0.66	
Avg. Revenue per Agent-Week	\$533.00	\$748.75	\$215.75	\$21.58
Total per Agent-Week	\$1,509.50	\$1,258.00	-\$251.50	-\$25.15

Notes: The purpose of this table is to better highlight some of the details of the commission schedule changes. We display agent-week level sales averages and transfer prices for different bundles of three separate products. While we have a partial record of the products sold, we do not have any way of knowing what products customers initial sought out when they called. As a result, we are unable to measure product-level conversion rates. Columns (1) and (2) show revenue transfer prices, average sales per agent-week, and average revenue per agent-week for different product bundles in the pre- and post-treatment periods, respectively. Column (3) displays the differences in transfer prices, average sales per agent-week, and average revenue per agent-week between between these two periods. Column (4) multiplies this difference by a hypothetical commission rate of 10% (which is at the top end of the commission rate distribution). As mentioned in Section 2.3, agents' commission rates were not mechanically changed, so an agent with a commission rate of 10% in the pre-treatment period likely maintained this commission rate in the post-treatment period.

Table A.2: Estimates of Effort Responses Using the Full Pre-Treatment Sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Adherence to Schedule							
Treated x Post	-0.001	0.003	0.014	0.017	0.009	0.014	
	(0.006)	(0.005)	(0.008)	(0.012)	(0.009)	(0.008)	
Treated x Post x Prod						-0.003	-0.003
						(0.004)	(0.004)
Observations	32,540	32,540	32,540	14,269	14,218	32,540	32,540
Panel B: Conversion Rate							
Treated x Post	0.005	0.001	0.006	0.001	0.005	0.011	
	(0.006)	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	
Treated x Post x Prod						-0.017***	-0.018***
						(0.004)	(0.003)
Observations	33,044	33,044	33,044	13,469	13,981	33,044	33,044
Panel C: Log RPC at Old Prices							
Treated x Post	0.076	0.053	-0.006	-0.006	-0.003	0.005	
	(0.046)	(0.033)	(0.033)	(0.034)	(0.041)	(0.034)	
Treated x Post x Prod						-0.054*	-0.044
						(0.027)	(0.023)
Observations	35,366	35,366	35,366	15,077	15,071	35,366	35,366
Panel D: Log RPC at New Prices							
Treated x Post	0.069	0.061*	0.063	0.037	0.062	0.077*	
	(0.043)	(0.030)	(0.032)	(0.032)	(0.041)	(0.033)	
Treated x Post x Prod						-0.061*	-0.049*
						(0.026)	(0.022)
Observations	35,366	35,366	35,366	15,077	15,071	35,366	35,366
Week Fixed Effects	✓	✓	✓	✓	✓	✓	
Agent Fixed Effects		✓	✓	✓	✓	✓	✓
Division Trend Controls			✓	✓	✓	✓	
Week x Division Fixed Effects							✓
Re-Weighted				✓			
Balanced Sample					✓		

Notes: This table is an analog of Table 4. The sample includes all current employees in Division 1 and the control divisions with non-missing data. The models in Columns 1–6 include fixed effects for week, division, and office location. All models include cubic splines for tenure and a cubic polynomial in age. The OLS regression in Column 1 includes dummies for ethnicity, gender, and marital status. The specifications in Columns 2–7 include individual fixed effects. Columns 3–6 include division-specific trend controls. The specification in Column 4 uses a re-weighting estimator based on the propensity score for being in Division 1 (see Appendix 7). The balanced panel in Column 5 restricts to workers who are present prior to July, 2016 and after April, 2017. Columns 6 and 7 consider heterogeneous responses based on worker productivity, and Column 7 omits week fixed effects and division-specific trend controls and instead includes week by division fixed effects. Differing numbers of observations across panels reflect differences in data availability. The sample include all pre- and post-treatment period data in the immediate sample. Reported standard errors are clustered by manager.

