

A Salesforce-Driven Model of Consumer Choice

By Bicheng Yang
University of British Columbia

Tat Chan and Raphael Thomadsen
Washington University in St. Louis

Abstract

This paper studies how salespeople affect the choices of which products consumers choose, and from that, how a firm should set optimal commissions as a function of the appeal, substitutability and profit margins of different products. We also examine whether firms are better off promoting products through sales incentives or price discounts. To achieve these goals, we develop a salesforce-driven consumer choice model to study how performance-based commissions incentivize a salesperson's service effort toward heterogeneous, substitutable products carried by a firm. The model treats the selling process as a joint decision by the salesperson and the consumer. It allows the salesperson's efforts to vary across different transactions depending on the unique preferences of each consumer, and incorporates the effects of commissions and other marketing mix elements on the selling outcome in a unified framework. We estimate the model using data from a car dealership. We find that the optimal commissions should be lower for popular items and for items that are closer substitutes with other products. We also find that for the car industry we study, the cost of selling more cars using sales incentives is cheaper than the cost of selling the same number of cars using price discounts.

Keywords: Salesforce management, Incentives, Consumer choice, Differentiated products

1. Introduction

A firm's salesforce is an important marketing tool for both B2C and B2B environments (Kotler and Keller 2008). Data from the Bureau of Labor Statistics show that in 2012 about 14 million people—more than 10% of the workforce—were employed in sales-related occupations. The amount that firms spend on their salesforce is three times as high as spending on advertising (Zoltners et al. 2008). These statistics suggest salespeople's effort can have a significant influence on demand; otherwise, firms would not invest as much in their salesforce.

The traditional consumer-choice literature typically assumes consumers make purchase decisions on their own, with a focus on how a firm's marketing mix (e.g. price promotion) influences their decisions. In reality, consumers may not have complete market information, and the effort from salespeople (e.g., product persuasion, recommendation, and demonstration) can shift consumers' product preferences and thus increase sales. We observe this impact from the sales-management literature, which measures how performance-based commissions help increase sales. The research focus in many of these studies is on designing the optimal compensation structure that solves the classic principal-agent problem. The impact of commissions on consumers' product choices, however, has not been investigated.

This paper studies how, in a market of differentiated products, commissions influence not only the total sales but also *which* products consumers choose. The commissions set for selling various products can incentivize salespeople to allocate different levels of efforts across products and, together with other marketing mix elements that directly influence the consumer preferences, impact the product consumers choose. To understand this deeper, we develop a model that simultaneously incorporates the decisions of salespeople and consumers. The selling process is structurally modeled as a joint decision that involves two sides: although the consumer makes the

final product choice, the salesperson's decision of how much service effort to invest in each product influences the consumer's choice. This *salesforce-driven consumer-choice* model enables us to infer the effects of commissions and other marketing mix elements on demand. We empirically estimate this model from a dataset provided by a car dealership in Japan, and use the results to investigate how commissions should be set differently across the dealer's multiple products. We then study the effectiveness of increasing commissions as a "push" strategy, versus discounting prices as a "pull" strategy, to increase the sales of a product. The findings are important for firms that sell diverse substitutable products.

Our study bridges the consumer-choice literature and the sales-management literature by studying decisions on both sides. The majority of the sales-management literature has focused on salesforce productivity, and the design of the compensation structure based on aggregate sales (see Mantrala et al. 2010 for a thorough overview). From a theoretical analysis, Basu et al. (1986) and Rao (1990) show that a combination of salary plus nonlinear commission with respect to sales is optimal using a common principal-agent framework that links a salesperson's effort to total sales. Empirically, Misra and Nair (2011) estimate a dynamic principal-agent model to estimate the effectiveness of commissions that are combined with quotas and ceilings. They find that the presence of ceilings on commissions demotivates effort. Chung et al. (2014) study a compensation scheme with quotas and bonuses using a dynamic structural model as well, but include latent-class heterogeneity and estimate the discount factors. Daljord et al. (2016) examine the implication of a uniform compensation policy on a firm with a heterogeneous salesforce team, and suggest that while the homogenous plan leads to a much lower profit relative to a fully heterogeneous plan, the ability to choose salespeople mitigates the loss.

A key missing component in all of these papers is consumers, as they only look at the link between compensation and aggregate sales. Kim (2014) incorporates consumers, but focuses on the impact of sales representatives on a single product (i.e., a car radiator). Our paper extends this literature by noting the fact that there are multiple differentiated products that can receive effort, and estimating how consumers ultimately make the choice of which product to buy in the presence of the effort.

A stream of theoretical research has examined the incentives of salespeople when selling multiple products. Farley (1964) assumes that a salesperson has a fixed amount of time to allocate to selling and decides how much time to devote to each product to maximize their earnings. Sales of each product depends on the time the salesperson allocates to that product. Farley concludes that the salesperson's incentives align with the firm's incentives when commissions are set as an equal percentage of the gross margins across all products. Weinberg (1975) extends this result to the case in which a salesperson also controls prices. Srinivasan (1981) notes that the equal commission policy is generally not optimal when a salesperson chooses the amount of selling time to devote to multiple products based on both the commission they get and the disutility of working, and he recommends setting commission rates higher for products with larger elasticities with respect to the service time. Basu et al. (1985) note that because these results rely on the assumption of a deterministic relationship between sales and salesforce effort, a firm could instead more efficiently impose requirements on final sales to maximize its profits. Lal and Srinivasan (1993) extend Farley (1964) by modeling a stochastic selling process. They show that commissions should be set higher for products with higher sales-effort responsiveness, lower marginal costs, and lower uncertainty in the selling process. However, this literature still ignores the role of consumers – and their heterogeneous preferences – in the transaction outcomes.

In contrast to all of the studies above, we model the selling process as a joint decision by a salesperson and a consumer. While the salesperson decides how much service effort to invest toward each product, the consumer makes the final purchase decisions. The service effort is valuable to consumers and helps match their heterogeneous needs with differentiated products. As such, the likelihood of selling a product is influenced not only by the salesperson's productivity and commissions but also by the consumer's preferences for the different products. Our model allows the salesperson's efforts to vary across different consumers depending on the unique product preferences of each consumer. In that sense, our model has some similarity to Copeland and Monnet (2003), who allow for a discrete level of high or low effort for each transaction. Furthermore, our unified framework incorporates the effect of commissions that motivate salespeople's efforts, as well as the effects of other marketing mix variables (e.g. price) that have a direct effect on consumers. This allows us to compare the tradeoff between increasing commissions vs. offering discounts for a product. In that sense, our paper also contributes to the consumer choice literature by studying how salesforce effort – in addition to other marketing mix elements – affects purchase decisions. In fact, we show that ignoring the role of salespeople in a choice model can create an omitted-variable problem that creates a bias in the estimated consumers' price sensitivity.

