



The Comprehensive Effects of Sales Force Management: A Dynamic Structural Analysis of Selection, Compensation, and Training

Citation

Chung, Doug J., Byungyeon Kim, and Byoung G. Park. "The Comprehensive Effects of Sales Force Management: A Dynamic Structural Analysis of Selection, Compensation, and Training." *Management Science* 67, no. 11 (November 2021): 7046–7074.

Published version

<https://doi.org/10.1287/mnsc.2020.3853>

Link

<https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37373868>

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles (OAP), as set forth at

<https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFIOWI2M2EwYTg>

Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#)

The Comprehensive Effects of Sales Force Management: A Dynamic Structural Analysis of Selection, Compensation, and Training

Doug J. Chung, Harvard University
Byungyeon Kim, Harvard University
Byoung G. Park, University at Albany*

April 2020

Forthcoming at *Management Science*

Abstract

This study provides a comprehensive model of an agent's behavior in response to multiple sales management instruments, including compensation, recruiting/termination, and training. The model takes into account many of the key elements that constitute a realistic sales force setting—allocation of effort; forward-looking behavior; present bias; training effectiveness; and employee selection and attrition. By understanding how these elements jointly influence agents' behavior, the study provides guidance on the optimal design of sales management policies. A field validation, by comparing counterfactual and actual outcomes under a new policy, attests to the accuracy of the model. The results demonstrate a trade-off between adjusting fixed and variable pay; how sales training serves as an alternative to compensation; a potential drawback of hiring high-performing, experienced salespeople; and how utilizing a leave package leads to sales force restructuring. In addition, the study offers a key methodological contribution by providing formal identification conditions for hyperbolic time preference. The key to identification is that, under a multi-period nonlinear incentive system, an agent's proximity to a goal affects only future payoffs in non-pecuniary benefit periods, providing exclusion restrictions on the current payoff.

Keywords: sales force compensation; training; selection; recruiting; termination; hyperbolic discounting; present bias; dynamic structural models; exclusion restriction; identification.

* Doug J. Chung is the MBA Class of 1962 Associate Professor of Business Administration and Byungyeon Kim is a doctoral student at Harvard University. Byoung G. Park is an Assistant Professor of Economics at the University at Albany. The authors thank the department editor, associate editor, and three anonymous reviewers for their constructive feedback during the review process. The authors also thank Murat Gumrukcu, Sunil Gupta, Rajiv Lal, Donald Ngwe, Hakan Sahin, Alvin Silk, Karin Soydan, K. Sudhir, Hasan Ulusoy, Yi Xin, Gamze Yucaoglu, and the seminar participants at Harvard University, McKinsey & Company as well as the conference participants at the 2018 INFORMS Marketing Science Conference, 2018 China India Insights Conference, 2018 Frontiers in Empirical Marketing Conference, and 2019 Marketing Dynamics Conference for their comments and suggestions. The authors also thank the anonymous firm for cooperating in this study.

1. Introduction

Effective management of the sales force is vital to the success of sales-driven organizations. Approximately 15 million salespeople in the United States, representing about 10% of the entire U.S. labor force, serve as links between the customer and the firm (U.S. Department of Labor, 2018). Investments in these salespeople are approximated to be 10% of sales revenues and can reach up to 40% in certain industries (Heide, 1999). The U.S. economy spends more than \$800 billion on sales forces each year, nearly four times the estimated \$208 billion spending on media (\$98 billion) and digital (\$110 billion) advertising (Zoltners et al., 2013; MAGNA, 2018). As these significant figures suggest, personal selling represents one of the most important elements of the firm’s marketing mixes, highlighting the importance of managing and motivating salespeople to achieve the organization’s objectives.

Since the earliest days of personal selling, organizations have utilized three main sales (force) management instruments to better control and motivate the sales force: compensation, recruiting/retention (of high-ability employees), and training. **Figure 1** illustrates the relation between these instruments and the organization’s sales performance. Performance is an outcome of salespeople’s behavior, and the sales management instruments are used for training and motivating proper behavior, as well as for selecting the right type of people. Not only do the three key instruments differ in cost and effectiveness across different types of people, but they also are interconnected in their effectiveness at changing behavior and, thus, attaining the desired performance outcome. This study aims to jointly examine the effectiveness of multiple sales management instruments in the selection and performance of heterogeneous salespeople.

A key, if not the most important, instrument in sales management is compensation. Organizations frequently use compensation to motivate and control the behavior of salespeople. A sales force compensation system typically consists of fixed- and variable-pay components, with each component playing a distinct role in managing sales force behavior. Fixed pay (base salary) compensates for risk and, thus, provides stability and security of income (Arrow, 1971; Basu et al., 1985; Harris & Raviv, 1978; Hölmstrom, 1979; Lal & Srinivasan, 1993). Variable pay, on the other hand, provides a direct link between the sales outcome and financial rewards, thereby inducing motivation to achieve superior performance. Examples of variable pay include commissions, given as a proportion of sales, and lump-sum bonuses, contingent on meeting a preset quota. According to Joseph and Kalwani (1998), 95% of U.S. firms utilize some form of variable pay to incentivize their salespeople, with the most frequently used forms being commissions and quota-based bonuses.

Organizations also change the components of their incentive systems frequently. About 80% of U.S. firms revise their compensation programs every two years or less to better motivate salespeople and to tailor their behavior to the goal of the sales organization (Zoltners et al., 2012).

However, simply providing the salespeople with an optimal menu of compensation is insufficient for an organization to achieve its desired sales outcome. To support the productivity of their sales force, organizations often rely on sales training, which serves to increase productivity, stimulate communication inside and outside the organization, reduce inter- and intra-department misunderstandings, enrich sales force morale, and decrease selling costs (Stanton & Buskirk, 1987; Churchill et al., 1993; Dubinsky, 1996). Organizations in the U.S. invest \$15 billion annually in sales training programs and devote more than 33 hours per year to training each salesperson (Lorge & Smith, 1998; Ingram et al., 2015). Thus, to effectively allocate resources across the sales management instruments, it is essential to properly assess and evaluate the outcomes of the organization’s sales training policy.

Whereas compensation and training serve to induce the right behavior, selection (recruitment/termination) affects the organization’s performance through changes in the sales force composition. Salespeople are known to exhibit a high rate of attrition: the estimated annual turnover rate of 27% is more than twice that of the average work force in the U.S. (Richardson, 1999). There are two types of employee selection: (i) firm-induced selection, which involves recruiting, retention, and termination; and (ii) employee-induced selection, or voluntary turnover. When properly managed, selection allows the organization to maintain a healthy sales force by retaining high-quality employees and terminating persistently low performers. However, selection—especially voluntary turnover—also involves substantial costs to the organization, including hiring and training expenditures, jeopardized customer relationships, and territory vacancies (Griffeth & Hom, 2001; Boles et al., 2012). Hence, deriving a proper policy to control for sales force selection is vital to the success of a sales organization.

Despite the ubiquitous use of the above sales management instruments, however, there is little insight into their joint effect on various behavioral outcomes. For instance, how should a firm design its compensation system to select the right salespeople—i.e., to retain the high, while discouraging the low, performers—over time? Which is more effective in motivating salespeople to meet their goals—increasing the level of monetary compensation or providing sales training opportunities? Can recruiting/termination policy replace the role of compensation and, if so, at what cost?

Separately identifying each of the above issues turns out to be problematic because various behavioral outcomes are often interrelated and occur simultaneously. Heterogeneous salespeople exhibit differences in productivity, time preference, and responsiveness to compensation components and training, which, in combination, determine an individual’s performance. The performance outcome results in compensation that influences employee attrition, which naturally leads to the selection of heterogeneous salespeople. This interrelated nature of behavioral outcomes necessitates an integrated model of sales force management.

There are two key challenges in modeling and identifying salespeople’s response to various management instruments. First, data at the salesperson level on various management practices are difficult to obtain, as many organizations treat HR information as confidential. As a result, previous studies have narrowed their focus to a single sales management instrument, such as compensation (Misra & Nair, 2011; Chung et al., 2014). Second, a researcher observes neither the agent’s effort nor his or her time preference—i.e., the degree to which immediate utility is favored over delayed utility. Rather, the researcher observes only the attrition decision and performance outcome over a specific period, both of which are likely correlated with the agent’s forward-looking allocation of effort and outside opportunities. This requires a behavioral assumption about the link between a sales agent’s motives (e.g., how close the person is to achieving quota at the end of the period) and his or her allocation of effort over time.

To overcome these challenges, we collaborate with a major multinational firm and formulate a comprehensive model of sales force behavior in response to various sales management practices. The model takes into account many of the key realistic elements of salespeople’s behavior, including allocation of effort, stay-or-leave decision, forward-looking behavior, present bias, and learning from training opportunities. Overall, we seek to gain insights into ways in which employee training, outside employment opportunities, and various elements of compensation jointly affect the selection and performance of heterogeneous salespeople.

This study also makes an important methodological contribution to the economics and marketing literatures. The study provides a formal proof on the identification conditions of a hyperbolic discounting model—a more general structure than an exponential discounting model. An agent’s distance-to-quota (DTQ), under a nonlinear incentive contract, affects only his or her future payoffs in non-pecuniary benefit periods, providing exclusion restrictions on current payoffs to identify the agent’s time preference. However, identifying time preference in a hyperbolic discounting model becomes challenging when confronted with an agent’s continuous choice (e.g., effort). Existing

studies are built largely upon a discrete choice setting (Fang & Wang, 2015; Abbring & Daljord, 2019), and, thus, the identification results do not fully translate. This study offers a formal discussion of the associated limitations and provides proper regularity conditions for identifying a hyperbolic discounting model under continuous choice. Building on the theoretical identification results, the empirical application shows support for agents’ hyperbolic discounting time preferences. A hyperbolic discounting model can potentially explain agents’ seemingly irrational behaviors (such as extreme procrastination) that are difficult to explain with a standard discounting model but are commonly observed in the real world (Ainslie, 1992; Kirby, 1997; Frederick et al., 2002).

The estimation results reveal the existence of different types of salespeople that possess heterogeneous utility and time preferences. A series of counterfactual experiments shows ways in which salespeople’s performance and selection change with alternative compensation plans, recruiting/termination policies, and sales training opportunities. The results demonstrate a trade-off relation between adjusting fixed and variable pay; a potential drawback of hiring only high-type salespeople; the short- and long-term outcomes of hiring rookies vs. experienced salespeople; how a collective leave package can lead to selective departure of the sales force; and how sales training can serve as an alternative to providing additional compensation.

A field validation—that compares the actual sales records (following changes in sales management instruments) with the simulated counterfactual outcomes—demonstrates the accuracy and applicability of the model. Hence, this study’s framework and model can provide a practical application for organizations to foresee the effect of multiple sales management instruments on the behavior of their sales force.

The remainder of the study’s structure is as follows: Section 2 summarizes the related literature. Section 3 describes the institutional settings and provides model-free evidence that facilitates the empirical analyses. Section 4 presents the modeling framework of sales management and an agent’s dynamic optimization problem. Section 5 illustrates the identification of dynamic models under exponential and hyperbolic time preferences. Section 6 describes the estimation procedure. Section 7 discusses the estimation results, counterfactual simulations, and field validation. Section 8 concludes.

2. Related Literature

This study on multi-dimensional sales force management contributes to several streams of research. First and foremost, the study relates to the strand of literature on sales force compensation.

The theoretical studies on this topic find conflicting results regarding components that constitute an optimal compensation system. Early works of Basu et al. (1985) and Rao (1990), under the principal-agent framework of Hölmstrom (1979), find that the optimal compensation system includes a salary and a nonlinear commission. More granularly, Hölmstrom and Milgrom (1987) and Lal and Srinivasan (1993) show that, in a multi-period setting, only a linear contract can achieve the first-best outcome. In contrast, Oyer (2000) finds that a compensation system with a quota-bonus and a linear over-achievement commission is uniquely optimal when the participation constraints are unbinding. More recently, Schöttner (2016) derives conditions under which a commission plan dominates a bonus plan and vice versa, depending on the degree of the agent’s responsiveness to incentives.

The findings of empirical studies also have discrepancies. Oyer (1998), using aggregate sales data, finds that quota-bonus compensation induces salespeople to manipulate the timing of sales, thereby negatively affecting productivity. On the contrary, Steenburgh (2008), analyzing individual-level data, finds that quota-bonus pay induces additional effort that provides net improvement in sales. A similar disparity is reported using dynamic models; for example, Misra and Nair (2011) and Chung et al. (2014) report contrasting findings regarding the effect of quota-bonuses on sales performance.

This study’s contribution to this literature is twofold. First, it expands the scope of outcomes to discuss the dynamic selection of sales agents, providing a better understanding of how an organization’s compensation system facilitates the restructuring of its sales force. In addition, the study examines agents’ effectiveness gain through training, along with their response to compensation. Both sales training and compensation serve as significant investments for organizations, and, thus, this study allows the evaluation of the relative effectiveness of these sales management instruments. To the best of our knowledge, this study is the first to jointly examine the effect of multiple sales management instruments on sales force selection and performance.

Since selection, by definition, accompanies employee turnover, this study relates to the strand of literature on the antecedents of sales force turnover. Existing studies have put emphasis on the negative aspects of salespeople’s departure. Richardson (1999) derives managerial measures for assessing the direct and indirect costs of turnover, and Darmon (2008) proposes a cost-benefit analysis of turnover for management efficiency. Using empirical analyses, Shi et al. (2017) finds that the negative effects vary, and Sunder et al. (2017) finds that turnover risk is the greatest for salespeople with moderate performance.

The abovementioned studies, however, are limited to evaluating the short-term effect of territory absence and potentially overlook the selection process that takes place simultaneously. That is, if a firm can select the right salespeople, then despite the short-term loss, employee turnover may result in greater long-term profitability. Hence, this study contributes to the appropriate evaluation of turnover by investigating salespeople’s latent future potential. In addition, the structural approach of the study allows for various counterfactual policy simulations, whereas descriptive studies limit this applicability.

This study also relates to the strand of literature on sales training effectiveness. Although various studies have emphasized the pivotal role of sales training on performance and have proposed conceptual frameworks (Walker et al., 1977; El-Ansary, 1993; Honeycutt et al., 1995; Attia et al., 2005), only a handful of empirical studies have followed this footprint, likely due to difficulties in collecting data. In addition, early empirical studies, relying mainly on survey measures, have generated strikingly mixed findings, ranging from a 50% increase in performance (Martin & Collins, 1991; Roman et al., 2003) to being largely uninfluential (Christiansen et al., 1996; Dubinsky, 1996).

More recently, Kumar et al. (2014), by examining the effect of voluntary training opportunity on the salesforce lifetime value, shows that sales training, indeed, has a positive effect in both the short term and the long term. However, the paper evaluates only the correlation between salespeople’s self-selected training participation and outcomes and refrains from developing a causal inference. To identify the causal effect of training, Atefi et al. (2018) conducts a controlled field experiment that varies training policies across retail stores. The paper finds a positive relation between sales outcomes and the proportion of salespeople who receive training; however, it analyzes sales training only at the aggregate (store) level due to the institutional and experimental settings.

This study provides several insights into the literature by measuring the comprehensive effects of sales training. First, by analyzing the training records at the individual level, the study examines the differential effect of training opportunities across heterogeneous salespeople. Second, the dynamic model allows us to analyze the long-term effects of training, which affects not only intertemporal performance outcomes, but also subsequent selection of the sales force. Third, the structural formulation of the model allows for cost-wise comparison between training and compensation policies, providing guidance to organizations on their resource allocation.

Finally, this study relates to the economics and psychology literatures on time preference and intertemporal decision-making. People discount future payoffs, and to capture the behavioral response to present versus future outcomes, researchers have used mainly two models of time

preference: exponential and hyperbolic discounting. The exponential discounting model assumes that people discount the future at a fixed rate per unit of time (Samuelson, 1937; Dhami, 2016), representing stationarity and time-consistent behavior. In contrast, the hyperbolic discounting model posits that people discount the immediate future from the present more than they do for the same time interval in the distant future (Ainslie, 1975; Thaler, 1981; Ainslie & Herrnstein, 1981; Loewenstein & Prelec, 1992; Laibson, 1997; O’Donoghue & Rabin, 1999), implying present bias and time-inconsistent behavior.