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Table A.3: Summary Statistics for Division 2 By Productivity Level

	Adjusted Worker Fixed Effects		
	Bottom Third	Middle Third	Top Third
	(1)	(2)	(3)
Commission	306.13	496.85	717.00
	(206.50)	(316.55)	(345.83)
RPC	73.85	101.41	129.99
	(36.80)	(44.14)	(46.00)
Adherence	0.75	0.78	0.81
	(0.21)	(0.15)	(0.10)
Conversion	0.23	0.29	0.34
	(0.11)	(0.11)	(0.12)
Phone Hours	16.96	17.70	17.45
	(6.68)	(6.61)	(5.39)
Total Calls	52.46	45.53	47.69
	(22.64)	(17.86)	(14.23)
Tenure (days)	324.13	679.67	1386.26
	(93.80)	(356.30)	(414.86)
Age	26.16	31.14	32.51
	(4.00)	(8.52)	(9.94)
Single	0.77	0.51	0.30
	(0.42)	(0.50)	(0.46)
White	0.92	0.25	0.74
	(0.28)	(0.43)	(0.44)
Male	0.75	0.75	0.61
	(0.44)	(0.43)	(0.49)
Survey Response to Firm Fairness	0.22	0.45	0.10
	(0.42)	(0.50)	(0.30)
Survey Response to Referral Likelihood	0.63	0.63	0.54
	(0.49)	(0.49)	(0.50)
Survey Response to Promotion Likelihood	0.66	0.57	0.46
	(0.48)	(0.50)	(0.50)
Agent-Weeks	95	97	90
Agents	13	13	12

Notes: This table presents cross-sectional summary statistics for Division 2 using data eight weeks prior to the Division 2 commission schedule changes. Each column represents an approximate tercile of the distribution of adjusted worker fixed effects in the pre-treatment period. Adjusted worker fixed effects are calculated from a regression of log commissions on worker dummy variables, division-by-week dummy variables, and a cubic spline in tenure. We then correct for sampling variation using the method in Lazear et al. (2015).

Table A.4: Linear Probability Model Estimates of Turnover Responses in Division 2

	Last Week in Firm		
	(1)	(2)	(3)
Treated x Post x Prod	-0.011	0.002	-0.008
	(0.007)	(0.007)	(0.007)
Treated x Post	0.013**	0.023*	
	(0.004)	(0.010)	
Time Fixed Effects	✓	✓	
Division x Week-of-Year Fixed Effects		✓	
Time x Division Fixed Effects			✓
Observations	45,328	45,328	45,328
Mean Turnover Prob in Treated Division		0.008	
p-value on Treated x Post x Prod	0.232	0.681	0.442
p-value on Treated x Post	0.125	0.383	

Notes: The dependent variable is an indicator that equals one if it is the worker's last week at the firm. The sample includes all current employees in Division 2 and the control divisions with non-missing data. Estimates come from a linear probability model that captures changes in the turnover probability for the existing workforce. Each model includes a 5th order polynomial for workers' tenure to account for a potentially arbitrary baseline relationship between tenure and turnover. *Prod* refers an agent's sales z-score, which is the standardized measure of an agent's pre-treatment productivity estimated as their adjusted worker fixed effect according to the procedure in Lazear et al. (2015). For additional details, see Section 2.5. The specification in Column 2 includes division by week-of-year fixed effects to account for seasonality. The specification in Column 3 includes week by division fixed effects. Two forms of inference are presented, one using standard errors clustered by manager (in parentheses) and the second using p-values with division-level clusters (see the final two lines) computed using the wild cluster bootstrap randomization inference procedure in MacKinnon and Webb (2018). We use the t-statistic version of the procedure that imposes the null hypothesis.

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

8 Online Appendix Materials

8.1 Spillovers to the Control Group

To test another identifying assumption, the lack of spillovers to control divisions (Obloj and Zenger, 2017), we conduct structural break tests for the control group. Figure OA.4 in the plots the parameter estimates from various specifications of these break tests. Structural break tests come from regressions using the control sample. The figure reports the post-treatment indicator parameter estimates and confidence intervals. We consider several different dependent variables, and each regression includes a post-treatment indicator for Division 1, the matrix of agent characteristics X_{it}, division fixed effects, and trends for each division. These results suggest that there are minimal spillovers to the control group.

Figure OA.3 plots the time series process for the control groups around the event date for Division 1. Within a month of the event date, there is minimal movement in the control group averages. Conversion rates and RPC do show some mild deterioration after the first month, which is likely due to seasonality based on the time of the year.