The estimation results show that not only do consumers have heterogeneous product preferences but also salespeople have heterogeneous sensitivity toward commissions. We use counterfactual analysis to illustrate how product-specific commissions should be set differently based on how attractive each product is for consumers, how substitutable the product is with other products, and the dealer's profit margin for the product. We also compare the effectiveness of using price promotions versus commission incentives to increase sales. We show that when the

commission sensitivity among salespeople (relative to the price sensitivity among consumers) is high, increasing commissions is likely to be more profitable for the dealer. When the commission sensitivity is at a moderate level, however, jointly discounting prices and increasing commissions will be more profitable than using either discounted prices or increased commissions alone.

The rest of the paper is organized as follows: we first develop the model in section 2. Model estimation and identification are also discussed in the section. Section 3 presents the results from the empirical application. Finally, we conclude in section 4.

2. A Salesforce-Driven Model of Consumer Choice

To study how the service effort from salespeople influences the consumer purchase decision, we develop a model that incorporates the decisions of both the customer and the salesperson. The model assumes a single firm sells J differentiated products. For a salesperson s , let $C(s)$ be the collection of all consumers who have been served by the salesperson.

When consumer $i \in C(s)$ visits the store, the consumer purchases at most one unit of one product, and makes the purchase decision that maximizes her indirect utility given the prices and options available in period t . The salesperson chooses which product to recommend, and also how much service effort to invest in the recommended product, to maximize his indirect utility that is a function of the commission of any product that is sold minus the cost of his selling effort. The consumer has an initial preference for each product, but the salesperson can demonstrate the benefits of the recommended product to increase the consumer's purchase utility. To make the selling efficient, the salesperson talks to the consumer and obtains perfect knowledge about the consumer's preferences before deciding which product to recommend.¹ After the salesperson

¹ We discuss this assumption in more detail in Section 3.4 and also discuss potential alternative assumptions.

exerts effort, the consumer makes the final decision of which product she will buy, or she may walk away without making a purchase. Because the salesperson’s service effort directly impacts the consumer’s decision, we call this model a salesforce-driven consumer-choice model, in contrast to traditional choice models that ignore the influence of salespeople.

2.1 Consumer Utility and Purchase Choice

Assume the salesperson exerts service effort e_{sijt} on product j . Consumer i ’s indirect utility is then

$$U_{isjt} = X_{jt}\beta_i + \gamma_i p_{jt} + e_{sijt} + \varepsilon_{ijt} \quad (1)^2$$

where X_{jt} is a matrix of product characteristics and other factors that influence the demand, p_{jt} is the price of the product at time t , and β_i and γ_i represent the individual-specific product preferences and price sensitivity, which are distributed as F_β and F_γ , respectively. The utility of the no-purchase outside option is normalized as $U_{is0t} = \varepsilon_{i0t}$. The stochastic component ε_{ijt} represents the individual heterogeneous product preference, with a joint distribution $\boldsymbol{\varepsilon}_{it} \equiv (\varepsilon_{i0t}, \dots, \varepsilon_{ijt})' \sim F_\varepsilon$. For simplicity, we assume that $\boldsymbol{\varepsilon}_{it}$ are i.i.d. across individuals; however, within-individual preferences are correlated across products. We also assume that $\boldsymbol{\varepsilon}_{it}$ is exogenous to which salesperson is assigned to the consumer. This is a reasonable assumption since $\boldsymbol{\varepsilon}_{it}$ represents the customer’s innate preferences, and because the matching between the salesperson and the consumer is random in our empirical setting.³ We assume that the service effort e_{sijt} that

² Each consumer only appears once in our data, so i also defines the time period t . We include t in the subscripts because the choice set, prices and commissions vary over time. Each i also corresponds to a unique s . Thus, the inclusion or omission of an s subscript does not denote any increase in the flexibility of the variable. We include an s index for e to denote that effort is chosen by the salesperson, while we do not include an s subscript for ε to denote that ε reflects a customer’s innate preference.

³ Note that this does not imply that the effort e_{sijt} is independent from $\boldsymbol{\varepsilon}_{it}$, as the decision of how much effort the salesperson will invest depends on the preferences of the consumer.

directly shifts the purchase utility is non-negative.⁴ We also assume that it does not influence the utility for other products $U_{isj't}$ for $j' \neq j$.⁵ If no service effort is invested in a product, $e_{sijt} = 0$ for that product.

Let y_{sijt} be an indicator function that equals 1 if the consumer chooses to purchase product j , and 0 otherwise. Conditional on the salesperson's service efforts, $e_{s11t}, \dots, e_{sijt}$, for the products sold by the firm, the likelihood that the consumer will purchase product j is

$$Pr(y_{sijt} = 1) = \int_{\epsilon} \int_{\beta} \int_{\gamma} 1 \left\{ X_{jt} \beta_i + \gamma_i p_{jt} + e_{sijt} + \epsilon_{ijt} = \max \left\{ \max \left\{ X_{j't} \beta_i + \gamma_i p_{j't} + e_{sij't} + \epsilon_{isj't}, j' = 1, \dots, J \right\}, \epsilon_{i0t} \right\} \right\} dF_{\gamma} dF_{\beta} dF_{\epsilon} \quad (2)$$

where $1\{\cdot\}$ is an indicator function that equals 1 if the logical expression inside the bracket is true, and 0 otherwise.

If service effort were not in the utility function, equation (2) would be the same as a standard discrete choice model, where consumers are assumed to make their own purchase decisions based purely on their preferences. In our model, however, e_{sijt} is the salesperson's endogenous decision that depends on the commissions of selling various products and the consumer type that is captured by $(\beta_i, \gamma_i, \epsilon_{it})$. Furthermore, service efforts are unobserved in data, so the service effort must be integrated out when evaluating the likelihood function.

2.2 Salesperson Utility and the Choice of Service Efforts

To evaluate the product choice likelihood in equation (2), we need to model the salesperson's choice of service efforts. We assume that the salesperson's decisions are based on maximizing his

⁴ Since e_{sijt} is not observed from data, we normalize the marginal utility of the service effort for the consumer to one. Therefore, one unit of the effort is equivalent to increase the consumer utility by the monetary value of $\$1/\gamma_i$.