In terms of identifying time preference, Rust (1994) shows that the discount factor is generally not identified from naturally occurring data without further restrictions. Magnac and Thesmar (2002) generalizes this idea to provide conditions on exclusion restrictions—variables that do not affect an agent’s current payoff but only his or her future payoff—that allow the identification of the discount factor. Empirical studies in economics and marketing have applied the exclusion restrictions to identify time preference (discount factor) across various contexts, including new and used durable goods (Chevalier & Goolsbee, 2009; Ishihara & Ching, 2019); cellular phone usage (Yao et al., 2012); hardware and software platforms (Lee, 2013); sales force compensation (Chung et al., 2014); and consumer learning and inventory (Ching et al. 2014; Akca & Otter, 2015; Ching & Osborne 2019).

In a debate in the literature, Fang and Wang (2015) and Abbring and Daljord (2019) discuss identification of time preference in dynamic discrete choice models. Fang and Wang (2015), by extending the exclusion restriction arguments in Magnac and Thesmar (2002), considers conditions to identify various discounting behaviors, including exponential, hyperbolic, and naïve time preferences. Abbring and Daljord (2019) considers exclusion restrictions on model primitives and suggests that the arguments presented in Fang and Wang (2015), under weaker conditions, may not allow for point identification of the discount factor.

Different from the above studies (in which choice is directly observed), this study’s identification allows the agent’s action (e.g., effort) to be unobserved and only indirectly inferred from the observable output (e.g., performance). Regarding identification of models involving unobserved choice variables¹, Hu and Xin (2019) provides a general framework under which the conditional choice probability and the law of motion for state variables are separately identified. The paper’s identification leverages exclusion restrictions that affect only the conditional choice probability, but

¹ Empirical studies that examine unobserved choice variables in a dynamic setting include Misra and Nair (2011) and Chung et al. (2014), which analyze sales force behavior with effort unobserved, and Cosguner et al. (2018), which estimates a dynamic oligopoly pricing model in which retail prices are observed, but wholesale prices are not.

not the state transition probability. Similar to Hu and Xin (2019), this study identifies the agent’s unobserved effort using variation in output (sales performance) in response to his or her state (DTQ). The agent’s DTQ, under a nonlinear incentive contract, provides exclusion restrictions by affecting only future payoff (through the evolution of state variables), but not current-period payoff, in non-pecuniary benefit periods.

This study contributes to this stream of literature by expanding the scope of identification to a quasi-hyperbolic discounting model that incorporates continuous choice of the agent’s actions. The study shows the limitations in applying the results of identification in a discrete choice setting (Magnac & Thesmar, 2002; Fang & Wang, 2015; Abbring & Daljord, 2019) to a hyperbolic discounting structure that accommodates continuous choice and provides regularity conditions for identification. Building on the identification arguments, the empirical application presents support for agents’ hyperbolic discounting time preferences that exhibit heterogeneity in both the present-bias and the long-term discount factors.

3. Institutional Details and Descriptive Analysis

This section presents the focal institution’s sales environment, its compensation plan, and model-free evidence on forward-looking behavior and allocation of effort, which justify the dynamic structural formulation of the model.

3.1. Sales Environment

The firm under study is a multinational generic pharmaceutical company, offering a portfolio of branded prescription products through its own direct sales force. The data come from the firm’s sales operations in Turkey. Some notable aspects of the Turkish pharmaceutical market are worth mentioning. First, the government heavily regulates prices. Second, the nation’s universal healthcare system induces a high level of competition among the generics companies. Third, the country’s regulations prohibit direct-to-consumer (DTC) advertising, as is the case in most other markets.² As such, personal selling plays an important—and the *only*—role in the firm’s go-to-market strategy. Thus, recruiting and maintaining a sustainable pool of salespeople and training and motivating them properly are critical factors for success.

² As of 2018, only Brazil, New Zealand, and the United States allowed direct-to-consumer advertising, with varying restrictions on mode and content.

The data consist of salespeople’s performance, turnover, and hours of training during a three-year period (2015-2017). **Table 1** shows the number of employees who joined and departed and the corresponding turnover rate for each year. The firm’s average (voluntary) turnover rate was 14.60% over the three years.³ We focus our attention on those individuals who have remained in the firm (stay) and those who have voluntarily separated (quit). To minimize the effect of the initial learning curve, we discard individuals less than or equal to three months since hire (i.e., who joined on or after October 2017). The data-cleaning process leaves us with 554 salespeople. **Table 2** shows the corresponding descriptive statistics. Employees who decided to stay with the firm tend to perform better, receive higher variable pay, and have longer tenure.

The firm offers three types of sales training programs: primary training session; year-end sales session; and new sales-employee orientation. The 12-hour primary training session took place twice during the data observation period: one in January 2015, targeted at the salespeople from the primary care division (representing half of the entire sales force), and the other in April 2016, targeted towards senior salespeople across all divisions. The three-hour year-end sales session, which took place in December 2016, was mandatory for all salespeople. In 2017, the firm introduced a new sales-employee orientation program (designed for salespeople with less than one year of tenure) that ran for three hours.⁴ The estimated hourly cost of sales training per salesperson was \$37 in 2016.⁵

The firm operates its sales activity by route call sales: each salesperson has a preplanned series of meetings with either physicians or pharmacists in his or her exclusive territory. On average, a salesperson makes 20 calls per day. During each meeting, the salesperson exerts effort to promote the firm’s range of products.

3.2. The Firm’s Compensation Plan

The firm’s compensation plan consists of three components: base salary, quota-based bonus, and overachievement commission. **Figure 2** illustrates an overview of the plan, and **Table 3** describes the specifics of the quota-bonus payment schedule. The salespeople, on average, receive a fixed

³ To focus on salespeople’s behavior towards selection (voluntary turnover), we treat layoffs as a separate strategic decision by the firm and do not consider involuntary departures. The involuntary turnover rate of the firm was 7% in 2017; the majority of those laid off were salespeople in their probation period (less than one year since hire).

⁴ Because the firm chose participants in the primary training sessions based on an entire division or seniority, there exists unique variation in training hours that is exogenous to individual performance. In addition, the salespeople that joined the firm during the observation period add to this variation, as they missed the training opportunities in the earlier periods.

⁵ All monetary figures in this study are in U.S. dollars, converted using the exchange rate at the beginning of the data period (January 2015).

monthly salary of \$1,500. At the end of the first three quarters, a salesperson receives a \$1,700 bonus if he or she has attained the respective quotas.⁶ At the end of the year, the firm gives a \$3,400 bonus if the salesperson has met the annual quota. In addition, salespeople receive an overachievement commission of approximately \$170 (2% of the combined bonuses of \$8,500) per any excess percentage points above the annual quota. The firm caps the overachievement commission at \$8,500, attained when the salesperson’s performance (sales/quota realization) reaches 150%.

In setting and updating the agents’ quotas, the firm uses a well-established outside consulting company (that gathers all of the pharmaceutical sales data in the country, including the firm’s competitors’) to incorporate market-level information, such as current share, growth potential, and territorial and seasonal fluctuations in demand. By adjusting quotas based on objective measures (rather than on past sales performance), the firm mitigates possible ratcheting concerns.

Some features of the firm’s quota-based bonus system are noteworthy. First, quotas are set to be cumulative from the beginning of the year. Second, the firm defers the unearned bonus amount in each quarter to the subsequent quarter. That is, if a salesperson misses the quota in a given quarter, the respective bonus amount is added to the total amount attainable in the next quarter. For example, if a salesperson meets both Q1 and Q2 quotas, he or she would receive \$1,700 in both March and June. However, if a salesperson meets only the Q2 quota and not the Q1 quota, he or she would receive \$3,400 only in June.

This payout structure creates unique dynamics in the forward-looking behavior of the salespeople. On the one hand, it motivates them to keep up the pace from the beginning of the year. If a salesperson performs adequately and achieves bonus in a given period, his or her motivation remains intact due to the attainable quarterly bonus in the next period. If the salesperson does not meet quota in a given quarter, the motivation to exert effort becomes greater in the subsequent periods, as the total bonus amount increases due to the deferred bonus amounts from the previous quarters.

On the other hand, the cumulative nature of the performance evaluation also raises potential concerns in which poor performers lose motivation and give up. Because the sales/quota realization accumulates from the beginning of the year, the effect of several negative sales shocks can have a lasting effect throughout the year. This could demotivate the salespeople with poor performance

⁶ As shown in **Table 3**, salespeople receive a small fraction of the bonus when sales are at 90-99% of quota, starting from the second quarter. Hence, in strict terms, ‘to meet the quota’ would mean that sales are at or above 90% of quota. However, the firm avoids using this definition to discourage underachievers from believing that they are performing adequately. This study follows the firm’s terminology, indicating that a salesperson meets quota when sales are at or above 100% of quota.

during the early part of the year, whereas they would have received a fresh start under an independent quarterly-quota system.

3.3. Model-Free Evidence: Forward-Looking Behavior

If salespeople’s proximity to bonuses (i.e., distance-to-quota, DTQ) affects their performance in non-bonus periods, this suggests forward-looking behavior (Chung et al., 2014). Specifically, the state of DTQ would affect the performance of salespeople who have a reasonable chance of achieving the bonus. Hence, to show evidence of forward-looking behavior, we divide salespeople by their cumulative quota achieved (%QA): when $\%QA > 0.8$, salespeople have a reasonable probability of attaining the bonus at the end of each quarter, but when $\%QA < 0.8$, the chance is slim. In addition, the probability likely decreases as time passes and performance accumulates. **Table 4** reports the results of a regression analysis, with each column having monthly performance as the dependent variable and %QA by the previous month as the explanatory variable, separately for each group of salespeople who are $\%QA > 0.8$ and $\%QA < 0.8$. Hereafter, the term ‘performance’ denotes sales normalized by the agent’s corresponding monthly quota, which are used to construct the cumulative interim and annual quotas. As indicated in Section 3.2, quotas are set by a well-established outside consulting firm, taking into account territorial and seasonal fluctuations in demand.

Consistent with forward-looking behavior, the (state) variable %QA is significant throughout the year for salespeople with $\%QA > 0.8$. However, for those with $\%QA < 0.8$, %QA is significant only during earlier periods of the year. This is the case because, despite some bad outcomes during the earlier months, there still exists some probability of meeting the quota by achieving high performance for the remainder of the year. However, the chance of achieving the quota decreases as months with low performance accumulate, and, by mid-year, the low-performing salespeople ($\%QA < 0.8$) start to give up.

For a graphical illustration of forward-looking behavior, **Figure 3** displays the scatterplot and the best-fitting nonparametric smoothed polynomial (and its 95% confidence interval) of the salespeople’s performance in bonus-paying months on the %QA by the previous month. Three items stand out from **Figure 3**. First, from March through September, a considerable number of salespeople with low %QA achieve monthly performance greater than 100%. However, in December, very few in the lower group exhibit excess performance. Consistent with the results in **Table 4**, salespeople far from quota likely give up in December because they cannot achieve the annual quota in just a month. Second, a salesperson’s effort increases as he or she is on track to meet quota but

flattens once the quota is met ($\%QA > 1$). The proximity to bonuses (DTQ) motivates the salesperson, but once he or she surpasses quota ($\%QA > 1$), the motivation is no longer intact. Lastly, a salesperson's marginal effort with regard to his or her state ($\%QA$ by the previous month) increases with time in a calendar year. That is, the slope of the fitted line is steeper in December than in March. There are two likely reasons: (i) the presence of the overachievement commission in December motivates salespeople to exert greater effort; and (ii) the large year-end bonus (including the overachievement commission) is less discounted due to temporal proximity and, thus, motivates salespeople more towards the end of the year.

4. Model

This section presents a comprehensive model of a sales agent's behavior based on the sales management framework illustrated in **Figure 1**. The discussion proceeds in three parts: (i) the agent's per-period utility and performance response functions; (ii) dynamic allocation of effort and stay-or-leave decisions; and (iii) time preference.

An agent derives utility from compensation and disutility from effort and faces intertemporal employment (stay-or-leave) decisions. Compensation is nonlinear and dependent on the history of performance (e.g., quarterly sales outcomes). Hence, the agent exhibits forward-looking behavior and dynamically allocates effort.

4.1. Per-Period Utility and Performance Response

Agent i in period t derives per-period⁷ utility based on his or her choice of actions—whether to stay with the firm (d_{it}) and, (if so), how much effort to exert (e_{it})—such that

$$\tilde{U}(W_{it}, d_{it}, e_{it}, \varepsilon_{dit}) = \begin{cases} M(W_{it}) - C(e_{it}) + \varepsilon_{1it} & \text{if } d_{it} = 1, \\ \rho_i + \varepsilon_{0it} & \text{otherwise.} \end{cases} \quad (1)$$

If the agent decides to stay with the firm ($d_{it} = 1$), he or she receives positive pecuniary utility $M(W_{it})$ as a function of compensation W_{it} . The amount of compensation $W_{it} = W(q_{it}, s_{it}; \psi_{it})$ is determined by the agent's performance q_{it} and state s_{it} , given the firm's compensation scheme ψ_{it} . Concurrently, the agent incurs disutility $C(e_{it})$ from exerting effort e_{it} , which affects the performance outcomes in the contemporaneous period. If the agent decides to quit ($d_{it} = 0$), he or she receives reservation value ρ_i in perpetuity. The reservation value represents the agent's outside option. The

⁷ The period in the empirical application is a month.

decision to leave the firm is an absorbing state (i.e., permanent), and, thus, the agent cannot return to the firm once the action is taken.⁸

In addition to the deterministic elements, the per-period utility includes a structural error term ε_{dit} , which represents the state unobserved by the researcher, but observed by the agent in his or her stay-or-leave decision d_{it} .⁹ The structural error follows a Type-I extreme value distribution with location parameter zero and scale parameter σ_ε and is assumed to be independently and identically distributed across choices and agents over time.

The agent's per-period performance q_{it} is a function of his or her individual effect α_i , effort e_{it} , and an idiosyncratic performance shock ξ_{it} such that

$$q_{it} = \exp(\alpha_i + e_{it} + \xi_{it}), \quad (2)$$

or in logarithmic terms, $\ln(q_{it}) = \alpha_i + e_{it} + \xi_{it}$. The log-linear specification allows the agent's sales performance to be always positive,¹⁰ consistent with the empirical setting.

The individual effect (heterogeneity) α_i represents the agent's baseline ability (i.e., performance attained without any effort).¹¹ The performance shifters x_{it} affect individual heterogeneity such that $\alpha_i = \alpha_0 + \alpha_1 x_{it}$, where x_{it} includes the agent's tenure, training, tenure-training interaction, and level of higher education. An agent's cumulative hours of sales training forms the training variable to capture the long-run persistence effect. By this structure, the training hours accumulate to form the agent's stock of expertise, which carries over to the subsequent periods and affects his or her performance over time. The distribution of the performance shock ξ_{it} (common knowledge to the agent¹²) follows $N(0, \sigma_\xi^2)$ and is independent of the agent's state $(s_{is}, \alpha_i, \varepsilon_{is})$ and effort e_{is} for any $s \leq t$.

⁸ No agent in the data returned after departing from the firm.

⁹ As is standard in the literature, the error term satisfies the conditional independence assumption (Rust, 1987) in that, in a given period, it is not a function of an agent's effort allocation decision (e_{it}). That is, the error term (unobserved state) realizes ex-ante of the agent's current-period effort decision.

¹⁰ An agent can obtain positive sales performance with no effort under various contexts, including customers' (i) need-based purchases without any salesperson interaction; and (ii) repeat purchases based on previously built relationships.