8.2 Substitution to Different Products

Whether agents could reduce the impact of the commission schedule changes because of substitution to other products is an empirical question. The approach is to estimate whether sales revenue becomes more heavily weighted to items with more favorable relative prices under the new commission schedule. Although there were some relative price changes that may have given rise to agent substitution, we find that agents could not offset the adverse effects of the commission schedule changes by changing their mix of products sold. That is, the overall change in commissions-per-call that we estimate closely follow the predicted reductions given the pre-treatment mix of products sold.

Figure OA.1: Re-weighted Commissions and Commissions-per-Call for Division 1

(a) **Log Commissions per Call (log scale)**

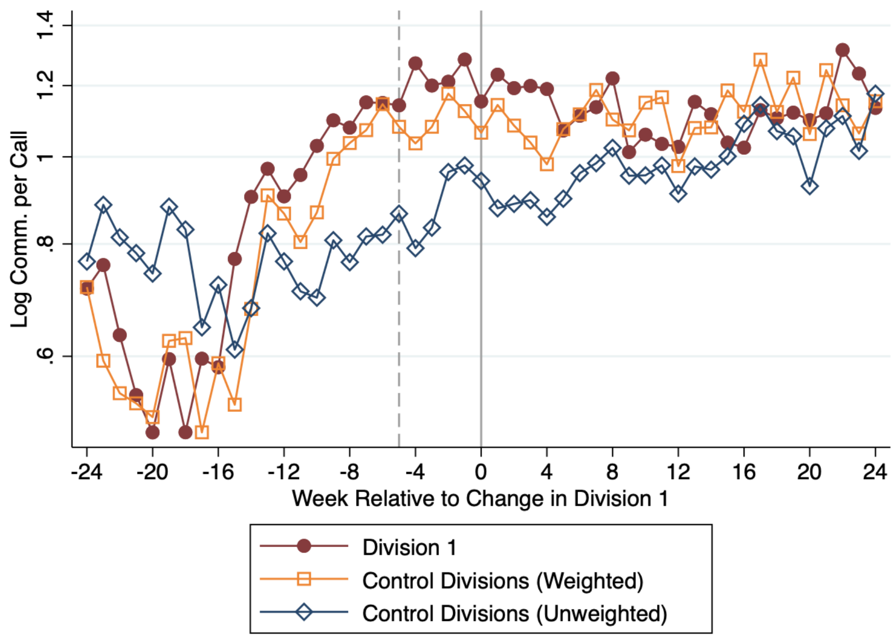

(b) **Commissions per Week (log scale)**

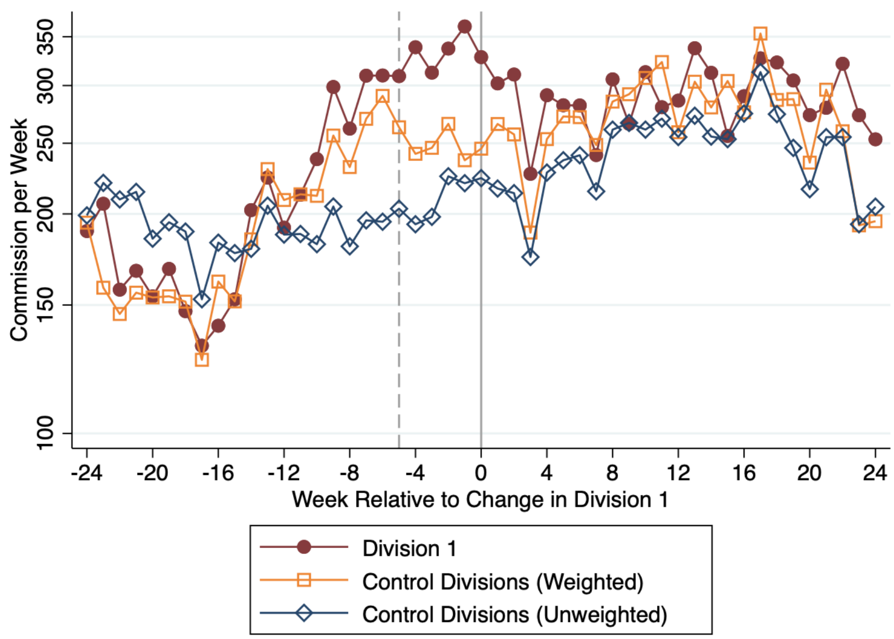

Notes: These figures display unweighted and propensity score weighted comparisons of agents in the control divisions and Division 1. The dashed line represents the end of the period used for estimating the propensity score weights. Output measures are displayed on a log scale.