⁵ For example, e_{sijt} could represent services that focus on the product, such as trial use, persuasion, and explaining in detail and demonstrating product functions.

transaction utility, which comes from the commission he obtains from the transaction and the effort he exerts. Suppose the salesperson exerts efforts $e_{si1t}, \dots, e_{sijt}$, and the consumer purchases product j . The salesperson's indirect utility is specified as

$$V_{sijt}(e_{si1t}, \dots, e_{sijt}) = \alpha_s \cdot comm_{jt} + \omega_{sijt} - \left(\sum_{k=1}^J e_{sikt}\right)^2 \quad (3)$$

where $comm_{jt}$ is the commission that the salesperson obtains from selling product j , and the coefficient α_s represents the salesperson's responsiveness to commissions. We allow α_s to be heterogeneous across salespeople with a distribution F_α . This coefficient captures the salesperson's marginal utility for commissions relative to the marginal cost of exerting one unit of service effort.⁶ This lack of responsiveness can be interpreted either directly as how much the salesperson values the commission or indirectly as the ability of salesperson. A less capable salesperson, for example, will be represented by a smaller α_s , implying that for a specific set of commissions he will sell less than his peers who have a larger α_s . Likewise, a salesperson who is less motivated by commissions will also sell less and have a smaller α_s .

We restrict effort to be positive and the disutility to be quadratic, consistent with the assumptions of the previous sales literature (e.g. Misra and Nair 2011, Chung et al. 2013).

The stochastic component ω_{sijt} captures the unobserved factors that may influence the salesperson's choice. As examples, a salesperson in a car dealership may be affected by his manager's nudge for selling certain car models, or the salesperson may not be able to let his customer test-drive a car because other customers are already test-driving it. The value of ω_{sijt}

⁶ A general utility function is $V_{sijt}(e_{si1t}, \dots, e_{sijt}) = \tilde{\alpha}_s \cdot comm_{jt} + \tilde{\omega}_{sijt} - \tilde{\rho}_s \cdot \left(\sum_{k=1}^J e_{sikt}\right)^2$, where $\tilde{\alpha}_s$ represents the salesperson's marginal utility for commissions, and $\tilde{\rho}_s$ the marginal (dis)utility for service efforts. This yields the same choices as in our model where the coefficient $\alpha_s = \tilde{\alpha}_s / \tilde{\rho}_s$ and $\omega_s = \tilde{\omega}_s / \tilde{\rho}_s$.

will affect the decision of which product the salesperson will recommend, but not how much effort they put in, as we discuss below.

If the consumer walks away without making a purchase, the salesperson does not receive any commission and his utility will be $V_{si0t}(e_{si1t}, \dots, e_{siJt}) = \omega_{si0t} - (\sum_{k=1}^J e_{sikt})^2$. We denote the joint distribution function of $(\omega_{si1t}, \dots, \omega_{siJt}, \omega_{si0t})$ as F_{ω} .

To solve the optimal service efforts that maximize the salesperson's utility in equation (3), recall that the salesperson is able to correctly identify the customer's utility function from their dialog with the customer. The following two propositions are useful to help reduce the computational burden:

Proposition 1: *When dealing with a consumer of type $(\beta_i, \gamma_i, \epsilon_{it})$, the vector of the optimal service efforts $(e_{si1t}^*, \dots, e_{siJt}^*)$ has at most one positive entry, and the rest are all zero. That is, the salesperson will invest his service effort on one product at most.*

The proof is straightforward. Suppose the salesperson puts in positive service efforts for two products. Since the consumer will choose only one product, one of the service efforts is wasted and thus it reduces the salesperson's utility. The salesperson will know which effort is wasted because they know the consumer's preferences. The salesperson therefore should only focus on the product that the consumer will choose.⁷ Given that the salesperson only puts effort into one product, we can then solve for the optimal effort:

⁷ If the salesperson has uncertainty about consumer preferences, he may put in service efforts for multiple products. The perfect knowledge assumption helps simplify the computation of the optimal service effort in the model and keeps the estimation tractable. We note that most of the bargaining and matching models which, like the model in this paper, also involve two-sided decisions in the economics and marketing literatures make the same assumption. We view such a model as the first step that approximates the complicated transaction process in a reasonable way.

Proposition 2: Suppose that the salesperson recommends product j when dealing with a consumer of type $(\beta_i, \gamma_i, \boldsymbol{\varepsilon}_{it})$. The optimal service effort e_{sijt}^* he will invest is

$$e_{sijt}^* = \max\{X_{1t}\beta_i + \gamma_i p_{1t} + \varepsilon_{i1t}, \dots, X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt}, \varepsilon_{i0t}\} - (X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt})$$

This proof is also straightforward. Suppose $X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt}$ is equal to $\max\{X_{1t}\beta_i + \gamma_i p_{1t} + \varepsilon_{i1t}, \dots, X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt}, \varepsilon_{i0t}\}$, i.e., product j is the preferred choice without any service. The salesperson does not need to put in any effort for j and it will still be purchased. If $X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt}$ is less than $\max\{X_{1t}\beta_i + \gamma_i p_{1t} + \varepsilon_{i1t}, \dots, X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt}, \varepsilon_{i0t}\}$, the effort $e_{sijt} = \max\{X_{1t}\beta_i + \gamma_i p_{1t} + \varepsilon_{i1t}, \dots, X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt}, \varepsilon_{i0t}\} - (X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt})$ will increase the utility to the level that the customer will purchase product j , but expend no additional effort that would lead to further disutility in the salesperson's utility function.

Note that if the service effort required to sell a product is too high (because the consumer's utility of purchasing any product is too low), the salesperson may decide not to recommend any product and let the consumer walk away without purchase. In this case, the optimal service efforts $(e_{s11t}^*, \dots, e_{sijt}^*)$ in Proposition 1 are all zero. The salesperson may also put in zero effort if they decide to recommend the customer's most-preferred product (where $X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt}$ is equal to $\max\{X_{1t}\beta_i + \gamma_i p_{1t} + \varepsilon_{i1t}, \dots, X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt}, \varepsilon_{i0t}\}$), but the product will still be sold (because the consumer's utility of purchasing the product is already high). Proposition 2 treats the latter case as the salesperson recommending product j , and distinguishes from the former case of which the salesperson does not recommend any of the products.