¹¹ Strictly speaking, individual heterogeneity can also be interpreted as the baseline level of effort. As we cannot directly observe effort, the two effects (baseline ability and baseline effort) are not distinguishable from each other. Thus, effort e_{it} represents the agent's additional contribution to performance from the baseline.

¹² More formally, the agent has rational expectation on the law of motion: the agent knows the distribution of the performance shock ξ_{it} , which affects the transition probability of future states.

By the performance response function in **Equation (2)**, the agent's *unobserved* effort e_{it} is (stochastically) linked to his or her *observed* performance q_{it} in a given period. The performance outcome q_{it} , in turn, has both (i) a direct effect on contemporaneous compensation $W_{it} = W(q_{it}, s_{it}; \psi_{it})$ in bonus/commission periods; and (ii) an indirect effect on future compensation through the evolution of the state variables $s_{i,t+1} = f(q_{it}, s_{it}; \psi_{it})$, where $f(\cdot)$ is the state transition function.

Equation (1) represents the *ex-post* utility of the agent, as the performance shock ξ_{it} in **Equation (2)**, which affects W_{it} through q_{it} in a given period, has yet to be realized when making the stay-or-leave and effort decisions. To form the basis of decision-making, the agent takes *expectation* over his or her compensation $W_{it} = W(q_{it}, s_{it}; \psi_{it})$, given effort e_{it} (which determines performance outcome q_{it} under current state s_{it}). In this manner, the *ex-ante* utility function of the agent is

$$U(d_{it}, e_{it}, s_{it}, \varepsilon_{dit}) = \begin{cases} E[M(W_{it}) | e_{it}, s_{it}] - C(e_{it}) + \varepsilon_{1it} & \text{if } d_{it} = 1, \\ \rho_i + \varepsilon_{0it} & \text{otherwise,} \end{cases}$$

where the functions M and C take on a parametric functional form. The pecuniary utility M of wealth W_{it} takes the form of mean-variance utility such that

$$E[M(W_{it}) | e_{it}, s_{it}] = E[W_{it} | e_{it}, s_{it}] - \gamma_i \text{Var}[W_{it} | e_{it}, s_{it}],$$

where $\gamma_i > 0$ represents the agent's risk preference.¹³ The disutility C is specified to be convex in effort e_{it} , such that

$$C(e_{it}) = \theta_i e_{it}^2,$$

where $\theta_i > 0$ denotes the agent's ease and flexibility in exerting effort. An implicit benefit of the mean-variance utility specification is that it provides, by construction, scale and location normalization of utility. This allows us to estimate, rather than to normalize, the agent's reservation value ρ_i and the scale parameter σ_ε of the structural errors (see Section 5.4 for a detailed discussion of identification).

Given these specifications, the ex-ante utility (hereafter simply referred to as utility) function can be represented as

$$U(d_{it}, e_{it}, s_{it}, \varepsilon_{dit}) = \begin{cases} E[W_{it} | e_{it}, s_{it}] - \gamma_i \text{Var}[W_{it} | e_{it}, s_{it}] - \theta_i e_{it}^2 + \varepsilon_{1it} & \text{if } d_{it} = 1, \\ \rho_i + \varepsilon_{0it} & \text{otherwise.} \end{cases} \quad (3)$$

¹³ The mean-variance utility represents a second-order approximation to a general concave utility function with constant absolute risk-aversion (CARA).

The reservation value shifters z_i affect the agent's reservation value ρ_i , such that $\rho_i = \rho_{0i} + \rho_1 z_{it}$, where ρ_{0i} represents agent i 's baseline reservation value. Reservation value shifters z_{it} include tenure and the level of higher education.

4.2. Compensation and State Variables

This study focuses on a class of nonlinear compensation schemes, the payout of which depends on aggregate performance over a specific time horizon. These schemes typically include variable pay components such as quarterly/annual bonuses or end-of-year commissions, which are commonly administered in practice (Joseph & Kalwani, 1998). By providing a reward at the end of a quota evaluation cycle consisting of multiple periods, the compensation scheme stimulates the sales agent's forward-looking behavior, as the agent's effort exerted today influences his or her future payoff. The accumulation of effort is captured by a subset of the state variables in s_{it} , whose subsequent-period values $s_{i,t+1}$ evolve as a function of current-period performance q_{it} and state s_{it} .

The firm's incentive scheme (ψ_{it}) includes the following components: (i) individual-specific monthly base salary w_{it} ; (ii) maximum attainable quarterly-bonus amount Q_t (including the deferred amount from previous periods), common across all agents but varying across years; (iii) quarterly-bonus payout rate R_{qt} ; and (iv) end-of-year overachievement commission rate R_{yt} .¹⁴

Formally, the components Q_t , R_{qt} , and R_{yt} are as follows:

$$Q_t = \begin{cases} 1,700 & \text{if } s_{i1t} = 3, \\ 3,400 & \text{if } s_{i1t} = 6, \\ 5,100 & \text{if } s_{i1t} = 9, \\ 8,500 & \text{if } s_{i1t} = 12, \\ 0 & \text{otherwise,} \end{cases}$$

¹⁴ Although illustrated based on the institutional setting, our model is applicable to a wide class of nonlinear compensation schemes.

$$R_{qt} = \begin{cases} 0.35 & \text{if } 0.90 \leq s_{i2,t+1} < 0.91, \text{ and } s_{i1t} = 12 \\ 0.41 & \text{if } 0.91 \leq s_{i2,t+1} < 0.92, \text{ and } s_{i1t} = 12 \\ 0.47 & \text{if } 0.92 \leq s_{i2,t+1} < 0.93, \text{ and } s_{i1t} = 12 \\ 0.53 & \text{if } 0.93 \leq s_{i2,t+1} < 0.94, \text{ and } s_{i1t} = 12 \\ 0.59 & \text{if } 0.94 \leq s_{i2,t+1} < 0.95, \text{ and } s_{i1t} = 12 \\ 0.65 & \text{if } 0.95 \leq s_{i2,t+1} < 0.96, \text{ and } s_{i1t} \in \{6, 9, 12\} \\ 0.72 & \text{if } 0.96 \leq s_{i2,t+1} < 0.97, \text{ and } s_{i1t} \in \{6, 9, 12\} \\ 0.79 & \text{if } 0.97 \leq s_{i2,t+1} < 0.98, \text{ and } s_{i1t} \in \{6, 9, 12\} \\ 0.86 & \text{if } 0.98 \leq s_{i2,t+1} < 0.99, \text{ and } s_{i1t} \in \{6, 9, 12\} \\ 0.93 & \text{if } 0.99 \leq s_{i2,t+1} < 1.00, \text{ and } s_{i1t} \in \{6, 9, 12\} \\ 1.00 & \text{if } 1.00 \leq s_{i2,t+1}, \text{ and } s_{i1t} \in \{3, 6, 9, 12\} \\ 0 & \text{otherwise,} \end{cases}$$

$$R_{yt} = \begin{cases} 1.00 & \text{if } 1.00 \leq s_{i2,t+1} < 1.01, \text{ and } s_{i1t} = 12, \\ 1.02 & \text{if } 1.01 \leq s_{i2,t+1} < 1.02, \text{ and } s_{i1t} = 12, \\ \vdots & \\ 1.98 & \text{if } 1.49 \leq s_{i2,t+1} < 1.50, \text{ and } s_{i1t} = 12, \\ 2.00 & \text{if } 1.50 \leq s_{i2,t+1}, \text{ and } s_{i1t} = 12, \\ 0 & \text{otherwise,} \end{cases}$$

where the state variable s_{i1t} denotes the month-type ($\{1, 2, \dots, 12\}$) and s_{i2t} denotes the percentage of cumulative quota achieved (%QA) by the end of the previous month. The above components are collected to represent the firm's incentive scheme by the vector $\psi_{it} = \{w_{it}, Q_t, R_{qt}, R_{yt}\}$.

Given the incentive scheme ψ_{it} , an agent receives compensation $W_{it} = W(q_{it}, s_{it}; \psi_{it})$ based on performance q_{it} and state s_{it} . Compensation W_{it} is comprised of three components: (i) monthly base salary w_{it} ; (ii) quarterly (and annual) bonus QB_{it} ; and (iii) end-of-year overachievement commission OC_{it} , in the following form:

$$W_{it} = w_{it} + QB_{it} + OC_{it},$$

whose elements QB_{it} and OC_{it} are expressed as follows:

$$QB_{it} = \max \left\{ Q_t \cdot R_{qt} \left(\frac{(s_{i1t} - 1) \cdot s_{i2t} + q_{it}}{s_{i1t}} \right) - s_{i3t}, 0 \right\},$$

$$OC_{it} = Q_t \cdot R_{yt} \left(\frac{(s_{i1t} - 1) \cdot s_{i2t} + q_{it}}{s_{i1t}} \right),$$

where s_{i3t} represents the amount of bonus accrued (%BA) in previous quarters (limiting the deferral of the quarterly-bonus amount if the agent previously received the bonus). Note that in non-bonus

periods, $QB_{it}=0$, and, thus, W_{it} depends solely on w_{it} ; note, also, that OC_{it} is distributed only in December.

The state variables directly linked to compensation include: (i) the month-type within the year, s_{1t} ; (ii) the percentage of cumulative quota achieved (%QA), s_{2t} ; and (iii) the amount of annual bonus accrued (%BA), s_{3t} .

The state variables evolve as follows:

1. Month-type

$$s_{1t} = \begin{cases} 1 & \text{if } t \text{ is the start of the year,} \\ s_{1(t-1)} + 1 & \text{otherwise.} \end{cases}$$

2. Percentage of cumulative quota achieved (%QA)

$$s_{2t} = \begin{cases} 0 & \text{if } t \text{ is the start of the year,} \\ \frac{s_{1(t-2)} \cdot s_{2(t-1)} + q_{i(t-1)}}{s_{1(t-1)}} & \text{otherwise.} \end{cases}$$

3. Percentage of annual bonus accrued (%BA)

$$s_{3t} = \begin{cases} 0 & \text{if } t \text{ is the start of the year,} \\ \max \left\{ Q_{(t-1)} \cdot R_{q(t-1)} \left(\frac{s_{1(t-2)} \cdot s_{2(t-1)} + q_{i(t-1)}}{s_{1(t-1)}} \right), s_{3(t-1)} \right\} & \text{otherwise.} \end{cases}$$

Whereas the month-type evolves in a self-contained manner, the latter two state variables evolve (stochastically) based on the agent's effort. The percentage of cumulative quota achieved (%QA) evolves every month, based on the performance in previous periods. The percentage of annual bonus accrued (%BA) evolves stepwise every quarter, based on receiving the quarterly bonus. The state variables that directly affect compensation are represented by the vector $s_{it} = \{s_{1t}, s_{2t}, s_{3t}\}$.

4.3. Dynamic Allocation of Actions

The per-period utility function in **Equation (3)**, when linked with the aforementioned course of actions, outcomes, and state transitions, naturally leads to a dynamic formulation of the model. An agent chooses actions that solve the dynamic optimization problem, maximizing the sum of current and future payoffs over discrete time periods ($t=1,2,\dots,\infty$). The value function is defined as the agent's discounted present value of the expected utility stream (given states s_{it} and ε_{dit}) such that

$$\begin{aligned}\tilde{V}(s_{it}, \varepsilon_{dit}) &= \mathbb{E} \left[\sum_{\tau=t}^{\infty} \phi(\tau - t) \left\{ \max_{d_{i\tau}, e_{i\tau}} U(d_{i\tau}, e_{i\tau}, s_{i\tau}, \varepsilon_{di\tau}) \right\} \middle| s_{it}, \varepsilon_{dit} \right] \\ &= \max_{d_{it}, e_{it}} \left\{ U(d_{it}, e_{it}, s_{it}, \varepsilon_{dit}) + \mathbb{E} \left[\sum_{\tau=t+1}^{\infty} \phi(\tau - t) \left\{ \max_{d_{i\tau}, e_{i\tau}} U(d_{i\tau}, e_{i\tau}, s_{i\tau}, \varepsilon_{di\tau}) \right\} \middle| d_{it}, e_{it}, s_{it}, \varepsilon_{dit} \right] \right\},\end{aligned}$$

where $\phi(j)$ denotes the discount function for utility from future j -periods forward ($j=0, 1, 2, 3, \dots$) and $\phi(0)=1$. Hence, the agent's value is represented by the expected utility flow upon making an infinite sequence of optimal decisions $(d_{i\tau}, e_{i\tau} : \tau=t, t+1, \dots)$, governed by the discount function $\phi(\cdot)$. The expectation is taken with regard to both the idiosyncratic performance shock ξ and the structural error ε for each period $\tau \geq t+1$.

The choice-specific value with respect to action pair (d_{it}, e_{it}) , which represents the discounted present value when the agent chooses actions d_{it} and e_{it} , is defined as

$$V(d_{it}, e_{it}, s_{it}, \varepsilon_{dit}) = U(d_{it}, e_{it}, s_{it}, \varepsilon_{dit}) + \mathbb{E} \left[\sum_{\tau=t+1}^{\infty} \phi(\tau - t) \left\{ \max_{d_{i\tau}, e_{i\tau}} U(d_{i\tau}, e_{i\tau}, s_{i\tau}, \varepsilon_{di\tau}) \right\} \middle| d_{it}, e_{it}, s_{it}, \varepsilon_{dit} \right]. \quad (4)$$

In each period, the agent incorporates the information contained in the current states $(s_{it}, \varepsilon_{dit})$ to evaluate the future outcome of current-period actions: employment (d_{it}) and effort (e_{it}) .

The agent's effort policy (the optimal level of effort), as a function of the state s_{it} and the stay-or-leave decision d_{it} , is given by

$$e_{it} = e(d_{it}, s_{it}) = \begin{cases} \arg \max_e \{V(1, e, s_{it}, \varepsilon_{1it})\} & \text{if } d_{it} = 1, \\ 0 & \text{otherwise.} \end{cases}$$

That is, the agent chooses the *optimal* level¹⁵ of effort e_{it} , which maximizes the discounted stream of expected utility flow, conditional on the current states and on staying with the firm. The temporal trade-off of exerting effort e_{it} (in non-bonus/commission periods) arises between the per-period disutility $C(e_{it})$ in **Equation (1)** and the state-transition $s_{i,t+1}$ (updated through the performance outcome q_{it} in **Equation (2)**) towards a higher probability of future pecuniary benefits.

The agent decides to continue with the firm if the choice-specific value of staying and exerting effort, $V(1, e_{it}, s_{it}, \varepsilon_{1it})$, is greater than the value of leaving, $V(0, 0, s_{it}, \varepsilon_{0it})$ ¹⁶. That is,

¹⁵ For brevity, we suppress the optimality notation (*) throughout the study.

¹⁶ Note that once the agent leaves the firm ($d_{it}=0$), the absorbing state implies that (i) effort $e_{it}=0$ in all subsequent periods; and (ii) the recursive formulation in **Equation (4)** degenerates to receiving the reservation value ρ_i in perpetuity.

$$d_{it} = \begin{cases} 1 & \text{if } V(1, e_{it}, s_{it}, \varepsilon_{1it}) \geq V(0, 0, s_{it}, \varepsilon_{0it}), \\ 0 & \text{otherwise.} \end{cases}$$

The summary of the model dynamics is as follows: After observing his or her current state, an agent exerts effort and incurs disutility. Exerted effort, in combination with an idiosyncratic shock, determines the agent's current-period sales performance. This performance affects both the current-period payoff and the probability distribution of state variables in the subsequent period. Hence, the agent's effort helps preserve his or her state in a healthy condition, increasing the chance of receiving a monetary payoff in later periods. However, if the current state shows a limited chance of receiving future payoffs (e.g., after several periods of low performance), the agent may stop exerting effort in order to reduce disutility. Furthermore, if the value of staying becomes lower than the outside option, the agent will decide to leave the firm.