Figure OA.2: Structural Break Tests in the Control Divisions

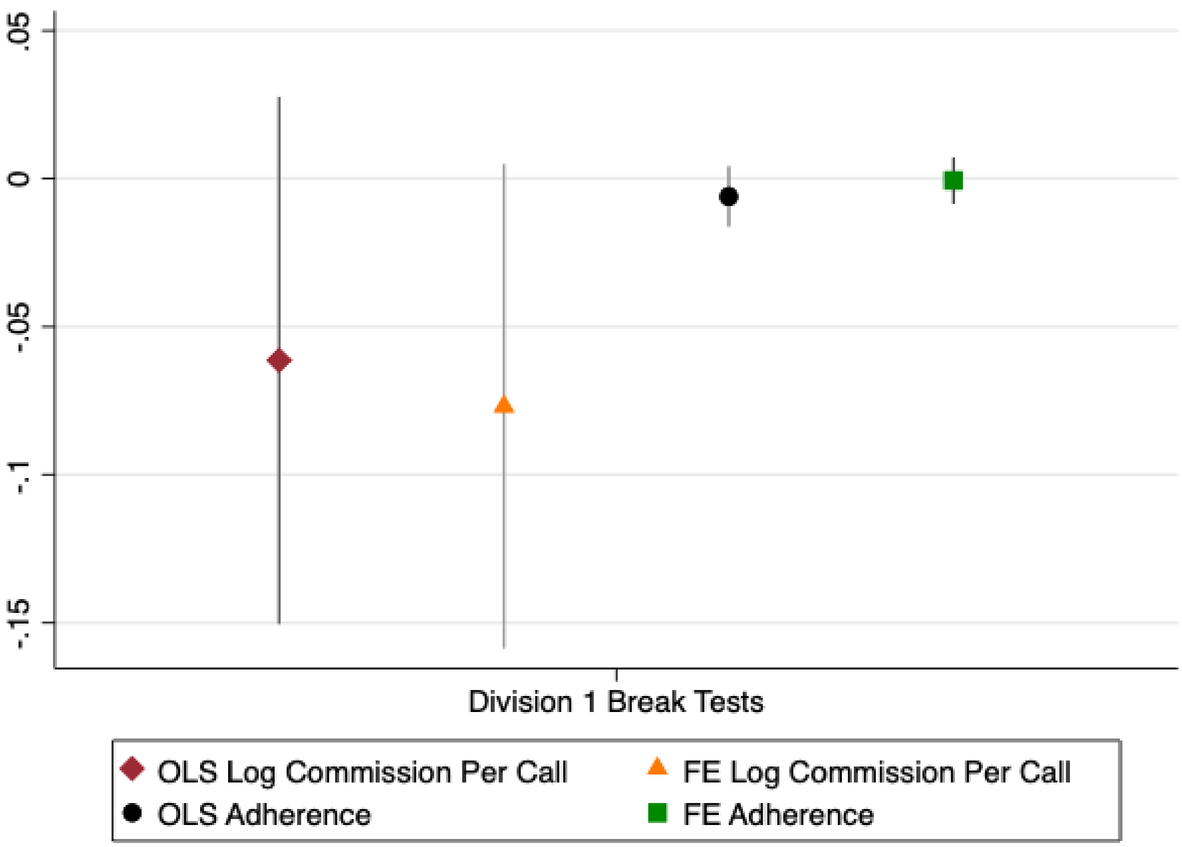

Notes: These structural break tests come from regressions using the control sample. The figure reports the post-treatment indicator parameter estimates and the corresponding confidence intervals. The dependent variable is in the legend, and each regression includes a post-treatment indicator for Division 1, the matrix of agent characteristics X_{it}, division fixed effects, and trends for each division. Specifications with "FE" add individual fixed effects.