Define \hat{e}_{sijt} as the effort that sets $\hat{e}_{sijt}(\beta_i, \gamma_i, \boldsymbol{\varepsilon}_{it}) = \max\{X_{1t}\beta_i + \gamma_i p_{1t} + \varepsilon_{i1t}, \dots, X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt}, \varepsilon_{i0t}\} - (X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt})$. The two propositions imply that the probability that the salesperson recommends product j can be derived from the utility function in equation (3) as

$$\begin{aligned}
Pr_{sijt}(s \text{ recommends product } j) &= \int_{\omega} \int_{\alpha} 1 \left\{ \alpha_s \cdot comm_{jt} + \omega_{sijt} - \hat{e}_{sijt}^2(\beta_i, \gamma_i, \boldsymbol{\varepsilon}_{it}) = \right. \\
&\quad \left. \max\{\alpha_s \cdot comm_{1t} + \omega_{si1t} - \hat{e}_{si1t}^2(\beta_i, \gamma_i, \boldsymbol{\varepsilon}_{it}), \dots, \alpha_s \cdot comm_{jt} + \omega_{sijt} - \right. \\
&\quad \left. \hat{e}_{sijt}^2(\beta_i, \gamma_i, \boldsymbol{\varepsilon}_{it}), \omega_{si0t}\} \right\} dF_{\alpha} dF_{\omega} \tag{4}
\end{aligned}$$

Under the model's assumptions, if the salesperson's optimal decision is to recommend product j , the product will be purchased by the consumer (even when the effort level is zero).

Therefore, the likelihood that the consumer will purchase product j can be represented by

$$\begin{aligned}
Pr(y_{sijt} = 1) &= \int_{\boldsymbol{\varepsilon}} \int_{\beta} \int_{\gamma} \int_{\omega} \int_{\alpha} 1 \left\{ \alpha_s \cdot comm_{jt} + \omega_{sijt} - \hat{e}_{sijt}^2(\beta_i, \gamma_i, \boldsymbol{\varepsilon}_{it}) = \max\{\alpha_s \cdot \right. \\
&\quad \left. comm_{1t} + \omega_{si1t} - \hat{e}_{si1t}^2(\beta_i, \gamma_i, \boldsymbol{\varepsilon}_{it}), \dots, \alpha_s \cdot comm_{jt} + \omega_{sijt} - \right. \\
&\quad \left. \hat{e}_{sijt}^2(\beta_i, \gamma_i, \boldsymbol{\varepsilon}_{it}), \omega_{si0t}\} \right\} dF_{\alpha} dF_{\omega} dF_{\gamma} dF_{\beta} dF_{\boldsymbol{\varepsilon}} \tag{5}
\end{aligned}$$

Since \hat{e}_{sijt} is a function of the consumer preferences captured by $(\beta_i, \gamma_i, \boldsymbol{\varepsilon}_{it})$, the above equation shows that the probability of selling a product is influenced by (1) the consumer's product preferences, (2) the salesperson's type and, (3) commissions across products. It is important to highlight how our model differs from standard consumer choice models and the recent literature on the salesforce incentives. Compared with the former, which has a similar likelihood function as equation (2), the indicator within the integrals in equation (5) is a function of the salesperson's utility instead of the consumer's utility. This illustrates the centrality of the salesperson's decision that impacts the consumer choice; thus, we call this model the salesforce-driven consumer-choice model. Compared with the literature on the salesforce incentives (e.g. Misra and Nair 2011 and Chung et al. 2014), our model allows the salesperson to allocate different levels of service efforts across differentiated products. Furthermore, it allows the efforts to vary across different consumers in different periods depending on the unique product preferences of each consumer as well as the

level of commissions across products.⁸ Finally, since there is no non-linearity in the compensation plan we study in this paper, our model is static in nature, without dealing with the dynamic decisions as in Misra and Nair (2011) and Chung et al. (2014).⁹

2.3 Model Estimation

Let Θ be the vector of all model parameters that determine the distribution functions F_α , F_ω , F_γ , F_β and F_ϵ . We estimate Θ by maximizing the likelihood function in equation (5). As the likelihood function involves high-dimensional integrals, we make a parametric distribution assumption regarding F_ω , and assume that $\omega_{sijt} \sim \text{Gumbel}(0, \theta_\omega)$ and i.i.d. across s , i , and j (and therefore implicitly, t). Thus, equation (5) can be rewritten as

$$Pr(y_{sijt} = 1 | \Theta) = \int_\epsilon \int_\beta \int_\gamma \int_\alpha \frac{\exp\left(\frac{\alpha_s \cdot comm_{jt} - \hat{\theta}_{sijt}^2(\beta_i \gamma_i, \epsilon_{it})}{\theta_\omega}\right)}{1 + \sum_l \exp\left(\frac{\alpha_s \cdot comm_{lt} - \hat{\theta}_{silt}^2(\beta_i \gamma_i, \epsilon_{it})}{\theta_\omega}\right)} dF_\alpha dF_\gamma dF_\beta dF_\epsilon \quad (6)$$

Next, we use a simulation method to evaluate equation (6). Conditional on a candidate Θ , the simulation procedure is outlined below:

1. For every observation (i, t) in the data, make NS draws of $(\epsilon_{it}^{sim}, \beta_i^{sim}, \gamma_i^{sim})$ from F_ϵ , F_β , F_γ , where $sim = 1, \dots, NS$. Also make NS draws of α_s^{sim} from F_α for every salesperson. The same α_s^{sim} will be used for the salesperson for repeated transactions.

⁸ Our model has some similarities to Copeland and Monnet (2003), who allow for a discrete level of high or low effort for each “transaction” (a worker that sorts checks decides whether effort should be exerted for each check).

⁹ Other incentives such as the future increase in base salary and career movements (e.g. Yang et al. 2017) may influence the salesperson’s decisions in a dynamic way. Because the focus of this paper is on commissions, we abstract away from these further complications.