4.4. Time Preference

The above forward-looking model naturally leads to a conceptual question: How does an agent discount the stream of future utility to derive the optimal policy? In other words, what is the agent's time preference, the degree to which immediate utility is favored over delayed utility? The question can be addressed through varying the structure of $\phi(j)$, the discount function in **Equation (4)**. We consider two models of time preference: exponential discounting and quasi-hyperbolic discounting.

4.4.1. Exponential Discounting

The exponential discounting model (Samuelson, 1937) postulates that an agent's discount function for the j -th future period takes the form

$$\phi(j) = \delta^j \text{ for } j = 0, 1, 2, \dots,$$

where $\delta \in (0, 1)$. The model implies time-consistent behavior by featuring stationary discounting (geometric decay) over expected future utility. Because of its analytical convenience, exponential discounting is frequently assumed in the economics and marketing literatures.

The dynamic optimization problem can be decomposed into an infinite sequence of single-period decisions. Assuming exponential discounting, the infinite sum of the discounted future utility flow in **Equation (4)** can be replaced by the subsequent-period value function such that

$$\begin{aligned} V(d_{it}, e_{it}, s_{it}, \varepsilon_{dit}) &= U(d_{it}, e_{it}, s_{it}, \varepsilon_{dit}) + \mathbb{E} \left[\sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \left\{ \max_{d_{i\tau}, e_{i\tau}} U(d_{i\tau}, e_{i\tau}, s_{i\tau}, \varepsilon_{di\tau}) \right\} \middle| d_{it}, e_{it}, s_{it}, \varepsilon_{dit} \right] \\ &= U(d_{it}, e_{it}, s_{it}, \varepsilon_{dit}) + \delta \mathbb{E} \left[\max_{d_{i,t+1}, e_{i,t+1}} V(d_{i,t+1}, e_{i,t+1}, s_{i,t+1}, \varepsilon_{di,t+1}) \middle| d_{it}, e_{it}, s_{it}, \varepsilon_{dit} \right]. \end{aligned}$$

Henceforth, for brevity of exposition, subscripts i and t are suppressed, and the subsequent period $(t+1)$ is denoted by a prime ($'$) symbol when possible.

Let $v(d, e, s)$ denote the deterministic portion of the choice-specific value in **Equation (4)** (i.e., $v(\cdot) = V(\cdot) - \varepsilon_{dit}$) and define it as the *choice-specific value function*. Similarly, let $u(d, e, s)$ denote the deterministic portion of the utility function (i.e., $u(\cdot) = U(\cdot) - \varepsilon_{dit}$). Assuming additive separability and serial independence of the structural errors, the above equation simplifies to

$$v(d, e, s) = u(d, e, s) + \delta E \left[\max_{d', e'} \{v(d', e', s') + \varepsilon'_d\} \middle| d, e, s \right]. \quad (5)$$

4.4.2. Quasi-Hyperbolic Discounting

The quasi-hyperbolic discounting model (Phelps & Pollak, 1968; Laibson, 1997) posits that an agent's discount function for the j -th future period takes the form

$$\phi(j) = \begin{cases} 1 & \text{if } j = 0, \\ \beta \delta^j & \text{if } j = 1, 2, 3, \dots, \end{cases}$$

where $\delta \in (0, 1)$ is the standard discount factor, and $\beta \in (0, 1]$ is the present-bias factor. Often referred to as the Beta-Delta preference, the model parsimoniously captures present bias and, thus, time-inconsistency. The standard discount factor δ captures long-term, time-consistent discounting; and the present-bias factor β captures short-term impatience and the discontinuity between the present and the future (O'Donoghue & Rabin, 1999). Note that exponential discounting is a special case of quasi-hyperbolic discounting when $\beta=1$ (i.e., the agent is not present-biased).

Given quasi-hyperbolic discounting, the choice-specific value in **Equation (4)** becomes

$$v(d, e, s) = u(d, e, s) + \beta \delta E \left[\max_{d', e'} \{\tilde{v}(d', e', s') + \varepsilon'_d\} \middle| d, e, s \right], \quad (6)$$

where

$$\tilde{v}(d, e, s) = u(d, e, s) + \delta E \left[\max_{d', e'} \{\tilde{v}(d', e', s') + \varepsilon'_d\} \middle| d, e, s \right]. \quad (7)$$

Unlike the case of exponential discounting, however, the quasi-hyperbolic discounting model does not allow a recursive representation of a single value function. The flow of future utility involves an additional value function $\tilde{v}(\cdot)$ due to the agent's time-inconsistency. Hence, the optimal choice of effort e in the present becomes different from that of the future (i.e., the agent is present-biased).

The structure in **Equations (6) and (7)** requires solving two equations for two functions. This leads to a challenge in identification, which we discuss, in detail, in the following section.

5. Identification

This section presents the formal identification arguments and proceeds in the following order. First, we discuss the primitives of static utility—performance response, pecuniary utility of wealth, and disutility of effort—and then the agents’ time preferences—both exponential and quasi-hyperbolic. Finally, we provide an intuitive discussion of identification. The formal arguments build upon those of Magnac and Thesmar (2002), who propose exclusion restrictions to identify the standard (exponential) discount factor. We expand identification of time preference to consider the present-bias factor in a quasi-hyperbolic discounting model that accommodates continuous choice. The Appendix provides proofs regarding formal arguments.

5.1. Static Utility

Suppose that the data consist of $(d_{it}, s_{it}, q_{it}, \psi_{it})^{17}$ for agents $i = 1, 2, \dots, N$ over time $t = 1, 2, \dots, T$, and that these observations are independent and identically distributed across agents. We first consider the identification of the performance response function in **Equation (2)**. The challenge in identifying the performance response of unobserved effort e_{it} is in controlling for individual heterogeneity α_i . Because one does not directly observe either construct, separately identifying effort from individual heterogeneity typically is infeasible without further restrictions. The issue becomes further complicated because an agent’s effort policy is likely a function of individual heterogeneity α_i . That is, an agent takes into account his or her own baseline productivity when making effort decisions.

The agent’s behavior under a nonlinear compensation scheme provides conditions for the effort policy to be separately identified from individual heterogeneity. The idea is to exploit observations in which the optimal effort is trivially a corner solution. Consider the following assumption:

¹⁷ Note that the agent’s state variables s_{it} are computed given his or her performance history q_{it} and the firm’s compensation scheme ψ_{it} .

Assumption 1 (Corner Solution). *Suppose that there exists a subset S_α with a positive probability measure in the support of state variables s such that, if $s \in S_\alpha$, the derivative of the value function with respect to e is non-positive—i.e., $\frac{\partial v(d, e, s, \Omega)}{\partial e} < 0$ for any $e \geq 0$.*

That is, if s takes a value in S_α , the agent exerts zero effort. The set S_α exists when the agent is far above the quota (i.e., the bonus is already attained) or far below the quota (i.e., the bonus is not within reach). In either state, the agent’s additional performance provides limited gains and, thus, the agent is better off not incurring any effort (and avoiding the associated disutility).

Proposition 1. *Under Assumption 1, (i) the agent’s effort policy e_{it} , (ii) individual heterogeneity α_i , and (iii) the distribution of performance shock ξ_{it} are identified.*

Proof. See Appendix A. The proof is based on a nonparametric regression approach in a similar vein to Hu and Xin (2019).

Proposition 1 governs the relation among unobserved effort, individual heterogeneity, and observed performance, and stipulates the forward-looking behavior of the agents. From Proposition 1, the agent’s effort policy conditional on staying with the firm, $e_{it} = e(1, s_{it})$, is identified. Even if an agent leaves the firm, the optimal effort is identified during the period in which the agent stayed with the firm.

Regarding the identification of the choice-specific value function in **Equation (5)**, the following lemma holds:

Lemma 1. *Under Assumption 1, the difference in choice-specific value functions, $v(1, e_{it}, s_{it}) - v(0, 0, s_{it})$, is nonparametrically identified up to scale at the optimal level of effort e_{it} (identified in Proposition 1).*

Proof. See Appendix B. The proof uses the conditional choice probability approach in Magnac and Thesmar (2002).

By the model specification in **Equation (3)**, the value of leaving, $v(0, 0, s_{it})$, is an unknown constant that does not depend on the agent’s effort choice or state variables. Hence, Lemma 1 implies that the value of staying, $v(1, e_{it}, s_{it})$, is identified up to location and scale at the optimal level of effort. Given nonparametric identification of the choice-specific value function v , what remains for identification are the primitives of the utility function and time preference (discount factor).

5.2. Exponential Discounting Model

The identification of the exponential discounting model materializes from the exclusion restriction (Magnac & Thesmar, 2002) provided by a nonlinear compensation scheme.

Assumption 2 (Exclusion Restriction). *Suppose that the state variables s can be partitioned into two vectors, s_1 and s_2 , where s_1 is a vector of variables that satisfies the following condition: there exists a subset S_1 in the support of s_1 such that, if $s_1 \in S_1$, both*

- (i) $u(d, e, s_1, s_2) = u(d, e, s_1, \tilde{s}_2)$ for any s_1 and $\tilde{s}_2$, and
- (ii) $v(d, e, s_1, s_2) \neq v(d, e, s_1, \tilde{s}_2)$ for some s_1 and $\tilde{s}_2$ hold.

That is, if s_1 takes a value in S_1 , s_2 does not affect the present utility. In the empirical application, the variables month type and an agent's DTQ play the role of s_1 and s_2 , respectively. For example, there is no performance-based lump-sum payment in October or November, so the DTQ does not affect the per-period utility. In these months, the per-period utility depends only on the disutility of effort, which does not include s . However, the future expected utilities would differ according to the DTQ.

Proposition 2. *Suppose that the agent's true time preference follows the exponential discounting model. Under assumptions 1-2, the instantaneous utility function u at the optimal effort and the discount factor δ are nonparametrically identified.*

Proof. See Appendix C. The proof uses exclusion restrictions similar to Corollary 3 and Proposition 4 in Magnac and Thesmar (2002).

Once u is nonparametrically identified, the parametric identification of the mean-variance and reservation value within u is straightforward. As there is no variation in pecuniary utility during non-bonus periods, we have θ identified in those months. Parameter γ is identified during the bonus-paying months. The remaining ρ and σ_ε are identified using the variation in agents' stay-or-leave decisions (see Section 5.4 for an intuitive discussion).

5.3. Quasi-Hyperbolic Discounting Model

Under the quasi-hyperbolic discounting model in **Equations (6) and (7)**, there exist two value functions, v and $\tilde{v}$, and two discount factors, β and δ . The quasi-hyperbolic discounting model is

more general than the exponential discounting model, as the latter is a special case of the quasi-hyperbolic discounting model when $\beta=1$.

As Lemma 1 equally applies to the quasi-hyperbolic discounting model, the choice-specific value function v is nonparametrically identified (up to location and scale). However, identification of $\tilde{v}$ is not as straightforward, as **Equations (6) and (7)** form a system of equations. Multiplying **Equation (7)** by β and subtracting it from **Equation (6)** yields

$$v(d, e, s) - \beta \tilde{v}(d, e, s) = (1 - \beta)u(d, e, s).$$

Because $\beta > 0$, the above equation simplifies to

$$\tilde{v}(d, e, s) = \frac{\beta - 1}{\beta} u(d, e, s) + \frac{1}{\beta} v(d, e, s).$$

Finally, inserting the above equation into **Equation (6)** establishes that

$$v(d, e, s) = u(d, e, s) + \delta E \left[\max_{d', e'} \{ (\beta - 1)u(d', e', s') + v(d', e', s') + \beta \varepsilon'_d \} \middle| d, e, s \right]. \quad (8)$$

In **Equation (8)**, the distribution of ε'_d and the value function v are known, whereas the per-period utility u and the discount factors, δ and β , are unknown. Because this equation summarizes the system of equations, u , δ , and β are identified if there exists a unique solution to **Equation (8)**.

In mathematics, the structure of **Equation (8)** is known as a nonlinear Fredholm integral equation of the second kind (Arfken & Weber, 1999; Polyanin & Manzhirov, 1998; Vetterling et al., 1992). Solving the integral equation for the unknown utility function $u(d, e, s)$ is an ill-posed inverse problem due to the maximum function and integration taken over the utility function. Because a lot of information is “integrated out” and naturally lost during the process, it is well known that the solution to this ill-posed inverse problem may not exist, or even if a solution exists, it may not be unique.¹⁸

The essence of this problem arises due to the *continuity* of the choice variable. If the choice variable is discrete, the integral equation in **Equation (8)** can be replaced by a matrix algebra, and the problem is simplified to finding the inverse of the matrix. For example, Abbring and Daljord (2019) and Fang and Wang (2015), in a discrete choice setting, rely on matrix algebra to find the inverse for identification. This is not applicable to our setting—in which the choice variable is continuous—as solving for the inverse of an integral equation is ill-posed. Thus, without further

¹⁸ Conceptually, obtaining the solution to this problem is equivalent to finding an inverse mapping of the nonlinear integral. Even if there exists a unique solution, it is known to be extremely difficult, if not impossible, to obtain.

restrictions, the utility function and the discount factors cannot be nonparametrically identified, even if the exclusion restrictions hold. Intuitively, the ill-posed problem is due to the fact that the continuous choice in the model requires the utility function to be an infinite dimensional object (absent parametric assumptions), whereas in a discrete choice model, the utility function is represented by a finite dimensional vector (as in Fang and Wang (2015)). Because of this difference, a finite number of exclusion restrictions is insufficient to nonparametrically identify the utility function of a continuous choice hyperbolic discounting model.

The exponential discounting model bypasses this issue because the utility function does not enter the integral due to its recursive nature. That is, the value function for the future payoffs is identified directly from the choice probabilities. In contrast, in the quasi-hyperbolic discounting model, the utility function enters the integral as in **Equation (8)**. This change in the value function creates complications in solving for the equation, leading to uncertainty about the existence of the solution and, if it does exist, its uniqueness.

A typical solution for an ill-posed inverse problem is “regularization.” In a broad sense, to regularize is to provide additional assumptions that can aid the existence, uniqueness, and numerical stability of a solution. Some common examples include discretization of variables (Magnac & Thesmar, 2002; Fang & Wang, 2015; Abbring & Daljord, 2019), eigenvalue-eigenfunction decomposition (Hu & Schennach, 2008; Hu & Xin, 2019), parameterization of functions, and Lasso-type penalization methods.

The parametric assumption in **Equation (3)** on the per-period utility function u serves as a regularization to identify quasi-hyperbolic time preference under continuous choice. To illustrate, given the parameter vector $\mu = (\gamma, \theta, \rho, \sigma_\varepsilon)$, the agent’s optimal effort (in the subsequent period) conditional on staying with the firm is

$$e'(s \mid \mu) = \arg \max_e \{(\beta - 1)u(1, e, s) + v(1, e, s)\}.$$

Note that prior to parametrization, this optimal effort for the subsequent period was intractable.

Given the extreme value distribution assumption, the future payoff component within the expectation in **Equation (8)**, conditional on s' , now becomes

$$\begin{aligned} & \max_{d', e'} \{(\beta - 1)u(d', e', s' \mid \mu) + v(d', e', s') + \beta \varepsilon_{d'}'\} \\ &= \beta \sigma_\varepsilon \log \left\{ \exp \left(\frac{(\beta - 1)u(1, e', s' \mid \mu) + v(1, e', s')}{\beta \sigma_\varepsilon} \right) + \exp \left(\frac{(\beta - 1)\rho_i + v(0, 0, s')}{\beta \sigma_\varepsilon} \right) \right\} \\ &\equiv \Lambda(s' \mid \mu, \beta). \end{aligned}$$

The expectation of the above future payoff over s' , given the current period state and choice variables, becomes

$$\mathbb{E} \left[\max_{d', e'} \{ (\beta - 1)u(d', e', s' | \mu) + v(d', e', s') + \beta \varepsilon' \} \mid d, e, s \right] = \int \Lambda(s' | \mu, \beta) f(s' | d, e, s) ds'.$$

Thus, the identification criteria in **Equation (8)** simplify to a function of the parameters (μ, δ, β) , where

$$u(d, e, s | \mu) - v(d, e, s) + \delta \int \Lambda(s' | \mu, \beta) f(s' | d, e, s) ds' = \Pi(d, e, s | \mu, \delta, \beta).$$

The true parameter vector (μ, δ, β) solves the above equation $\Pi(d, e, s | \mu, \delta, \beta) = 0$ for all (d, e, s) .