Figure OA.3: Trends in Observable Outcomes for Control Divisions

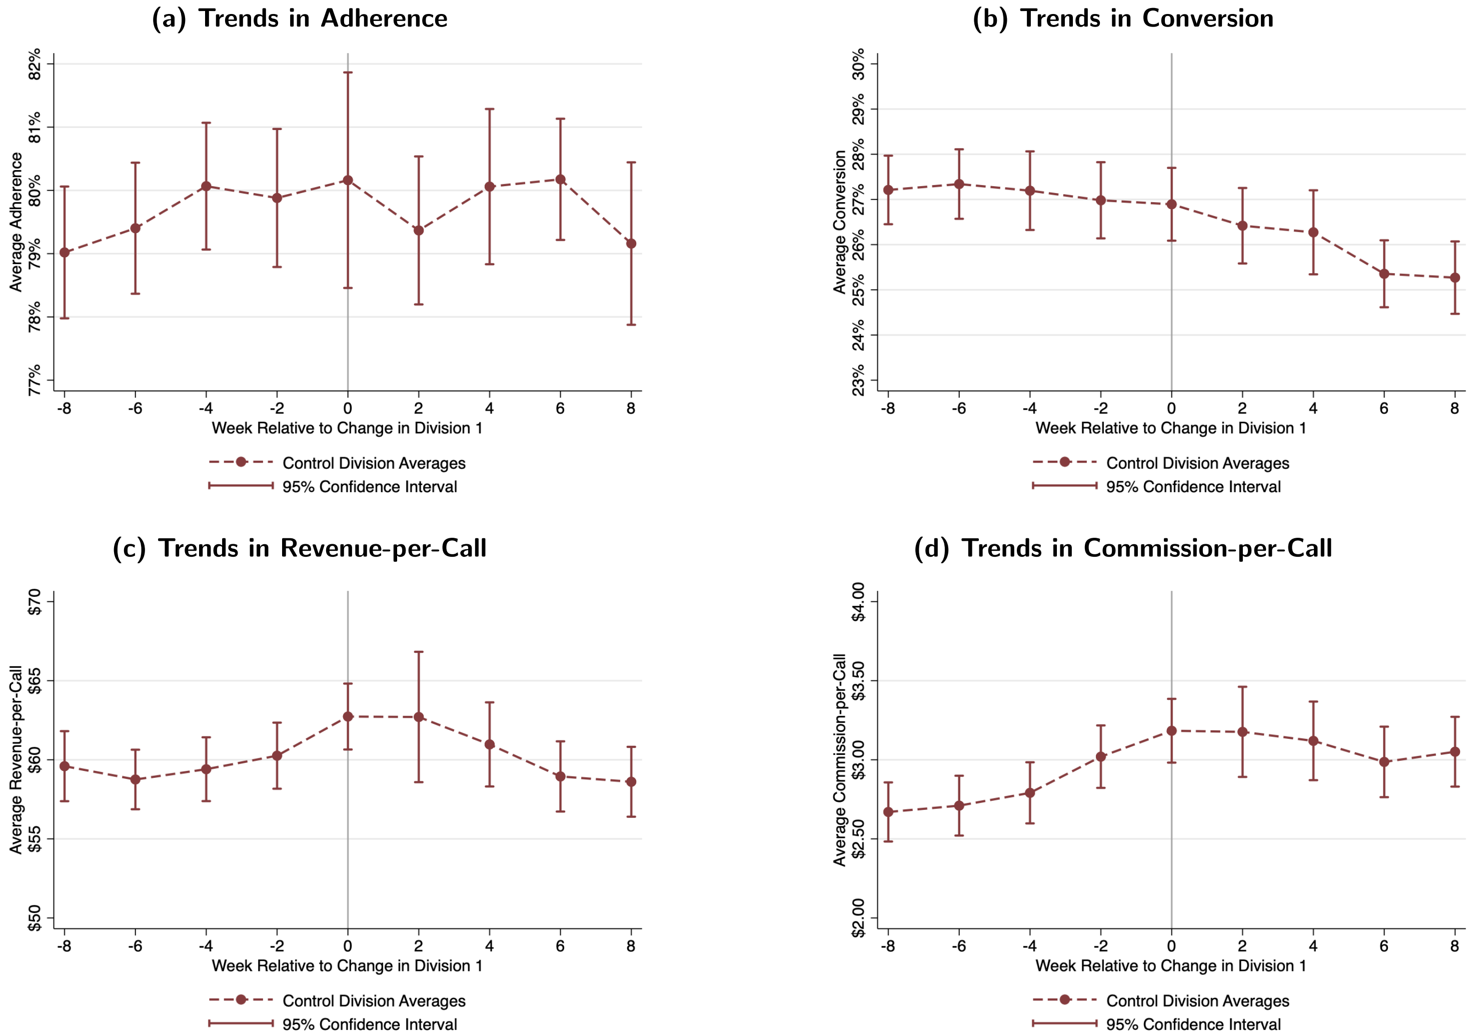

Notes: These figures show raw averages in different outcome variables for the control divisions. Adherence and conversion are the two proxies for an agent's supply of effort. Revenue-per-call and commission-per-call are two additional measure of output. To improve the readability of these figures, we aggregate data into bi-weekly clusters. Confidence intervals are based on the standard errors of the means.