2. Conditional on the simulated $(\boldsymbol{\varepsilon}_{it}^{sim}, \beta_i^{sim}, \gamma_i^{sim})$, calculate the effort level

$$\hat{e}_{sijt}(\boldsymbol{\varepsilon}_{it}^{sim}, \beta_i^{sim}, \gamma_i^{sim}) = \max\{X_{1t}\beta_i^{sim} + \gamma_i^{sim}p_{1t} + \varepsilon_{ijt}^{sim}, \dots, X_{jt}\beta_i^{sim} + \gamma_i^{sim}p_{jt} + \varepsilon_{ijt}^{sim}, \varepsilon_{i0t}^{sim}\} - X_{jt}\beta_i^{sim} - \gamma_i^{sim}p_{jt} - \varepsilon_{ijt}^{sim}.$$

3. Conditional on $\hat{e}_{sijt}(\boldsymbol{\varepsilon}_{it}^{sim}, \beta_i^{sim}, \gamma_i^{sim})$ and α_s^{sim} , calculate the likelihood $Pr(y_{sijt} = 1 | \boldsymbol{\varepsilon}_{it}^{sim}, \beta_i^{sim}, \gamma_i^{sim}, \alpha_s^{sim})$ in equation (6).

4. Calculate the simulated choice probability $Pr(y_{sijt} = 1 | \boldsymbol{\Theta}) = \frac{1}{NS} \sum_{sim=1}^{NS} Pr(y_{sijt} = 1 | \boldsymbol{\varepsilon}_{it}^{sim}, \beta_i^{sim}, \gamma_i^{sim}, \alpha_s^{sim})$.

We search for the estimate $\hat{\boldsymbol{\Theta}}$ to maximize the following simulated maximum likelihood function:

$$L = \sum_t \sum_s \sum_i \sum_j \log \left(Pr(y_{sijt} = 1 | \boldsymbol{\Theta}) \right) \cdot y_{sijt} \quad (7)$$

2.4 Identification

We only observe the consumers' product choices. The main identification of the commission sensitivity (α_s) among salespeople versus the price sensitivity (γ_i) among consumers comes from an exclusion restriction: a change in commissions causes salespeople to reallocate service efforts among products but does not directly impact consumers' utility. Below, we provide more details regarding the identification of different parameters in the model.

The likelihood function in equation (6) indicates what variation in the data is required to identify the model. First, let commissions be fixed. Note that $\hat{e}_{sijt}(\beta_i, \gamma_i, \boldsymbol{\varepsilon}_{it}) = \max\{X_{1t}\beta_i + \gamma_i p_{1t} + \varepsilon_{i1t}, \dots, X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt}, \varepsilon_{i0t}\} - (X_{jt}\beta_i + \gamma_i p_{jt} + \varepsilon_{ijt})$ is a monotonic function of X_{jt} and p_{jt} , and the probability function $Pr(y_{sijt} = 1)$ is monotonically decreasing in the amount of effort needed to make a sale, $\hat{e}_{sijt}(\beta_i, \gamma_i, \boldsymbol{\varepsilon}_{it})$. Thus, parameters β_i and γ_i are identified by how product j 's market share changes following changes in X_{jt} and p_{jt} . The identification of F_ε comes

from how products are substitutable with each other when their prices and product attributes change, the same argument that has been well-established in the standard consumer choice literature. For example, if a price reduction of product j cannibalizes the sales of product k more than other products, it implies ε_{ijt} is more positively correlated with ε_{ikt} .

Next, let prices and other product attributes be fixed and assume that F_ω is known. Equation (6) shows directly how α_s is identified from how the probability of salesperson s selling product j changes with different commissions. The distribution of α_s , F_α , is identified from the heterogeneity in the responsiveness from commissions across salespeople. Note that we assume that p_{jt} and $comm_{jt}$ are not perfectly correlated; otherwise the model cannot separately identify the price coefficient from the commission coefficient. Whether this assumption is valid or not depends on the commission scheme. If the firm adopts a revenue-based commission, i.e., $comm_{jt} = r \cdot p_{jt}$, where r is a fixed commission rate, then the model cannot be identified.¹⁰

In addition to changes in commissions and prices, other data variations can also help with the model identification. In particular, product entries and exits change the choice sets for consumers and salespeople over time. Changes in the market share of other products following the entries and exits can identify not only F_ε but also the distributions of β_i , γ_i and α_s . As an example, suppose we observe that after an existing product that provides salespeople with a high commission exits, the market share of the products that also provide high commissions increase more than those that provide low commissions even as this pattern is not replicated through shifting prices before the exit. This implies that salespeople are very sensitive to the commission levels.

¹⁰ As we describe later, the firm in our empirical application adopts a commission that is based on the profit margin. The margins vary over time and are not fully aligned with changes in prices since prices and costs to the dealer do not vary too closely across time. Therefore, the price and commission coefficients can be identified in our empirical model.

The above argument about how to identify α_s assumes that F_ω is known. To separately identify α_s from F_ω , however, proves to be difficult because service efforts are not observed from the data. For an illustration, assume that $\omega_{sijt} \sim \text{Gumbel}(0, \theta_\omega)$. Equation (6) shows that (i) a very positive α_s , or (ii) a very small θ_ω , can both predict that the probability of selling product j increases drastically following an increase in $comm_{jt}$. However, the way that α_s and θ_ω are functionally specified in equation (6) implies that the magnitude of the change in the selling probability as commissions vary in the data differs between case (i) and case (ii). Consequently, in theory a very positive α_s can be separately identified from a very small θ_ω . However, in practice accurately estimating these parameters separately turns out to be challenging in our data due to the similarity in the effect that both of these parameters have on choices. Thus, we set $\theta_\omega = 1$.

We use a simulation study to further show that our model can recover the structural parameters. We first draw the prices and commissions of four hypothetical products from a large range over 24 months, and simulate the sales of the products under different sets of parameter values. For simplicity, we assume that each product's utility can be summarized through a unique brand intercept, and normalize the intercept for the fourth product as 0. For this simulation study we also assume that consumer heterogeneity is present only in the ε terms, and that the salesforce is homogeneous but draws ω 's for individual transactions. We further assume that the ε 's are normally distributed with mean zero, and that the covariances are all zero except for the covariance of products 1 and 2 (Σ_{12}) and the covariance of products 3 and 4 (Σ_{34}). Finally, ω_{sijt} is assumed to have a Gumbel distribution with θ_ω fixed to one. We then use the simulated data to estimate the model parameters using the likelihood function in equation (7). Table 1 presents the results for a variety of parameters. The results show that the model is able to recover the true parameters very well, with the price and the commission coefficients being especially close to the true values.