Thus, for identification, the assumption of a full-rank condition is sufficient.

Assumption 3 (Rank Condition). *Denote the agent's decision and state variables by $x = (d, e, s)$. There exists a subset $X = \{x_j : j = 1, 2, \dots, J\}$ in support of x such that $\left\{ \frac{\partial \Pi(x_j)}{\partial (\mu, \delta, \beta)} : j = 1, 2, \dots, J \right\}$ has a rank that is greater than or equal to the number of parameters.*

The assumption rules out the case in which different values of parameters yield identical observations in the model. Mathematically, the assumption holds if no parameters are linearly dependent. The nonlinear nature of the model (the agent's effort enters the mean-variance utility nonlinearly and the disutility function quadratically) readily satisfies the rank condition assumption. The sufficient conditions for Assumption 3 are formally stated in Appendix D.

Theorem 1. *Suppose that the agent's true time preference follows the quasi-hyperbolic discounting model. Under Assumptions 1-3, the standard discount factor δ , present-bias factor β , and parameters of the utility function μ are parametrically identified.*

Proof. See Appendix E. The proof uses the local identification approach to find the unique solution to $\Pi(d, e, s | \mu, \delta, \beta) = 0$.

5.4. Intuitive Discussion of Identification

In addition to the formal identification arguments described above, we discuss model identification in our empirical context. First, we provide intuition regarding the identification of static utility. Then, we discuss identification regarding the discount factor(s).

A key challenge to identifying unobserved effort and utility parameters arises from limited variation in the agent's compensation contract. There exists some variation in the compensation contract—specifically in the quarterly and annual bonus amounts across years. Thus, variations in

performance across the different compensation regimes enable identification. In addition, the relation between an agent’s sales performance and his or her state variables help identification. The agent likely exerts more effort when close to quota than when far from quota. Thus, systematic differences in sales performance at different DTQs identify effort and, thus, facilitate identification of the disutility of effort (Misra & Nair, 2011; Chung et al., 2014). Suppose that there are two agents with the same states, but one has higher performance than the other. Then, we can infer that the agent with higher performance has lower disutility of effort. Similarly, suppose that there are two agents, both of whom have no chance of meeting quota (DTQ is very low), but one has higher performance than the other. Then, we can infer that the agent with higher performance has higher baseline ability. The extent to which an agent over- or underperforms on quota identifies the risk-aversion parameter. A risk-averse agent would constantly over- or underachieve in bonus periods, whereas a risk-neutral agent would just meet quota. The variation in sales in the same states within an agent identifies the distribution of the performance shocks. The variation in sales with variation in performance shifters identifies the performance response parameters.

As described in Section 4.1, the parametric functional form on the agent’s payoff (specifically, the mean-variance utility function) provides location and scale normalization to facilitate identification. The mean-variance utility specification implicitly presumes that the constant term of utility is zero, and the parameter associated with the mean of wealth is unity. Thus, the mean (ρ_i) and variance (σ_ϵ) of the outside option is identified under this specification. Intuitively, if, given a level of income, salespeople are frequently leaving the firm, we can infer that the value of the outside option is high. Relatedly, the observed attrition behavior at different levels of income identifies the variance of the outside option. For example, if salespeople’s attrition behavior does not change much with changes in income, we can infer high variance in the value of the outside option. Naturally, the variation in reservation shifters identifies the corresponding parameters.

As explained in Section 5.2, an agent’s DTQ in non-bonus periods acts as an exclusion restriction to identify discount factor(s). Suppose that there are two agents with the same characteristics who display the same behavior (and, thus, performance) at the end of the year (final period of a compensation cycle). However, suppose that, in non-bonus periods, one agent performs better than the other, even though both are in the same state (DTQ). We can infer that the agent with high performance in non-bonus periods has a higher discount factor (or a lower discount rate). The hyperbolic discounting model, under the functional form specification of utility, is identified if there

exist more than two periods with exclusion restrictions. The performance of an exponential discounter would be more consistent and smoother throughout the year compared to that of a hyperbolic discounter.

6. Estimation

The estimation procedure follows the full-solution method (Rust, 1987) using maximum likelihood, rather than the conditional choice probability approach (Hotz & Miller, 1993; Bajari et al., 2007), since the two-step estimation procedure can generate biases if the state variables in the policy function are correlated with the first-stage errors. In addition, the maximum likelihood approach has the minimum variance achievable by a consistent and asymptotically normally distributed estimator.

6.1. Individual Likelihood

Given the value function in **Equation (4)** and the empirical specification of the per-period utility function in **Equation (3)**, one can obtain the expected value function through the inner loop in the conventional nested fixed-point algorithm (NFXP) such that

$$\text{EV}(d, e, s) \equiv E_{\varepsilon, \xi}[V(s', \varepsilon_{d'} \mid d, e, s)] = \int_{\xi'} \sigma_{\varepsilon} \log \left\{ \sum_{d \in \{0,1\}} \exp \left[\frac{\max_{e \in C_e} \{u(d', e', s') + \phi(1)\text{EV}(d', e', s')\}}{\sigma_{\varepsilon}} \right] \right\} d\xi'.$$

Then, the choice probability of stay-or-leave, $\pi_{d \in \{0,1\}}$, conditional on the agent's state, is obtained by solving the agent's dynamic optimization problem

$$\pi_{dit} = \Pr(d_{it} \mid s_{it}) = \frac{\exp \left(\left(\max_{e \in C_e} \{u(d, e, s) + \phi(1)\text{EV}(d, e, s)\} \right) / \sigma_{\varepsilon} \right)}{\sum_{d \in \{0,1\}} \exp \left(\left(\max_{e \in C_e} \{u(d, e, s) + \phi(1)\text{EV}(d, e, s)\} \right) / \sigma_{\varepsilon} \right)}. \quad (9)$$

In the process, the optimal level of effort e_{it} , given agent i 's state in period t , is inferred by the level at which the expected value function is maximized. The attained effort enters the performance response function in **Equation (2)**.

By combining **Equations (2)** and **(9)**, one can compute the likelihood of the agent's observations. Given the history (data) of an agent with observations over T -periods, the agent's likelihood is

$$L_i(\Omega_i; \mathbf{q}_i, \mathbf{d}_i, \mathbf{s}_i) = \prod_{t=1}^T (\phi_{\xi, i}(\ln(\hat{q}_{it}) - \ln(q_{it})) \cdot \pi_{1it})^{\hat{d}_{it}} \cdot \pi_{0it}^{(1-\hat{d}_{it})},$$

where the vector $\Omega_i = \{\delta_i, \beta_i, \gamma_i, \theta_i, \rho_i, \sigma_\varepsilon, \alpha_i, \sigma_\xi\}$ is the set of parameters of time preference, and the utility and performance response functions; $\hat{d}_{it}$ denotes the observed stay-or-leave decision; $\hat{q}_{it}$ is the observed performance; and $\phi_{\xi,i}$ denotes the probability density function of a normal distribution with mean zero and variance σ_ξ^2 .

6.2. Unobserved Heterogeneity

Discrete segments accommodate *unobserved* heterogeneity (Kamakura & Russell, 1989). Assume that salesperson i belongs to one of K segments $k \in \{1, \dots, K\}$, with relative probabilities

$$m_k = \frac{\exp(\lambda_k)}{\sum_{k'} \exp(\lambda_{k'})}.$$

Let $L_{ikt} = L(\Omega_k | k; q_{it}, d_{it}, s_{it})$ be the likelihood of parameters for individual i at time t , conditional on unobservable segment k , given the agent's data. Then, the likelihood of the segment-level parameters upon observing an individual's history is

$$L_k(\Omega_k; \mathbf{q}_i, \mathbf{d}_i, \mathbf{s}_i) = m_k \left(\prod_{t=1}^T L_{ikt} \right).$$

By summing over all of the unobserved states $k \in \{1, \dots, K\}$, the overall likelihood of individual i becomes:

$$L(\Omega; \mathbf{q}_i, \mathbf{d}_i, \mathbf{s}_i) = \sum_{k=1}^K L_k(\Omega_k; \mathbf{q}_i, \mathbf{d}_i, \mathbf{s}_i),$$

where $\Omega = \{\Omega_1, \dots, \Omega_K\}$ contains the segment-level parameters. Hence, the log-likelihood over the N sample of individuals becomes:

$$\sum_{i=1}^N \log \left(L(\Omega; \mathbf{q}_i, \mathbf{d}_i, \mathbf{s}_i) \right) = \sum_{i=1}^N \log \left(\sum_{k=1}^K m_k \left(\prod_{t=1}^T L_{ikt} \right) \right).$$

7. Results

This section presents the results in the following order. First, we show the results of the exponential and quasi-hyperbolic discounting models and discuss their implications. Then, we show how changes in the compensation plan have led to sales force selection across heterogeneous agents. Next, we show the results of counterfactual simulations that address the substantive questions of this study—how sales management instruments (compensation, recruiting/termination, and training policies) affect the performance and selection of salespeople. Finally, we compare simulated

performance and attrition with actual outcomes in the post-data-analysis period—accompanied by *real* changes in the firm’s sales management instruments—to validate the accuracy of our model.

7.1. Parameter Estimates

Table 5 shows the parameter estimates of the exponential and quasi-hyperbolic discounting models. Based on the Bayesian information criterion (BIC), the three-segment model shows the best fit.¹⁹

Regarding time preference, in the exponential discounting model, the discount factor (δ) is 0.895, 0.975, and 0.983, respectively, for segments 1, 2, and 3. The range of the standard discount factor is consistent with the behavioral and empirical studies on time preference (Frederick et al., 2002; Yao et al., 2012; Chung et al., 2014). In the quasi-hyperbolic discounting model, the standard discount factor (δ) is 0.996, 0.976, and 0.994, respectively, and the present-bias factor (β) is 0.477, 0.999, and 0.980, respectively, for segments 1, 2, and 3. **Figure 4** depicts a graphical illustration of time preference (and, thus, the amount of discounting towards the future) by segment.²⁰ The solid lines represent exponential and the dotted lines represent quasi-hyperbolic discounting. Segment 1 shows myopic and present-biased behavior, while segments 2 and 3 show forward-looking and time-consistent discounting behavior.

The quasi-hyperbolic discounting model is a more general model than the exponential discounting model. Furthermore, the BIC values of the two models indicate that the quasi-hyperbolic discounting model fits the data better, implying that some sales agents are present-biased in their time preferences. Hence, we use the results of the quasi-hyperbolic discounting model for inference. For the structural parameters of the utility function, the disutility parameter ranges from 0.556 to 23.911. The disutility parameter is small for segment 3 (hereafter referred to as the high type), representing the agents’ ease and flexibility in exerting effort. Conversely, the estimate is large for segments 1 and 2 (hereafter referred to as the low and moderate types). Hence, the following pattern appears in terms of segmentation. Segments 2 and 3 exhibit forward-looking behavior, which is expected of moderate-to-high-performing agents who seek the end-of-year bonus and overachievement commission. The low-type agents, on the other hand, show myopic behavior.

¹⁹ The BIC values for one- and two-segment quasi-hyperbolic discounting models are 11,793.79 and 11,169.31, respectively.

²⁰ Although depicted in a single plot for visual illustration, the respective segments in the exponential and quasi-hyperbolic models are not directly comparable, as the segment members—despite being similar—are not fully identical because they are from different model specifications.

The reservation value is low for segments 1 and 2, reflecting limited outside opportunities for these types of agents, and high for segment 3, implying the various potential opportunities outside the firm. For reservation value shifters, education and tenure are statistically insignificant, potentially reflecting the nature of personal selling, in which interpersonal and relational skills are more important than such observable characteristics.²¹

Regarding the parameters of the performance response function, tenure (both the mid-level and the senior dummy variables) has a positive effect on performance. In addition, training improves performance; however, the interaction effect of mid-level and senior dummies with training is negative and significant. Hence, training does not benefit the senior salespeople as much as it does the junior salespeople.

7.2. Selection

The change in the firm’s compensation plan likely has led to the selection of its sales force. **Table 6** shows the share of the three segments and their descriptive characteristics across the full three years of the data. Segment 1, the myopic low type, has the smallest share, 13.29%; segments 2 and 3 represent bigger shares of 64.24% and 22.47%, respectively. Consistent with the parameter estimates in **Table 5**, segment 3 achieves the highest performance, with the largest portion meeting the annual quota. The average base salary is high for this segment, reflecting their tenure. Segment 2, the forward-looking moderate type, falls short on performance compared to segment 3 and has lower tenure. The myopic low type, segment 1, falls short in every performance dimension and exhibits a stark difference in annual variable pay compared to the other segments.

In terms of selection induced over time, **Table 7** shows the segment portfolio (in percentages) of the total sales force by each year-end—reflecting how the portfolio of salespeople has changed over time. The share of segment 2, the forward-looking moderate type, has constantly increased. In contrast, the shares of segments 1 and 3 have decreased over the years. This reflects the frequent quitting within these segments, likely due to insufficient compensation from lack of productivity for segment 1 and good outside opportunities for segment 3.

7.3. Counterfactual Simulations

This section shows the results of several counterfactual simulations that address the key substantive question of this study: how can a firm manage, motivate, and sustain a healthy sales

²¹ In addition, tenure is within the focal firm, which may lead to underestimating its effect on the outside option (compared to industry tenure).

force using the sales management instruments outlined in **Figure 1**? The counterfactuals evaluate agents’ performance and selection according to changes in: (i) compensation structure, (ii) training hours, and (iii) recruiting and termination policies.

The counterfactuals suppose that the firm is undertaking its policy design at the beginning of 2018 (i.e., following the data observation period) with its remaining portfolio of salespeople ($N=400$). The basis for the changes is the 2017 policy. For each new regime, we simulate 200 paths per each individual-segment pair, using the parameter estimates of the quasi-hyperbolic discounting model. Then, we allocate individuals into each segment, based on segment probabilities. Finally, we aggregate performance and selection.

7.3.1. Alternative Compensation Structures

The challenge in designing a compensation plan is to determine the optimal ratio of fixed and variable pay. The theory predicts that when a firm increases the portion of fixed pay, employee attrition will likely decrease. But how would heterogeneous salespeople react differently to the change in terms of both performance and selection? Hence, the first counterfactual exercise examines a change in fixed versus variable pay, while keeping other components constant. **Table 8** depicts the performance, attrition, and compensation amount outcomes of the counterfactual simulations under the new regimes.

First, we increase the base salary by 5, 10, and 15% and keep everything else constant. As anticipated, the attrition rate decreases across all segments. However, a notable aspect is that sales productivity also decreases. This is driven mainly by the retention effect: being granted higher rent, the low-performing agents, who otherwise would have left the firm, are now more likely to stay with the firm. The retention effect is greater for the low types, reflected by the more pronounced decrease in productivity. Next, we increase the bonus amount by 5, 10, and 15% and keep everything else constant. Again, employee attrition decreases; however, compared to the case of the increase in base salary, the reduction is smaller. Moreover, the effect on productivity is positive, especially with the high-type agents, as an increase in the bonus amount helps motivate these agents to a greater extent.

The experiment demonstrates the trade-off between adjusting fixed versus variable pay on employee performance and selection. While increasing the fixed salary could serve as a simple remedy to reduce employee attrition, it could, on average, hurt the overall performance of the sales force. In contrast, an increase in variable pay does not harm performance but has a smaller effect on

employee attrition. The effect of policy changes applies heterogeneously across segments, which affects the resulting portfolio of the remaining salespeople.