Figure OA.4: Placebo Tests for Effort Estimations

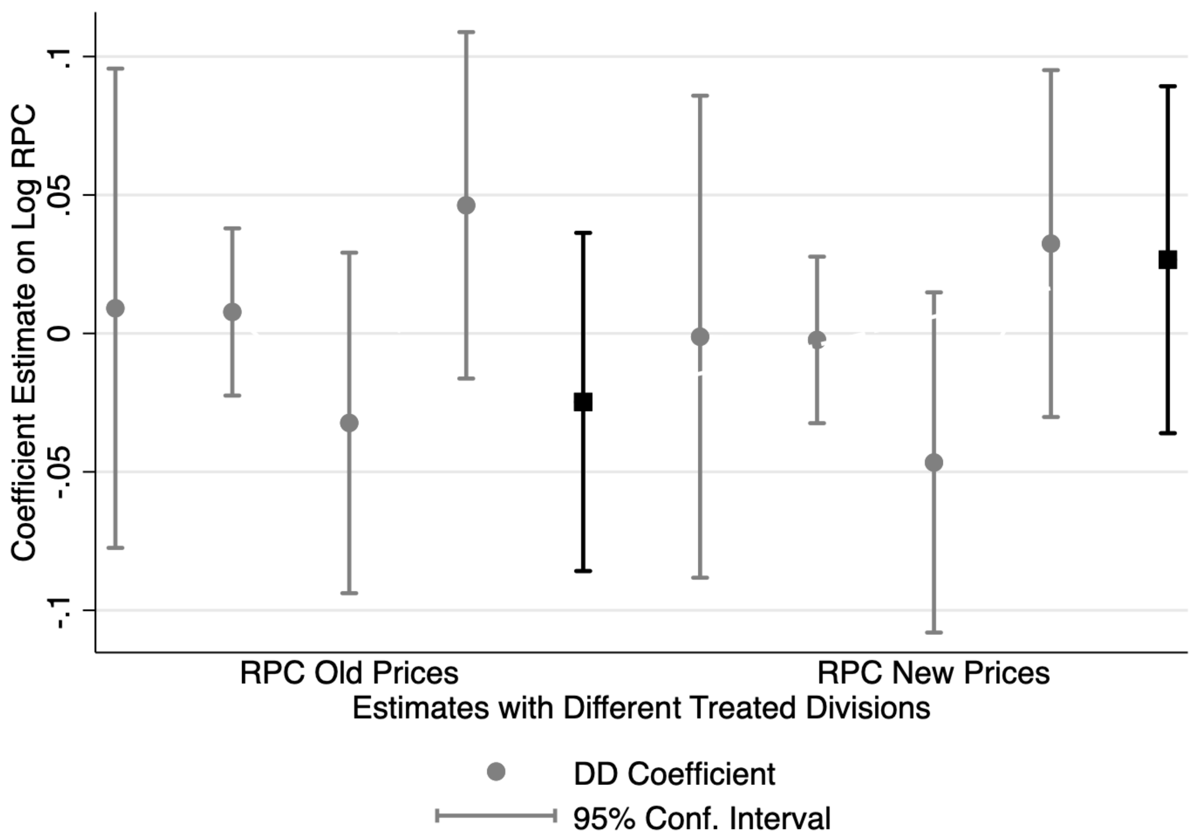

Notes: This figure plots placebo simulations for the effort response estimation using different divisions as the "treated" division, while the other divisions (including the actual treated division) make up the control group. These are marked by the gray dots. The actual results from Column 1 of Table 4 are depicted by the black squares. The left five estimations use log RPC based on the old prices as the dependent variable, whereas the right five estimations use log RPC based on the new prices as the dependent variable.

Table OA.1: Linear Probability Model Estimates of Turnover Responses (Low-Ordered Polynomials)

	Last Week in Firm				
	(1)	(2)	(3)	(4)	(5)
Treated x Post x Prod	0.021**	0.015**	0.016*	0.012**	0.012*
	(0.007)	(0.005)	(0.006)	(0.005)	(0.005)
Treated x Post	-0.006	-0.006		-0.002	
	(0.004)	(0.007)		(0.010)	
Treated x Placebo x Prod	-0.006		-0.002		
	(0.004)		(0.004)		
Treated x Placebo	0.000				
	(0.004)				
Week Fixed Effects	✓	✓		✓	
Division x Week-of-Year Fixed Effects		✓			
Week x Division Fixed Effects			✓		✓
Post-Territory Shock Period				✓	✓
Observations	51,497	51,497	51,497	19,689	19,689
Mean Turnover Probability in Division 1			0.037		

Notes: The dependent variable is an indicator that equals one if it is the worker's last week at the firm. The sample includes all current employees in Division 1 and the control divisions with non-missing data. Estimates come from a linear probability model that captures changes in the turnover probability for the existing workforce. These models include only Age, Age^2, $Tenure$, and $Tenure^2$, removing the higher-ordered polynomial terms on *Age* and *Tenure*. *Prod* refers an agent's sales z-score, which is the standardized measure of an agent's pre-treatment productivity estimated as their adjusted