Table 1. Test of model performance

	β_1	β_2	β_3	γ	α	Σ_{12}	Σ_{34}
True	0.00	0.00	0.00	-0.80	10.00	0.50	0.70
Estimated	0.00	0.01	0.00	-0.80	9.87	0.52	0.71
Standard error	(0.01)	(0.01)	(4.97E-03)	(2.99E-03)	(0.06)	(0.01)	(0.01)
True	3.00	2.00	4.00	-0.70	9.00	0.70	0.50
Estimated	3.07	2.08	4.06	-0.70	8.96	0.73	0.47
Standard error	(0.06)	(0.06)	(0.06)	(3.47E-03)	(0.05)	(0.02)	(0.04)
True	3.00	2.00	4.00	-1.00	5.00	0.50	0.70
Estimated	3.16	2.17	4.16	-1.00	4.98	0.51	0.63
Standard error	(0.09)	(0.09)	(0.09)	(4.34E-03)	(0.04)	(0.02)	(0.05)
True	1.00	2.00	3.00	-0.60	7.00	0.80	0.40
Estimated	0.95	1.96	2.96	-0.59	6.92	0.82	0.46
Standard error	(0.03)	(0.03)	(0.03)	(3.93E-03)	(0.04)	(0.01)	(0.02)
True	3.00	0.00	4.00	-0.90	6.00	0.30	0.50
Estimated	3.13	0.10	4.13	-0.90	5.93	0.15	0.42
Standard error	(0.09)	(0.08)	(0.09)	(4.24E-03)	(0.05)	(0.07)	(0.05)

3. Data and Analysis

3.1. Data Description

We estimate the model using a dataset provided by one of the largest regional chains of automobile dealers in Japan. This dealership sells cars produced by one of the largest car manufacturers in the world. It owns over 70 outlets and sells multiple brands produced by the same manufacturer. All the outlets sell the same set of car models. The dataset spans two years. There are altogether 828 salespeople. The mean and standard deviation of salespeople's tenure is 12.9 and 9.7, respectively. 59.8% of the salespeople have a college degree, and all of them are male. The data are well suited for our study because we know not only each salesperson's transactions but also the commissions the salespeople receive from each transaction. Furthermore, for each car we know both the selling price and the cost borne by the dealer, which the dealer calculates based on the price it pays the manufacturer for the car as well as administrative and inventory costs.

In each month, a salesperson receives a guaranteed base salary, which can change over time depending on his tenure and accumulated past work performance, plus commissions from the transactions he completes that month. The commission has a fixed and a profit-based component, with commissions calculated as $q_t + r_{jt} \cdot (p_{jt} - c_{jt})$, where q_t is a commission paid each time a car is sold and is fixed for all car models, r_{jt} the commission rate specific for car model j , p_{jt} the selling price, and c_{jt} the cost for the dealer (thus the commission for selling the model depends on the margin for the dealer). The dealer determines the selling prices and commission rates of each car model. In the data, commission rates differ across car models but remain constant over time except for two car models. However, prices and costs fluctuate across months and do not move in lock-step together, and have an average correlation of 0.52. Commissions make up around half of a salesperson's income, so the salespeople have a strong economic incentive to sell cars. We do not model the participation constraint of salespeople in this paper because we are not allowed to use salary data for this project. However, we were told by the company that very few people left the company during this time as Japan was in recession during that period.

We focus on the three most popular classes of cars purchased by individual consumers: multi-purpose vehicles (MPV), sedans, and mini sport-utility vehicles (mini SUV), which make up 25,731 transactions over the two years. We focus on the top 16 models with sales greater than 300 units over the 24 months, where each model has a distinct combination of brand, engine and transmission.¹¹ Table 2 provides summary statistics broken down by car model. We report all monetary values in U.S. dollars using an exchange rate of 118 Japanese Yen per USD. For each

¹¹ Within each model, sub-models differ by exterior and interior colors and other features. We are told the above three attributes are most important in consumers' decision-making.

car model, we present the mean and standard deviation of its price, cost and commissions on a per unit basis, as well as the average monthly total unit sales and profit.

Table 2. Summary statistics

Class	MPV						Mini-SUV	
Car model	1	2	3	4	5	6	7	8
Price (US\$)	12491 (92)	15223 (234)	18172 (369)	15308 (72)	16816 (455)	14440 (797)	11090 (282)	20580 (256)
Cost (US\$)	12344 (45)	14233 (193)	16850 (218)	14844 (22)	14950 (377)	12853 (439)	9868 (137)	19405 (206)
Commission per car (US\$)	191 (10)	318 (12)	434 (41)	262 (13)	393 (52)	360 (51)	353 (36)	311 (31)
Total unit sales per month	25 (11)	29 (10)	335 (154)	171 (54)	34 (29)	146 (53)	43 (24)	65 (31)
Total profit per month (US\$)	3954 (2830)	28656 (9808)	457322 (252393)	76841 (16830)	71132 (73887)	245418 (136119)	56793 (37324)	79441 (49729)

Class	Sedan							
Car model	9	10	11	12	13	14	15	16
Price (US\$)	11740 (66)	12868 (162)	16344 (217)	12829 (215)	14447 (164)	14540 (213)	10859 (296)	9249 (211)
Cost (US\$)	11195 (22)	12125 (96)	14802 (254)	11253 (64)	13442 (134)	12868 (101)	9297 (128)	8035 (76)
Commission per car (US\$)	306 (19)	355 (18)	478 (72)	484 (44)	421 (21)	504 (32)	497 (66)	425 (62)
Total unit sales per month	35 (11)	111 (26)	45 (23)	117 (56)	43 (13)	245 (99)	40 (15)	77 (40)
Total profit per month (US\$)	18610 (4817)	82626 (21502)	73516 (48759)	191118 (113655)	43594 (13774)	416357 (195736)	63121 (27718)	98745 (62940)

*Numbers in parentheses are standard deviations that are calculated from the variation across months.

The standard deviations reported in the table indicate that for each car model there are significant fluctuations in prices and costs across months. Similarly, the range of prices and costs across cars within each car class is also large. Thus the profits of selling a car vary across car models and months. Furthermore, the total unit sales averaged over months range from merely 25 (model 1) to 335 (model 3), reflecting a large variation in terms of the appeal to consumers across

car models. On average, the dealer pays salespeople about 30% of the margin as commission. The commissions also fluctuate across car models and months.