7.3.2. Sales Training

The next counterfactual simulation involves changes in sales training. As discussed in Section 7.1, sales training positively affects productivity, but mostly for junior salespeople. The increase in performance affects not only current-period utility but also future outcomes, which bring about changes in the dynamic optimization of effort over time. To evaluate the role of sales training in performance and selection, we provide 6, 12, and 24 hours of sales training in January for all agents. **Table 9** shows the results.

As anticipated, sales training leads to increased performance across all segments. In addition, the employee attrition rate decreases across all segments. The general trend of providing sales training is similar to that of increasing the bonus amount—training helps agents obtain better performance, which, in turn, raises the probability of attaining the bonus (evidenced by the increase in compensation amount).

Although the counterfactual results shown in the previous two subsections provide a practical tool for evaluating alternative compensation and training policies, in practice, managers often face a limited budget for implementing a new sales management policy. Therefore, the objective becomes finding a policy that offers the best possible outcome, subject to the budget constraint. To conduct such cost-benefit comparison requires normalizing the cost factor, as each policy entails a different cost to implement (in both compensation amount and training investment). Hence, the following counterfactual analysis evaluates a scenario in which the allocated budget is held fixed at \$2,400 per salesperson across all policies. This amount is about 10% of total annual compensation, including fixed and variable pay.²²

Table 10 depicts the relative effectiveness of corresponding policies: increasing salary by 12.60%, increasing bonus by 13.75%, and providing 19.50 hours of training, which all satisfy the given annual budget of \$2,400 per salesperson. Regarding sales training, both the fixed investment of providing training sessions at \$37/hour and an increase in compensation from better performance are included as costs. Similar to the findings from the above counterfactual exercises, a fixed salary proves effective for employee retention, whereas variable pay and sales training are more effective for performance growth. Regarding the latter two instruments, variable pay is used to boost

²² **Table 2** provides the data to derive this figure: $\$2,400 \approx (\$1,513.58 \times 12 + \$5,611.07) \times 10\%$

performance of the high-type salespeople (by providing greater upside potential), while sales training supports low and moderate types (through an increase in base productivity). Hence, the analysis demonstrates how the model of agent’s behavior is used to conduct a cost-benefit analysis and, thus, support firms in deciding the sales management policy that can best suit their desired outcome under a constrained budget.

Lastly, we elaborate on the risk of a simple cost-wise comparison across different sales management instruments. Without the model, a manager may have conceived a 19.5-hour sales training to be cost-wise equivalent to a 4% annual salary increase.²³ However, the analysis reveals that, due to the increase in compensation amount, the figures are closer to a 12.60% salary increase. The discrepancy arises due to changes in salespeople’s behavior: training increases salespeople’s productivity, which improves performance and, thus, requires more compensation. This example illustrates the limitation of a simple cost-wise comparison based on accounting figures and highlights the value of a structural model that captures the causal behavior change under a counterfactual scenario.

7.3.3. Recruiting and Termination Policy

A firm can induce selection of its sales force through its recruiting and termination policies. We consider two cases: (i) changes in the firm’s recruiting policy; and (ii) changes in its termination policy.

The recruiting policy relates to the type of salespeople that the firm should target during its hiring process. First, suppose that the firm can observe the agents’ latent types. Should the firm focus on targeting the high types, who are more likely to be skilled but require greater compensation to keep? Alternatively, should the firm target the moderate or low types at lower costs? To evaluate the outcomes of the recruiting policy, we simulate the firm to hire 50 of each type (low, moderate, and high) and compare the differences in performance and attrition over a five-year horizon. **Figure 5a** depicts the annual performance (solid lines) and cumulative attrition rate (columns) by segment.²⁴ As anticipated, the high types show better performance than the low types. However, the high types’ high performance comes at a cost: these agents are more apt to depart the firm due to better outside opportunities, leaving their territories vacant. Hence, in **Figure 5b**, we report the

²³ Providing 19.5 hours of training costs $37 \times 23 = \$721.5$; 4% of annual salary is $\$1,500 \times 12 \times 4.7\% = \720 .

²⁴ The performance figures tend to be lower than the segment characteristics in **Table 6**, as the respective new hires have zero tenure and no training.

attrition-adjusted performance, which accounts for the territory vacancy (treated as zero outcome). Although in the short run, hiring high-type agents leads to greater performance, in the long run, territory vacancy can be detrimental to the firm’s objectives. Therefore, without any changes in effort to retain high-type salespeople, simply recruiting a large number of them can have limited positive effects on the firm’s productivity.²⁵

The above experiment, however, is not directly applicable in practice—as firms cannot observe the agents’ hidden types. Hence, we examine a scenario in which the firm possesses information about the candidates’ tenure, which provides ex-ante understanding of the agents’ experience and, thus, the underlying type. The dilemma is whether to poach rivals’ experienced salespeople, who are more likely to be pre-equipped with sales techniques but require greater compensation to hire, or to target inexperienced rookies, who require a lower base salary and have the potential to be trained from the outset of their career.²⁶ Hence, to answer the question, we simulate the firm to hire 50 salespeople—either experienced (tenure 3-7 years) or rookies (tenure 0-2 years).²⁷ To capture the nurturing opportunity, the firm provides the rookie salespeople with a 24-hour sales training each year (equivalent to the salary difference with experienced salespeople). **Figure 6a** shows the performance and attrition results. As expected, experienced salespeople perform better than rookies in the short run. However, the gap narrows as sales training accumulates for the rookies, and they eventually outperform the experienced. Further, this gain in productivity lowers the rookies’ attrition rate, and, as shown in **Figure 6b**, rookie salespeople exhibit better net productivity (i.e., attrition-adjusted performance) by the fifth year. Therefore, a firm should consider the outcome priority (e.g., short- vs. long-term performance) and the associated efforts (e.g., nurturing vs. retention) when setting its recruitment policy.

On the flip side of recruiting is the firm’s policy for terminating its salespeople. In most nations, including Turkey, firm-initiated employee termination (layoff) is limited due to labor force regulations. Hence, to terminate a salesperson by discretion, the firm must provide a leave package that the employee will agree to. We evaluate the effect of a leave package by providing a lump sum

²⁵ This counterfactual exercise can also be viewed as the impact of not targeting any particular type of salespeople during the recruitment process. We thank the Associate Editor for providing this intuition.

²⁶ Whether to recruit experienced versus rookie salespeople was one of the main concerns for the firm.

²⁷ To compute the segment probabilities conditional on tenure, we apply Bayes’ theorem. For example, the probability of an experienced salesperson to belong in segment 3 is given by $\Pr(\text{Segment } 3 \mid \text{Mid-level}) = \Pr(\text{Mid-level} \mid \text{Segment } 3) \cdot \Pr(\text{Segment } 3) / \Pr(\text{Mid-level}) = 0.284 \cdot 0.226 / 0.244 = 0.264$. An implicit assumption is that the tenure within the focal firm reflects the sales force characteristics at the industry level.

of \$4,500, \$9,000, and \$18,000 (equivalent to a 3-, 6-, and 12-month base salary, respectively) that agents can opt into. **Table 11** shows that the leave package affects the low- and moderate-type agents (segments 1 and 2) more, as the marginal value of the package is higher for these segments. Average firm productivity increases, as agents with less potential tend to be the ones to accept the package and leave. Hence, the termination policy counterfactual reveals that strategically providing leave packages can potentially lead to better outcomes and to the firm’s desired selection of salespeople.

7.4. Field Validation

In the beginning of 2018, the focal firm, based on the results of the above counterfactual analyses, decided to raise its bonus amount by 20% and to offer an additional 12 hours of training (six hours each in the first and second halves of the year). To validate the accuracy of our model, we obtain the performance and attrition records under the new regime (January-June 2018) and compare the actual data and the counterfactual outcomes, simulated based on changes in the sales management instruments.

Figure 7 compares the actual (solid line) and projected (dotted black line) performance outcomes over the six-month period. The model simulation projects the general trend, though it is less cyclical. Overall, the projected performance results fit the actual outcomes well, with a mean absolute percentage error of 0.97% on aggregate and 3.74% in monthly sales. In terms of employee attrition, the model predicts that 11 salespeople would leave the firm during the six-month period. In reality, nine salespeople actually left the firm.

The comparison shows the competence of the model to predict and, thus, evaluate the outcomes under a new policy that includes multiple sales management instruments. We also simulate performance outcomes in the case that the firm had not made any changes (i.e., kept the 2017 plan) to its sales management instruments. The results (dotted gray line) show that the firm’s performance increased by 8.51% as a result of the changes in its sales management instruments.

Organizations should approach with caution when changing their sales management policy, as it can be quite costly. The cost includes not only the direct cost of amending administrative functions, but also opportunity costs and the cost of “getting it wrong.” For example, when an organization initially gives a bonus but takes it away later, salespeople’s performance can be lower than having not given the bonus in the first place because of erosion in intrinsic motivation (Lepper et al., 1973; Chung and Narayandas, 2017). In addition, the organization’s management can lose credibility with

its employees when management policies repeatedly change. Hence, the framework and model of this study provide rigorous yet practical means for organizations to foresee the result of a change in alternative sales management policies.

8. Conclusion

Managing a sales force is an intricate task with multidimensional outcomes. If properly managed, organizations can induce greater performance from their sales force while retaining their top performers. This study develops and estimates a dynamic structural model of comprehensive response to multiple sales management instruments, including compensation, training and recruiting/termination policies. The agent’s model takes into consideration many elements that constitute a realistic working environment—allocation of continuous effort; forward-looking behavior, present bias; effectiveness of sales training; and employee attrition. Substantively, the study provides guidance to firms on (i) evaluating the differential outcomes of various compensation policies; (ii) assessing the selection of different types of employees in relation to changes in recruiting and termination policies; and (iii) addressing the value of sales training.

The following summarizes the study’s results. An increase in fixed salary positively affects employee retention but may decrease aggregate sales because low-type agents, who otherwise would have left the firm, are likely to stay. In contrast, an increase in variable pay enhances sales productivity but has limited effect on employee retention. Because of the focal firm’s selection process over time, high performers steadily left the firm, while mid performers remained. However, if the firm were to focus mainly on recruiting high-performing experienced salespeople, sales would increase in the short term but would likely decrease in the long term due to territory vacancies created by salespeople’s attrition. Hence, firms should focus on retention efforts along with their recruiting efforts of high performers. In addition, providing adequate leave packages can lead to an appropriate selection of salespeople to maintain a healthy sales force. Furthermore, sales training, a novel management instrument that both academics and practitioners have often overlooked, is an effective long-term performance driver that aids salespeople in their early careers to improve their performance and, in turn, their retention. A field validation, comparing post-analysis actual and counterfactual outcomes, verifies the accuracy of the model. The field validation supports the practical applicability of the model in the real world—a model that can predict changes in behavior (and, thus, sales and employee attrition outcomes) under various sales management policies using multiple instruments.

Methodologically, the study introduces a new insight into the marketing and economics literatures by providing a formal proof regarding the identification of discount factors in a hyperbolic discounting model, accompanying continuous and unobserved choices. The key to identification is the aggregation of performance over a specific time horizon when evaluating compensation: an agent’s distance-to-quota for obtaining a bonus payment (in non-bonus periods) serves as an exclusion restriction that affects only future utility and not current utility. The study provides conditions under which both an exponential and a hyperbolic discounting model are identified, and through the empirical application, find evidence of present bias in salespeople’s behavior.

This study has some limitations that open avenues for future research. First, it does not consider multidimensional effort regarding different products (Chung et al., 2020) or customer types (Kim et al., 2019), where agents could exhibit dynamic substitution across products, customers, or both. For example, in the early periods of a quota-evaluation cycle, an agent might focus on high-ticket products that, if sold, could satisfy a large portion of his or her quota. However, as periods pass, an agent might gradually shift to low-ticket and easy-to-sell products. Second, free goods as a sales promotion tool, which is common in the pharmaceutical industry, can induce additional dynamics in a sales agent’s behavior. While free goods reduce the agent’s short-term returns on performance, they can induce greater long-term outcomes by building a stronger relationship with a customer. Relatedly, an agent’s effort, in addition to the immediate short-term effect, can also have a long-term effect on sales through augmented customer relationships. Finally, this study considers time-invariant unobserved heterogeneity; however, time-variant unobserved factors (Arcidiacono and Miller, 2011; Hu and Shum, 2012; Chou et al., 2019) may affect the agent’s effort decision. Although not addressed in this study due to data limitations and model parsimony, the abovementioned topics would provide exciting avenues for future research.

In summary, this study offers a comprehensive, practical, yet rigorous application for understanding the roles of multiple sales management instruments—compensation, training, recruiting and termination—in the selection and performance of salespeople. We believe that the results will guide organizations in their sales management practices to help recruit, compensate, train, and, thus, maintain a healthy sales force to achieve their desired outcomes.

Appendix

A. Proof of Proposition 1. Because the agent chooses effort after observing the state variables (s_{it}), individual heterogeneity (α_i), and the utility shock (ε_{dit}), his or her optimal effort policy e_{it} is a function of s_{it} , α_i , and ε_{dit} (i.e., $e_{it}=e(s_{it}, \alpha_i, \varepsilon_{dit})$). However, ε_{dit} does not affect effort because it is invariant to the effort choice conditional on the stay-or-leave decision, d_{it} . Hence, the performance response function can be represented as

$$\ln(q_{it}) = \alpha_i + e(s_{it}, \alpha_i) + \xi_{it},$$

where q_{it} and s_{it} are observed, but α_i and ξ_{it} are not.

By Assumption 1, when $s_{it} \in S_\alpha$, the value function is a decreasing function of effort. Hence, the optimal effort is zero (i.e., $e(s_{it}, \alpha_i) = 0$ for $s_{it} \in S_\alpha$). Therefore, we have $\ln(q_{it}) = \alpha_i + \xi_{it}$. Independence between ξ_{it} and s_{it} implies

$$E(\ln(q_{it}) \mid s_{it} \in S_\alpha) = \alpha_i + E(\xi_{it} \mid s_{it} \in S_\alpha) = \alpha_i,$$

from which α_i is identified.²⁸

Once α_i is identified (from observations $s_{it} \in S_\alpha$), the performance response (when $s_{it} \notin S_\alpha$) takes the form of a nonparametric regression with a known intercept: $e(s_{it}, \alpha_i)$ is a regression function of $\ln(q_{it}) - \alpha_i$ on s_{it} and α_i . Thus, the optimal effort e_{it} is identified from $E(\ln(q_{it}) - \alpha_i \mid s_{it}, \alpha_i)$ using nonparametric regression methods. The distribution of the residuals is a consistent estimator for the distribution of ξ_{it} .

(Q.E.D.)

B. Proof of Lemma 1.²⁹ Consider an agent with $(s_{it}, \varepsilon_{0it}, \varepsilon_{1it})$. The agent chooses to stay with the firm only if

$$v(1, e, s_{it}) + \varepsilon_{1it} \geq v(0, 0, s_{it}) + \varepsilon_{0it}.$$

Although ε_{0it} and ε_{1it} are unobserved, their joint distribution is assumed to be known up to the scale parameter σ_ε . Thus, the probability of staying with the firm can be written as

²⁸ The individual-specific fixed effect α_i is identified (and, thus, consistently estimated) in the case of a large T (i.e., $T \rightarrow \infty$). In the empirical analysis, the fixed-effects parameters are estimated at the aggregate level (rather than at the individual level) using panel data with fixed T .