worker fixed effect according to the procedure in Lazear et al. (2015). For additional details, see Section 2.5. The specification in Column 2 includes division by week-of-year fixed effects to account for seasonality. The specification in Column 3 includes week by division fixed effects. Columns 4 and 5 use a shortened pre-treatment period that only includes the weeks of data after the territory shock period. *Placebo* is an indicator for the date 52 weeks prior to the treatment date.

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Table OA.2: Estimates of Effort Responses (Low-Ordered Polynomials)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Adherence to Schedule							
Treated x Post	0.006	0.012	0.046**	0.032	0.005	0.047**	
	(0.009)	(0.010)	(0.015)	(0.021)	(0.010)	(0.015)	
Treated x Post x Prod						-0.008	-0.007
						(0.005)	(0.005)
Observations	8,119	8,119	8,119	7,064	3,706	8,119	8,119
Panel B: Conversion Rate							
Treated x Post	0.005	0.001	0.009	0.006	0.008	0.015*	
	(0.008)	(0.004)	(0.007)	(0.008)	(0.007)	(0.006)	
Treated x Post x Prod						-0.019***	-0.019***
						(0.005)	(0.005)
Observations	8,283	8,283	8,283	6,903	3,743	8,283	8,283
Panel C: Log RPC at Old Prices							
Treated x Post	-0.029	-0.043	0.001	-0.015	0.007	0.024	
	(0.032)	(0.022)	(0.033)	(0.036)	(0.048)	(0.034)	
Treated x Post x Prod						-0.043	-0.044
						(0.022)	(0.022)
Observations	9,229	9,229	9,229	7,840	4,126	9,229	9,229
Panel D: Log RPC at New Prices							
Treated x Post	0.023	0.004	0.023	0.003	0.018	0.049	
	(0.033)	(0.021)	(0.036)	(0.039)	(0.054)	(0.036)	
Treated x Post x Sales Z Score						-0.050*	-0.050*
						(0.022)	(0.023)
Observations	9,229	9,229	9,229	7,840	4,126	9,229	9,229
Week Fixed Effects	✓	✓	✓	✓	✓	✓	
Agent Fixed Effects		✓	✓	✓	✓	✓	✓
Division Trend Controls			✓	✓	✓	✓	
Week x Division Fixed Effects							✓
Re-Weighted				✓			
Balanced Sample					✓		