We use a classic multinomial logit choice model to investigate the impacts of prices and commissions on the consumers' choices, controlling for car-model fix effects and year-month fix effects. For comparison purposes, we also estimate a traditional consumer choice model without including commissions. The results for both models are reported in Table 3. Price coefficients are significant regardless of whether commissions are not included (column 1) or are included (column 2) in the estimation. The commission coefficient in the second model is significantly positive, demonstrating that it affects the consumers' choices. Furthermore, the price coefficient in column 2 is more negative than in column 1. To understand this pattern, note that commissions are set as $q_t + r_{jt} \cdot (p_{jt} - c_{jt})$, so the commission for a car increases as its price goes up. In the data, the average correlation between commissions and prices is 0.82 across car models. Thus, when prices increase for a car model salespeople's incentive to invest in service efforts also increase for that car model due to the higher commissions, offsetting much of the decrease in sales due to the higher prices. Ignoring the role of salespeople's incentives in transactions will therefore bias the estimated price sensitivity of consumers towards zero. This result demonstrates the importance of accounting for salespeople's incentives even when the researcher's focus is to understand price sensitivity.

Table 3. Results from a Logit Regression

	Without commission	With commission
Price	-0.118 (0.025)	-0.167 (0.029)
Commission coefficient		0.882 (0.247)
Car model indicators	Included	Included
Time indicators	Included	Included
Loglikelihood	-131724.646	-131718.274
BIC	263805.108	263804.634

3.2. Details of the Full Estimation Model

While Section 3.1 demonstrates the importance of controlling for commissions, merely putting commissions in the choice model offers no guidance on why commissions affect sales even though they do not benefit consumers. Without modeling the role of salespeople’s effort in the transaction process, we will not be able to study how to set optimal commissions (and prices) under new market conditions, which is one of our key research objectives. Thus, we estimate the proposed structural model described in Section 2.

For X_{jt} in equation (1), we include car model indicators, as well as year-month indicators (23 indicators for the 24 months) to control for seasonality. We assume that the price coefficient is distributed such that $\gamma_i \sim N(\bar{\gamma}, \sigma_\gamma^2)$. For simplicity we assume that there is no heterogeneity in other model parameters, and estimate a multinomial probit model, with $\varepsilon_{it} \sim N(0, \Sigma)$. The diagonal elements of Σ are normalized to 1 and, to simplify the estimation, we assume the correlation of ε_{ijt} and $\varepsilon_{ij't}$ is τ_c for all car models j and j' belonging to the same class c (MPV, sedan, or mini-SUV), and 0 otherwise.¹² We restrict the covariances with each class to be non-negative (i.e., $0 < \tau_c < 1$). Consequently, all MPV car models, for example, are closer substitutes for one another than for sedan or mini-SUV models.

For salespeople, we allow them to be heterogeneous in the commission coefficient α_s . We assume that there are two latent types of salespeople. For each type, k , we estimate the coefficient on commissions as

$$\alpha_k = \exp(\alpha_{k,1} + \alpha_{k,2} * tenure + \alpha_{k,3} * tenure^2 + \alpha_{k,4} * college) \quad (8)^{13}$$

¹² Note that this structure removes IIA concerns from our model.

¹³ The exponential function specification guarantees that the commission coefficient is always positive.

Under this specification, salespeople are differentiated by the latent type and, within each, differentiated by how long they have worked for the dealership (*tenure*, measured by months) and their education level (*college*, which is equal to 1 if the salesperson has a college degree, and 0 otherwise). The heterogeneity in α reflects the differences across salespeople in the marginal utility for commissions and/or the marginal (dis)utility for service efforts.

We also assume that $\omega_{sijt} \sim \text{Gumbel}(0, \theta_\omega)$ and is i.i.d. across s , i , and j . Therefore, the likelihood $\Pr(y_{sijt} = 1)$ is specified as in equation (6). Given that it is difficult to identify both θ_ω and the other parameters of the model, as we discussed in Section 2.4, we fix the value of θ_ω to 1 when estimating the model. We then vary θ_ω and re-estimate the model based on each unique value to test the robustness of the model estimates.¹⁴

In the model, consumers may choose the outside no-purchase option. Those consumers are not observed in the data; therefore, we need a proxy for the market potential. We use an approach based on Albuquerque and Bronnenberg (2012) that assumes that the number of people who consider buying a car in a given year, y , is

$$M_y = \frac{\text{Total Number of Households in the Region}_y}{7} \cdot \frac{\text{Observed Sales}_y}{\text{Regional Sales}_y}. \quad (9)$$

In this formula, “7” is the average number of years between car purchases we obtain from industry reports, and the first ratio represents the total number of potential consumers who are looking to buy a new car in the region in the year. For the second ratio, “*Observed Sales_y*” is the total number of cars sold by the dealership in data, and “*Regional Sales_y*” is the total number of cars sold in the region during the year, which we obtain from the Japanese census data. This ratio

¹⁴ We find that there are minor changes in the likelihood function value (within a magnitude of .02%), and that the model estimates are very similar under different values of θ_ω . This suggests that our results are robust to the assumption of θ_ω . Detailed estimation results are available from the authors upon request.

represents the market share of the dealer in our data. M_y is thus used as a proxy for the number of consumers who visited the dealer in the year.

We further calculate the total number of potential consumers for the dealer in a month-year as $M_{my} = M_y \cdot \frac{\text{Observed Sales in } m}{\text{Observed Annual Sales}}$, where “*Observed Sales in m*” is the average monthly sales in month-of-the-year m , and “*Observed Annual Sales*” is the average annual sales of the dealership. This way, the number of potential consumers for the dealer is assumed to be proportional to the monthly sales averaged across years.¹⁵ Finally, we calculate the number of potential buyers for each salesperson by dividing M_m by the total number of salespeople employed by the dealer. In other words, we assume that all salespeople have equal selling opportunities. Based on interviews with the dealer, we understand that the dealer uses a territorial system: When a consumer walks into a dealer, the salesperson who first sees her will greet her and ask some basic questions including where she lives. The customer will then be assigned to the salesperson who is in charge of the territory where their residence is located. If the salesperson who is in charge of the territory happens to be unavailable, the person who greets the consumer will take care of her. For fairness, the dealership assigns the territories with the goal of ensuring that each salesperson will serve the same number of customers; therefore, the assumption of equal selling opportunities seems reasonable. Because the territorial system is based on the consumer’s residence only and not on other consumer attributes, we assume that there is no strategic matching between a salesperson and a consumer.

¹⁵ Note that, since we include month indicators in the estimation model, our results are less sensitive to the way that we calculate the potential consumers. Suppose we over-estimate the number of potential consumers in a month. The estimated fixed effect of that month will adjust downward to reflect that the over-estimated potential consumers will walk away without purchasing from the dealer.