²⁹ The proof builds upon that of Lemma 1 in Magnac and Thesmar (2002).

$$\begin{aligned}
\Pr(d_{it} = 1 \mid s_{it}) &= \Pr(v(1, e, s_{it}) - v(0, 0, s_{it}) \geq \varepsilon_{0it} - \varepsilon_{1it}) \\
&= \Pr\left(\frac{v(1, e, s_{it}) - v(0, 0, s_{it})}{\sigma_\varepsilon} \geq \frac{\varepsilon_{0it} - \varepsilon_{1it}}{\sigma_\varepsilon}\right) \\
&= F\left(\frac{v(1, e, s_{it}) - v(0, 0, s_{it})}{\sigma_\varepsilon}\right),
\end{aligned}$$

where $F(\cdot)$ is the cumulative distribution function of $(\varepsilon_{0it} - \varepsilon_{1it})/\sigma_\varepsilon$. As the probability of staying with the firm can be computed from the observable data (d_{it}) , one can obtain the difference in choice-specific value function via

$$v(1, e, s_{it}) - v(0, 0, s_{it}) = \sigma_\varepsilon F^{-1}(\Pr(d_{it} = 1 \mid s_{it})). \quad (\text{Q.E.D.})$$

C. Proof of Proposition 2.³⁰ The value functions at states $(s_1 \in S_1, s_2)$ and $(s_1 \in S_1, \tilde{s}_2)$, which satisfy Assumption 2, can be evaluated such that

$$v(d, e, s_1, s_2) = u(d, e, s_1, s_2) + \delta E[\max_{d', e'} \{v(d', e', s') + \varepsilon'_d\} \mid d, e, s_1, s_2], \quad (\text{A.1})$$

$$v(d, e, s_1, \tilde{s}_2) = u(d, e, s_1, \tilde{s}_2) + \delta E[\max_{d', e'} \{v(d', e', s') + \varepsilon'_d\} \mid d, e, s_1, \tilde{s}_2]. \quad (\text{A.2})$$

Subtracting **Equation (A.2)** from **(A.1)** cancels out the per-period utility:

$$v(d, e, s_1, s_2) - v(d, e, s_1, \tilde{s}_2) = \delta \left\{ E[\max_{d', e'} \{v(d', e', s') + \varepsilon'_d\} \mid d, e, s_1, s_2] - E[\max_{d', e'} \{v(d', e', s') + \varepsilon'_d\} \mid d, e, s_1, \tilde{s}_2] \right\}.$$

From Lemma 1, $v(1, e_{it}, s_{it})$ is identified up to location and scale, and, thus, the difference in value functions (on the left-hand side) is identified up to scale. Using the identified difference in value functions and the law of motion, the difference in expected value functions (on the right-hand side) can be computed up to scale. Because the unidentified scale parameter on both sides cancels out, the discount factor δ is uniquely identified. The per-period utility u is identified from either **Equation (A.1)** or **(A.2)**, given the value function and the discount factor. The scale of the per-period utility function is normalized by the functional form of the mean-variance utility.

(Q.E.D.)

D. Sufficient Conditions for Assumption 3 (Rank Condition). Observe that the derivatives of Π with respect to the parameters are given by

³⁰ The proof uses a similar argument as in the proofs for Corollary 3 and Proposition 4 in Magnac and Thesmar (2002).

$$\begin{aligned}
\frac{\partial \Pi(x | \omega)}{\partial \gamma} &= -\text{Var}(W | x) - \delta(\beta - 1) \int \text{Var}(W' | x') \cdot \exp\left[\frac{(\beta-1)\rho + v'_0}{\beta\sigma_\varepsilon}\right] f(s' | x) ds', \\
\frac{\partial \Pi(x | \omega)}{\partial \theta} &= -e^2 - \delta(\beta - 1) \int (e')^2 \cdot \exp\left[\frac{(\beta-1)\rho + v'_0}{\beta\sigma_\varepsilon}\right] f(s' | x) ds', \\
\frac{\partial \Pi(x | \omega)}{\partial \rho} &= \delta(\beta - 1) \int \exp\left[\frac{(\beta-1)u' + v'_1}{\beta\sigma_\varepsilon}\right] f(s' | x) ds', \\
\frac{\partial \Pi(x | \omega)}{\partial \sigma_\varepsilon} &= \frac{1}{\sigma_\varepsilon} \int \left\{ \Lambda(s' | \mu, \beta) - \frac{[(\beta-1)u' + v'_1] \cdot \exp\left[\frac{(\beta-1)u' + v'_1}{\beta\sigma_\varepsilon}\right] + [(\beta-1)\rho + v'_0] \cdot \exp\left[\frac{(\beta-1)\rho + v'_0}{\beta\sigma_\varepsilon}\right]}{\exp\left[\frac{\Lambda(s' | \mu, \beta)}{\beta\sigma_\varepsilon}\right]} \right\} f(s' | x) ds', \\
\frac{\partial \Pi(x | \omega)}{\partial \delta} &= (\beta - 1) \int \Lambda(s' | \mu, \beta) f(s' | x) ds', \\
\frac{\partial \Pi(x | \omega)}{\partial \beta} &= \frac{1}{\beta} \int \left\{ \Lambda(s' | \mu, \beta) + \frac{(v'_1 - u') \cdot \exp\left[(\beta-1)u' + v'_1\right] + (v'_0 - \rho) \cdot \exp\left[(\beta-1)\rho + v'_0\right]}{\exp\left[\frac{\Lambda(s' | \mu, \beta)}{\beta\sigma_\varepsilon}\right]} \right\} f(s' | x) ds',
\end{aligned}$$

where $u' = u(1, e', s' | \mu)$, $v'_1 = u(1, e', s')$ and $v'_0 = v(0, 0, s')$.

The following conditions are sufficient for Assumption 3 (Rank Condition):

- (i) For any x and $\tilde{x}$ in X , there is a first-order stochastic dominance relationship between $f(s' | x)$ and $f(s' | \tilde{x})$.
- (ii) $\text{Var}(W | x)$ is weakly monotone in s on X but is not a constant function.
- (iii) e^2 and $\text{Var}(W | d, e, s)$ are linearly independent on X .
- (iv) $\Lambda(s | \mu, \beta)$ is strictly increasing in s on X .

The above conditions are readily satisfied in the study's empirical setting: condition (i) holds when x affects only the mean of s' ; conditions (ii) and (iii) are implied by the nonlinear structure of the compensation scheme; and condition (iv) follows from the extreme value distribution assumption.

If $f(s' | x)$ first-order stochastically dominates $f(s' | \tilde{x})$, we have that $\int (a(s')f(s' | x)ds' \geq \int (a(s')f(s' | \tilde{x})ds'$ for any weakly increasing function a . Thus, conditions (i)-(iv) jointly imply that the expected future payoffs vary across x and that the derivatives with respect to σ_ε and β are linearly independent. The derivatives with respect to γ, θ, β are linearly independent by condition (iii). By condition (iv), the derivative with respect to δ is linearly independent of the other derivatives.

E. Proof of Theorem 1. Let $\omega = (\mu, \delta, \beta)$ denote the vector of parameters, and suppose that ω_0 is the true value of these parameters. The vector of true parameters ω_0 is said to be locally identified³¹

³¹ The local identification approach is the standard definition of identification in the economics literature (e.g., Chen et al., 2014). Global identification can be achieved by assuming that the second derivative is globally convex or concave.

if there exists a positive number ζ such that no other parameter value $\tilde{\omega} \neq \omega_0$ satisfies $\Pi(\tilde{\omega}) = 0$ and $\|\tilde{\omega} - \omega_0\| < \zeta$. That is, ω_0 is the unique solution to Π within a certain radius.

The true parameter ω_0 solves $\Pi(x | \omega_0) = 0$ for all x . Although x has infinite support, the information necessary for identification is up to the number of the parameters. Let $\{x_1, x_2, \dots, x_J\}$ be a subset of the support of x satisfying Assumption 3. Denote the equations evaluated at the subset by

$$\pi(\omega) = \begin{pmatrix} \Pi(x_1 | \omega) \\ \Pi(x_2 | \omega) \\ \vdots \\ \Pi(x_J | \omega) \end{pmatrix},$$

and let π' denote the derivative of π with respect to ω , evaluated at ω_0 . A sufficient condition for local identification of α_0 is $\text{rank}(\pi') = \dim(\omega)$ (Chen et al. 2014), which directly follows from Assumption 3.

(Q.E.D.)

References

- Abbring JH, Daljord Ø (2019) Identifying the discount factor in dynamic discrete choice models. *Quantitative Economics*, Forthcoming.
- Ainslie G (1975) Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological bulletin*, 82(4): 463-496.
- Ainslie G, Herrnstein RJ (1981) Preference reversal and delayed reinforcement. *Animal Learning & Behavior*, 9(4): 476-482.
- Ainslie G (1992) Picoeconomics: The strategic interaction of successive motivational states within the person. *Cambridge: Cambridge University Press*.
- Akca S, Otter T (2015) Identifying the discount factor of forward looking consumers based on consumption from inventory. Working paper.
- Arcidiacono P, Miller R (2011) Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79(6): 1823-1867.
- Arfken GB, Weber HJ (1999) Mathematical methods for physicists. *San Diego, CA: Academic Press*.
- Arrow KJ (1971) Essays in the theory of risk bearing. *Chicago: Markham Publishing Company*.
- Atefi Y, Aharne M, Maxham JG, Donovan DT, Carlson BD (2018) Does selective sales force training work? *Journal of Marketing Research*, 55(5): 722-737.
- Attia AM, Honeycutt ED, Leach MP (2005) A three-stage model for assessing and improving sales force training and development. *Journal of Personal Selling & Sales Management*, 25(3): 253-268.
- Bajari P, Benkard CL, Levin J (2007) Estimating dynamic models of imperfect competition. *Econometrica*, 75(5): 1331-1370.
- Basu AK, Lal R, Srinivasan V, Staelin R (1985) Salesforce compensation plans: An agency theoretic perspective. *Marketing science*, 4(4): 267-291.
- Boles JS, Dudley GW, Onyemah V, Rouziès D, Weeks WA (2012) Sales force turnover and retention: A research agenda. *Journal of Personal Selling & Sales Management*, 32(1): 131-140.
- Chen X, Chernozhukov V, Lee S, Newey WK (2014) Local identification of nonparametric and semiparametric models. *Econometrica*, 82(2): 785-809.
- Chevalier J, Goolsbee A (2009) Are durable goods consumers forward-looking? Evidence from college textbooks. *The Quarterly Journal of Economics*, 124(4): 1853-1884.

Ching AT, Erdem T, Keane MP (2014) A simple method to estimate the roles of learning, inventories and category consideration in consumer choice. *Journal of choice modelling*, 13: 60-72.

Ching AT, Osborne M (2019) Identification and estimation of forward-looking behavior: The case of consumer stockpiling. *Marketing Science*, Forthcoming.

Chou C, Derdenger T, Kumar V (2019) Linear estimation of aggregate dynamic discrete demand for durable goods without the curse of dimensionality. *Marketing Science*, 38(5): 888-909.

Christiansen T, Evans KR, Schlacter JL, Wolfe WG (1996) Training differences between services and goods firms: impact on performance, satisfaction, and commitment. *Journal of Professional Services Marketing*, 15(1): 47-70.

Chung DJ, Steenburgh T, Sudhir K (2014) Do bonuses enhance sales productivity? A dynamic structural analysis of bonus-based compensation plans. *Marketing Science*, 33(2): 165-187.

Chung DJ, Narayandas D (2017) Incentives versus reciprocity: Insights from a field experiment. *Journal of Marketing Research*, 54(4): 511-524.

Chung DJ, Narayandas D, Chang D (2020) The effects of quota frequency: Sales performance and product focus. *Management Science*, Forthcoming.

Churchill GA, Ford NM, Walker OC (1993) Sales force management. *Homewood, IL: Richard D. Irwin, Inc.*

Cosguner K, Chan TY, Seetharaman S (2018) Dynamic pricing in a distribution channel in the presence of switching costs. *Management Science*, 64(3): 1212-1229.

Darmon RY (2008) The concept of salesperson replacement value: A sales force turnover management tool. *Journal of Personal Selling & Sales Management*, 28(3): 211-232.

Dhami S (2016) The foundations of behavioral economic analysis. *Oxford: Oxford University Press*.

Dubinsky AJ (1996) Some assumptions about the effectiveness of sales training. *Journal of Personal Selling & Sales Management*, 16(3): 67-76.

El-Ansary AI (1993) Selling and sales management in action: sales force effectiveness research reveals new insights and reward-penalty patterns in sales force training. *Journal of Personal Selling & Sales Management*, 13(2): 83-90.

Fang H, Wang Y (2015) Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions. *International Economic Review*, 56(2): 565-596.

Frederick S, Loewenstein G, O'Donoghue T (2002) Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2):351-401.

- Griffeth RW, Hom PW (2001) Retaining Valued Employees. *Thousand Oaks, CA: Sage*.
- Harris M, Raviv A (1978) Some results on incentive contracts with applications to education and employment, health insurance, and law enforcement. *American Economic Review*, 68(1): 20-30.
- Heide CP (1999) Dartnell's Sales Force Survey. *Chicago, IL: Dartnell Corporation*.
- Hölmstrom B (1979) Moral hazard and observability. *The Bell journal of economics*: 74-91.
- Hölmstrom B, Milgrom P (1987) Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, 55(2): 303-328.
- Honeycutt ED, Ford JB, Rao CP (1995) Sales training: executive's research needs. *Journal of Personal Selling & Sales Management*, 15(4): 67-71.
- Hotz VJ, Miller RA (1993) Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3): 497-529.
- Hu Y, Schennach SM (2008) Instrumental variable treatment of nonclassical measurement error models. *Econometrica*, 76(1): 195-216.
- Hu Y, Shum M (2012) Nonparametric identification of dynamic models with unobserved state variables. *Journal of Econometrics*, 171(1): 32-44.
- Hu Y, Xin Y (2019) Identification and Estimation of Dynamic Structural Models with Unobserved Choices. Working paper.
- Ingram TN, LaForge RW, Avila RA, Schwepker CH, Williams MR (2015) *Sales Management: Analysis and Decision Making*, 9th ed. New York: Routledge.
- Ishihara M, Ching AT (2019) Dynamic demand for new and used durable goods without physical depreciation: The case of japanese video games. *Marketing Science*, 38(3): 392-416.
- Joseph K, Kalwani MU (1998) The role of bonus pay in salesforce compensation plans. *Industrial Marketing Management*, 27(2): 147-159.
- Kamakura WA, Russell G (1989) A probabilistic choice model for market segmentation and elasticity structure. *Journal of Marketing Research*, 26(4): 379-390.
- Kim M, Sudhir K, Uetake K, Canales R (2019) When Salespeople Manage Customer Relationships: Multidimensional Incentives and Private Information. *Journal of Marketing Research*, 56(5): 749-766.
- Kirby KN (1997) Bidding on the future: evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General*, 126(1): 57-71.
- Kumar V, Sunder S, Leone RP (2014) Measuring and managing a salesperson's future value to the firm. *Journal of Marketing Research*, 51(5): 591-608.

- Laibson D (1997) Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2): 443-478.
- Lal R, Srinivasan V (1993) Compensation plans for single-and multi-product salesforces: An application of the Holmstrom-Milgrom model. *Management Science*, 39(7): 777-793.
- Lee RS (2013) Vertical integration and exclusivity in platform and two-sided markets. *American Economic Review*, 103(7): 2960-3000.
- Lepper M, Greene D, Nisbet R (1973) Undermining children's intrinsic interest with extrinsic reward: A test of "overjustification" hypothesis. *Journal of Personality and Social Psychology*, 28(1): 129-137.
- Loewenstein G, Prelec D (1992) Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics*, 107(2): 573-597.
- Lorge S, Smith K (1998) Getting into their heads. *Sales & Marketing Management*. 150(2): 58-64.
- MAGNA (2018) MAGNA Advertising Forecasts. Available at <https://magnaglobal.com/magna-advertising-forecasts-winter-2018-update> (2020/1/21).
- Magnac T, Thesmar D (2002) Identifying dynamic discrete decision processes. *Econometrica*, 70(2): 801-816.
- Martin WS, Collins BH (1991) Interactive video technology in sales training: a case study. *Journal of Personal Selling & Sales Management*, 11(3): 61-66.
- Misra S, Nair HS (2011) A structural model of sales-force compensation dynamics: Estimation and field implementation. *Quantitative Marketing and Economics*, 9(3): 211-257.
- O'Donoghue T, Rabin M (1999) Doing it now or later. *American Economic Review*, 89(1): 103-124.
- Oyer P (1998) Fiscal year ends and nonlinear incentive contracts: The effect on business seasonality. *The Quarterly Journal of Economics*, 113(1): 149-185.
- Oyer P (2000) A theory of sales quotas with limited liability and rent sharing. *Journal of labor Economics*, 18(3): 405-426.
- Phelps ES, Pollak RA (1968). On second-best national saving and game-equilibrium growth. *The Review of Economic Studies*, 35(2): 185-199.
- Polyanin AD, Manzhirov AV (1998) Handbook of integral equations. *CRC press*.
- Rao RC (1990) Compensating heterogeneous salesforces: Some explicit solutions. *Marketing Science*, 9(4): 319-341.