Notes: The sample includes all current employees in Division 1 and the control divisions with non-missing data. The models in Columns 1–6 include fixed effects for week, division, and office location. These models include only Age, Age^2, $Tenure$, and $Tenure^2$, removing the higher-ordered polynomial terms on Age and $Tenure$. The OLS regression in Column 1 includes dummies for ethnicity, gender, and marital status. The specifications in Columns 2–7 include individual fixed effects. Columns 3–6 include division-specific trend controls. The specification in Column 4 uses a re-weighting estimator based on the propensity score for being in Division 1 (see Appendix 7). The balanced panel in Column 5 restricts to workers who are present prior to July, 2016 and after April, 2017. Columns 6 and 7 consider heterogeneous responses based on worker productivity, and Column 7 omits week fixed effects and division-specific trend controls and instead includes week by division fixed effects. Differing numbers of observations across panels reflect differences in data availability. The sample used restricts to eight weeks of pre-treatment data and eight weeks of post-treatment data. The results are similar when all available pre- and post-treatment data is used (See Table A.2). Reported standard errors are clustered by manager. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Table OA.3: Estimates of Effort Responses (Log RPH and RPC Levels)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Log RPH at Old Prices							
Treated x Post	0.051	0.041	0.035	0.007	0.053	0.036	
	(0.046)	(0.035)	(0.046)	(0.044)	(0.059)	(0.046)	
Treated x Post x Prod						-0.016	-0.017
						(0.032)	(0.032)
Observations	9,145	9,145	9,145	7,761	4,075	9,145	9,145
Panel B: Log RPH at New Prices							
Treated x Post	0.102*	0.086*	0.057	0.026	0.065	0.061	
	(0.047)	(0.035)	(0.045)	(0.044)	(0.062)	(0.045)	
Treated x Post x Prod						-0.024	-0.024
						(0.032)	(0.032)
Observations	9,145	9,145	9,145	7,761	4,075	9,145	9,145
Panel C: Level RPC at Old Prices							
Treated x Post	-5.012	-3.381	-1.155	-2.117	-1.108	0.885	
	(2.614)	(2.015)	(2.593)	(3.091)	(3.596)	(2.491)	
Treated x Post x Prod						-4.362**	-4.377**
						(1.554)	(1.599)
Observations	9,229	9,229	9,229	7,840	4,126	9,229	9,229
Panel D: Level RPC at New Prices							
Treated x Post	-1.426	0.057	-0.068	-1.407	-0.781	1.813	
	(2.179)	(1.775)	(2.258)	(2.907)	(3.544)	(2.211)	
Treated x Post x Prod						-4.199**	-4.181**
						(1.457)	(1.493)
Observations	9,229	9,229	9,229	7,840	4,126	9,229	9,229
Week Fixed Effects	✓	✓	✓	✓	✓	✓	
Agent Fixed Effects		✓	✓	✓	✓	✓	✓
Division Trend Controls			✓	✓	✓	✓	
Week x Division Fixed Effects							✓
Re-Weighted				✓			
Balanced Sample					✓		

Notes: The sample includes all current employees in Division 1 and the control divisions with non-missing data. The models in Columns 1–6 include fixed effects for week, division, and office location. All models include cubic splines for tenure and a cubic polynomial in age. The OLS regression in Column 1 includes dummies for ethnicity, gender, and marital status. The specifications in Columns 2–7 include individual fixed effects. Columns 3–6 include division-specific trend controls. The specification in Column 4 uses a re-weighting estimator based on the propensity score for being in Division 1 (see Appendix 7). The balanced panel in Column 5 restricts to workers who are present prior to July, 2016 and after April, 2017. Columns 6 and 7 consider heterogeneous responses based on worker productivity, and Column 7 omits week fixed effects and division-specific trend controls and instead includes week by division fixed effects. The sample used restricts to eight weeks of pre-treatment data and eight weeks of post-treatment data. Reported standard errors are clustered by manager.

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$