- Richardson R (1999) Measuring the impact of turnover on sales. *The Journal of Personal Selling and Sales Management*: 53-66.
- Roman S, Ruiz S, Luis Munuera J (2002) The effects of sales training on sales force activity. *European Journal of Marketing*, 36(11/12): 1344-1366.
- Rust J (1987) Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, 55(5): 999-1033.
- Rust J (1994) Structural estimation of Markov decision processes. Engle RF, McFadden DL, eds. *Handbook of econometrics Vol.4, Amsterdam: North-Holland*.
- Samuelson PA (1937) A note on measurement of utility. *The Review of Economic Studies*, 4(2): 155-161.
- Schöttner A (2016) Optimal sales force compensation in dynamic settings: commissions vs. Bonuses. *Management Science*, 63(5): 1529-1544.
- Shi H, Sridhar S, Grewal R, Lilien G (2017) Sales Representative Departures and Customer Reassignment Strategies in Business-to-Business Markets. *Journal of Marketing*, 81(2): 25-44.
- Stanton WJ, Buskirk RH (1987) Management of the sales force. *Homewood, IL: Richard D. Irwin, Inc.*
- Steenburgh TJ (2008) Effort or timing: The effect of lump-sum bonuses. *Quantitative Marketing and Economics*, 6(3): 235.
- Sunder S, Kumar V, Goreczny A, Maurer T (2017) Why Do Salespeople Quit? An Empirical Examination of Own and Peer Effects on Salesperson Turnover Behavior. *Journal of Marketing Research*, 54(3): 381-397.
- Thaler RH (1981) Some empirical evidence on dynamic inconsistency. *Economics Letters*, 8(3): 201-207.
- U.S. Department of Labor, Bureau of Labor Statistics (2018). *National Occupational Employment and Wage Estimates*, May.
- Vetterling WT, Teukolsky SA, Press WH, Flannery BP (1992) Numerical recipes: the art of scientific computing (Vol. 2). *Cambridge: Cambridge University Press*.
- Walker OC, Churchill GA, Ford NM (1977) Motivation and performance in industrial selling: present knowledge and needed research. *Journal of Marketing Research*, 14(2):156-68.
- Yao S, Mela CF, Chiang J, Chen Y (2012) Determining consumers' discount rates with field studies. *Journal of Marketing Research*, 49(6): 822-841.

Zoltners AA, Sinha P, Lorimer SE (2012) Breaking the sales force incentive addiction: A balanced approach to sales force effectiveness. *Journal of Personal Selling & Sales Management*, 32(2): 171-186.

Zoltners AA, Sinha P, Lorimer SE (2013) Are you paying enough attention to your sales force? Harvard Business Review (April 12). <https://hbr.org/2013/04/are-you-paying-enough-attention-to>.

Table 1: Sales Force Turnover

	Number of Salespeople		
	2015	2016	2017
Beginning	303	330	367
Joined	58	102	91
Departed	31	65	58
Year-end	330	367	400
Turnover Rate	9.79%	18.65%	15.12%

Table 2: Descriptive Statistics

		All	Decision	
			Stay	Quit
Number of Salespeople		554	400	154
Monthly Base Salary (USD)	Mean	1,513.58	1,549.13	1,421.27
	SD	268.01	271.49	235.65
Annual Variable Pay (USD)	Mean	5,611.07	6,242.38	3,294.32
	SD	3,796.17	3,474.77	4,035.43
Tenure (Years)	Mean	4.08	4.63	2.66
	SD	4.56	4.75	3.67
Sales Training (Hours Per Year)	Mean	3.46	3.79	2.61
	SD	0.66	0.63	0.74
Higher Education (%)	Mean	93.68	93.50	94.16
Annual Performance (%)	Mean	95.12	97.50	86.40
	SD	15.35	13.01	19.64
Meet 100% Quota (%)				
Q1	Mean	29.13	30.73	24.82
Q2	Mean	28.74	30.15	24.94
Q3	Mean	26.93	29.50	18.78
Annual	Mean	28.36	29.71	23.39

Notes. The numbers are approximate for confidentiality.

Table 3: Variable Compensation Payout Ratio (2017)

	Q1	Q2	Q3	Q4
Performance (Sales/Quota)				
0 – 89.99 %	-	-	-	-
90 – 90.99 %	-	-	-	35 %
91 – 91.99 %	-	-	-	41 %
92 – 92.99 %	-	-	-	47 %
93 – 93.99 %	-	-	-	53 %
94 – 94.99 %	-	-	-	59 %
95 – 95.99 %	-	65 %	65 %	65 %
96 – 96.99 %	-	72 %	72 %	72 %
97 – 97.99 %	-	79 %	79 %	79 %
98 – 98.99 %	-	86 %	86 %	86 %
99 – 99.99 %	-	93 %	93 %	93 %
100 %	100 %	100 %	100 %	100 %
Greater than 100 %	100 %	100 %	100 %	101-200 %
Quota Evaluation Period	Jan-Mar	Jan-Jun	Jan-Sep	Jan-Dec
Bonus Amount	\$1,700	\$1,700	\$1,700	\$3,400

Notes. The quarterly variable compensation is determined by multiplying the allocated bonus amount (bottom row) by the payout rate, respective to the performance (sales/quota) over the evaluation period.

Table 4: Relation between Sales Performance and Distance-to-Quota

State	Variable	Monthly Performance											
		Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
%QA < 0.8	Intercept	-	0.71	0.83	0.63	0.49	0.47	0.84	0.67	0.49	0.50	0.43	0.38
		-	(0.04)	(0.11)	(0.06)	(0.09)	(0.07)	(0.08)	(0.10)	(0.11)	(0.13)	(0.15)	(0.15)
	%QA	-	0.14	0.13	0.25	0.29	0.40	-0.16	0.15	0.14	0.22	0.29	0.31
		-	(0.07)	(0.17)	(0.09)	(0.13)	(0.10)	(0.12)	(0.15)	(0.16)	(0.19)	(0.21)	(0.22)
%QA > 0.8	Intercept	-	0.70	0.74	0.55	0.39	0.49	0.26	0.12	0.60	0.21	-0.27	-0.04
		-	(0.06)	(0.09)	(0.07)	(0.07)	(0.09)	(0.08)	(0.09)	(0.09)	(0.09)	(0.09)	(0.11)
	%QA	-	0.31	0.30	0.37	0.56	0.50	0.70	0.98	0.28	0.71	1.24	1.10
		-	(0.06)	(0.09)	(0.07)	(0.07)	(0.08)	(0.07)	(0.09)	(0.09)	(0.08)	(0.09)	(0.10)

Notes. For each column, monthly performance is regressed on the cumulative quota achieved (%QA) by the previous month (respectively for each group of salespeople who are %QA > 0.8 and %QA < 0.8). Standard errors are reported in parentheses. Significance (at the 0.05 level) appears in boldface.

Table 5: Parameter Estimates

	Exponential Discounting			Quasi-Hyperbolic Discounting		
	Seg 1	Seg 2	Seg 3	Seg 1	Seg 2	Seg 3
<i>Time Preference</i>						
Standard Discount Factor	0.895	0.975	0.983	0.996	0.976	0.994
Present-Bias Factor				0.477	0.999	0.980
<i>Utility Function</i>						
Disutility of Effort	21.469	1.510	0.569	23.911	1.383	0.556
Risk-Aversion	0.477	0.843	0.065	0.731	1.130	0.016
Reservation Value						
Baseline	-1.890	-0.691	-0.145	-0.016	-1.015	0.079
Mid-Level (Tenure 3-7 Years)		-0.110			-0.103	
Senior (Tenure > 7 Years)		-0.266			-0.270	
Higher Education		0.107			0.032	
S.D. of Utility Shock		6.300			12.713	
<i>Performance Response Function</i>						
Baseline Productivity		-0.313			-0.328	
Mid-Level (Tenure 3-7 Years)		0.014			0.014	
Senior (Tenure > 7 Years)		0.024			0.027	
Training		0.002			0.002	
Mid-Level $\times$ Training		-0.001			-0.001	
Senior $\times$ Training		-0.002			-0.002	
Higher Education		0.002			0.000	
S.D. of Performance Shock		0.334			0.335	
Log-likelihood		-5405.964			-5385.765	
BIC		11059.254			11047.392	

Notes. Significance (at the 0.05 level) appears in boldface. Standard errors are approximated by inverse of the Hessian of the log-likelihood and are omitted from the table for brevity.

Table 6: Segment Characteristics

	Segment 1	Segment 2	Segment 3
Segment Size (%)	13.29	64.24	22.47
Monthly Base Salary (USD)	1,424.15	1,485.80	1,645.88
Annual Variable Pay (USD)	671.46	4,829.50	10,114.43
Tenure (Years)	2.95	3.99	4.99
Sales Training (Hours Per Year)	2.45	3.35	4.38
Higher Education	0.95	0.93	0.94
Annual Performance (%)	74.42	92.50	112.19
Meet 100% Quota (%)			
Q1	19.21	26.47	41.82
Q2	12.84	26.33	44.42
Q3	4.44	24.85	44.69
Annual	1.75	25.79	48.03

Notes. The numbers are approximate for confidentiality.

Table 7: Selection of Sales Force over Time

	<i>N</i>	Sales Force Portfolio		
		Segment 1	Segment 2	Segment 3
Beginning of 2015	303	10.26%	62.73%	27.01%
End of 2015	330	8.29%	63.34%	28.37%
End of 2016	367	8.71%	66.17%	25.12%
End of 2017	400	6.88%	66.53%	26.59%

Table 8: Counterfactual Simulation: Compensation Structure

Counterfactual Simulation	Total	Within Segment		
		Segment 1	Segment 2	Segment 3
1-1. Increase Salary: 5%				
Average Performance	−0.024	−0.069	−0.010	−0.037
Attrition Rate	−0.891	−1.675	−0.622	−1.197
Compensation Amount	3.948	4.968	4.146	3.210
1-2. Increase Salary: 10%				
Average Performance	−0.029	−0.091	−0.002	−0.069
Attrition Rate	−1.679	−3.331	−1.162	−2.179
Compensation Amount	7.911	9.931	8.317	6.422
1-3. Increase Salary: 15%				
Average Performance	−0.044	−0.137	−0.011	−0.080
Attrition Rate	−2.393	−4.911	−1.681	−2.941
Compensation Amount	11.862	14.885	12.463	9.645
2-1. Increase Bonus: 5%				
Average Performance	2.111	−0.005	1.762	4.378
Attrition Rate	−0.314	−0.054	−0.161	−0.902
Compensation Amount	3.538	0.037	3.460	4.894
2-2. Increase Bonus: 10%				
Average Performance	3.786	−0.010	3.442	7.048
Attrition Rate	−0.587	−0.097	−0.305	−1.681
Compensation Amount	6.950	0.075	7.177	8.867
2-3. Increase Bonus: 15%				
Average Performance	5.909	−0.010	5.840	9.653
Attrition Rate	−0.856	−0.103	−0.489	−2.348
Compensation Amount	11.186	0.113	12.214	12.972

Notes. The change in average performance and attrition rate is in percentage points; the change in compensation amount is in percentages.

Table 9: Counterfactual Simulation: Sales Training

Counterfactual Simulation	Total	Within Segment		
		Segment 1	Segment 2	Segment 3
<i>Increase Sales Training: 6Hrs</i>				
Average Performance	1.831	0.519	2.330	1.188
Attrition Rate	−0.269	−0.167	−0.201	−0.525
Compensation Amount	1.992	0.205	2.901	0.823
<i>Increase Sales Training: 12Hrs</i>				
Average Performance	3.598	1.066	4.499	2.528
Attrition Rate	−0.538	−0.278	−0.420	−1.027
Compensation Amount	4.049	0.434	5.854	1.753
<i>Increase Sales Training: 24Hrs</i>				
Average Performance	6.596	2.193	8.079	4.978
Attrition Rate	−1.101	−0.629	−0.911	−1.925
Compensation Amount	7.675	1.041	10.987	3.461

Notes. The change in average performance and attrition rate is in percentage points; the change in compensation amount is in percentages.

Table 10: Counterfactual Simulation: Cost-Benefit Analysis

Counterfactual Simulation	Total	Within Segment		
		Segment 1	Segment 2	Segment 3
<i>Increase Salary: 12.60%</i>				
Average Performance	−0.035	−0.096	−0.010	−0.072
Attrition Rate	−2.070	−4.143	−1.464	−2.576
<i>Increase Bonus: 13.75%</i>				
Average Performance	5.278	−0.013	5.098	8.965
Attrition Rate	−0.783	−0.108	−0.432	−2.185
<i>Increase Sales Training: 19.50Hrs</i>				
Average Performance	5.613	1.765	6.917	4.179
Attrition Rate	−0.883	−0.479	−0.725	−1.575

Notes. All alternative policies involve a cost of \$2,400 per salesperson (about 10% of total annual compensation, including fixed and variable pay). The changes are in percentage points.

Table 11: Counterfactual Simulation: Termination Policy

Counterfactual Simulation	Total	Within Segment		
		Segment 1	Segment 2	Segment 3
<i>Leave Package: \$4,500 (3-mo Salary)</i>				
Average Performance	0.009	0.000	0.011	0.008
Attrition Rate	0.854	1.129	0.940	0.444
<i>Leave Package: \$9,000 (6-mo Salary)</i>				
Average Performance	0.030	0.004	0.036	0.028
Attrition Rate	1.752	2.105	1.931	1.029
<i>Leave Package: \$18,000 (12-mo Salary)</i>				
Average Performance	0.056	0.030	0.062	0.058
Attrition Rate	3.796	4.097	4.266	2.275

Notes. The changes are in percentage points.

Figure 1: Sales Management Instruments






Figure 2: Firm's Compensation Plan




Notes. The numbers are approximate for confidentiality.

Figure 3: Performance and Cumulative Quota Achieved (%QA)


(a) March




(b) June



(c) September




(d) December



Notes. The y-axis depicts performance (sales/quota) in the bonus-paying months, and the x-axis shows the corresponding agent's cumulative quota achieved (%QA) by the previous month.


Figure 4: Time Discounting by Segment




Notes. The solid lines represent exponential discounting and the dotted lines represent hyperbolic discounting, respective to each segment. The y-axis depicts the rate of discounted future value (as compared to the present value), and the x-axis depicts time horizon (months forward).

Figure 5: Counterfactual Simulation: Recruiting (by Latent Segment)

(a) Performance and Attrition




(b) Attrition Adjusted Performance




Notes. The y-axis shows annual performance (lines) and cumulative attrition rate (columns) by segment. The x-axis depicts time horizon (years).

Figure 6: Counterfactual Simulation: Recruiting (Experienced vs. Rookies)



Notes. The y-axis shows annual performance (lines) and cumulative attrition rate (columns) by type. The x-axis depicts time horizon (years).

Figure 7: Field Validation



Notes. The solid line represents the actual outcome, and the dotted line represents the model projections (simulations). The y-axis depicts the monthly performance across all salespeople (who stay with the firm) and the x-axis depicts months in 2018 (post-analysis period).