Finally, we are concerned with the issue of price endogeneity in the model estimation. The price of a car may be correlated with the stochastic components (ε, ω) in the model. One issue one may worry about is that a consumer may negotiate price with the salesperson in the transaction process. During interviews with the dealer, we were told that, in contrast to US practices, consumer-initiated price negotiations are very rare in Japan. Even when price negotiations occur, salespeople cannot decide what price to offer; instead, they have to let the manager decide the final price. Therefore, price endogeneity is potentially less of a concern in this study than in some other studies. However, we further address the issue of endogeneity by adopting a control function approach proposed in Petrin and Train (2010). We estimate the model in two stages. In the first stage, we regress the price of every car model in each month on product-level indicators and a set of price instruments, which include the total number of car models in the same class and in other classes. We observe new car models enter and old car models exit in various months, and prices of existing car models significantly fluctuate following the entry and exits (as shown from the regression). These variables are valid instruments for prices under the assumption that the entries and exits are not correlated to the individual ε 's and ω 's. We have already included month and year and car model fixed effects, so the ε 's and ω 's reflect short-term shocks that are individual specific. Since the entry and exit decisions are decided much earlier (depending on the success of R&D) and do not depend on the demand shocks at the individual level, the assumption that they are not correlated with ε 's and ω 's is reasonable. Changes in the number of car models, on the other hand, are likely to be correlated with prices because the dealer will have incentives to adjust the prices of the other products in response to those changes to maximize its profits.

We obtain a residual ξ_{jt} for each car model in each month from the price regression. In the second stage, when estimating the proposed structural model, ξ_{jt} is included as an additional

covariate in the consumer utility function in equation (1). Petrin and Train (2010) show that this method helps correct the potential endogeneity problem in discrete choice models.¹⁶

In addition to the proposed model, we also estimate three alternative models with simpler specifications. The first (Model 1) does not account for heterogeneity across salespeople in α_s or for potential price endogeneity. The second model (Model 2) uses the control function approach to correct for price endogeneity, but still does not allow for heterogeneity across salespeople. The third model (Model 3) assumes that the heterogeneity across salespeople only comes from the tenure and education, which observed from data (see equation 9), but that there are no latent type differences. Comparing the results across these models helps us understand how each of the above components can increase the model fit, as well as how the estimation results are robust to different model specifications.

3.3. Estimation results

Table 4 reports the estimates from the four model specifications. All of the results show that the mean price coefficients are negative and significant; however, after controlling for the potential price endogeneity, the coefficients are more negative in the latter three models. The estimated heterogeneity (standard deviations) of the price coefficients are not statistically significant in any of the model specifications, although the standard errors are high in our proposed model, demonstrating that it is hard for us to pin down this effect.

The estimated correlation coefficients of the three classes of cars have a direct effect on the own- and cross-price elasticities, which in turn affects the way that the salespeople shift their

¹⁶ One may also be concerned that the wholesale price cost for the dealer, c_{jt} , is also endogenous. However, the wholesale price is the same for every transaction in the same month t . In the model, we include a fixed effect for every car model, as well as fixed effects for years and months, which should alleviate the endogeneity concern.

service efforts in response to changes in commissions. We observe high correlation in preferences in the MPV and mini-SUV segments, but not for the sedan segment. Using the results from the proposed model, we calculate the average price and commission elasticities and report them in Table 5. We find that the own- and cross-price elasticities are the smallest for the sedan models. The own-elasticity for commissions, however, are the highest for this class, suggesting that increasing commissions may be especially effective in increasing sales for this class of vehicles.

Table 4. Estimation results

	Model 1	Model 2	Model 3	Proposed full model	
Consumers					
Price coefficient: mean	-0.124 (0.002)	-0.244 (0.014)	-0.248 (0.004)	-0.253 (0.028)	
Price coefficient: s.d.	0.006 (0.005)	0.006 (0.013)	0.007 (0.400)	0.003 (0.203)	
Correlation:MPV	0.929 (0.054)	0.929 (0.055)	0.766 (0.032)	0.316 (0.078)	
Correlation:Sedan	0.014 (0.285)	0.011 (0.414)	0.015 (2.587)	0.017 (1.061)	
Correlation: Mini-SUV	0.691 (0.062)	0.690 (0.061)	0.665 (0.097)	0.434 (0.170)	
Coefficient for the residual in the price regression		0.120 (0.008)	0.123 (0.042)	0.131 (0.026)	
Salespeople					
Average commission sensitivity	3.060	3.020	3.154	2.578	
				Seg 1	Seg 2
Constant	1.117 (0.114)	1.106 (0.111)	0.843 (1.089)	1.174 (0.040)	-1.203 (1.599)
Tenure			0.049 (0.041)	0.032 (0.003)	-5.616 (8.640)
Tenure^2			-0.001 (0.001)	-0.001 (8.92E-05)	-0.010 (0.185)
College			0.012 (0.013)	-0.026 (0.016)	-0.006 (2.344)
Segment Size				0.668 (0.018)	0.332 (0.018)
Car model indicators	Included	Included	Included	Included	
Month indicators	Included	Included	Included	Included	
Loglikelihood	-131426.257	-131425.947	-131272.978	-129971.165	
BIC	263404.641	263416.292	263147.162	260604.883	

*Prices and commissions are converted in \$1,000 for estimation

Table 5. Price and commission elasticities

		Price	Commission
MPV	Own	-6.103	0.768
	Cross	0.148	-0.024
Sedan	Own	-5.336	1.011
	Cross	0.026	-0.006
SUV	Own	-6.235	0.824
	Cross	0.311	-0.044

We next turn to how responsive salespeople are to commissions. Because the estimates of the coefficient on commissions follows the form of equation (8), it can be hard to interpret the overall effect of commissions on salesperson utility. For this reason, we report the commission coefficient averaged across all salespeople in the row labeled “Average commission sensitivity” in Table 4.

In terms of salesperson heterogeneity, our proposed model also allows the responsiveness to commissions to differ among salespeople based on unobserved types. The last model in the table shows that there are two distinct latent segments. The majority segment (67% of salespeople) has a high sensitivity to commissions, with salespeople with a higher tenure being more responsive to commissions. The smaller segment has a much lower commission sensitivity. This may imply that salespeople in that segment have a low selling ability and therefore the cost of service effort is high.¹⁷ Finally, college education does not have a significant effect for either segment.

We also compare the behaviors of the salespeople identified as belonging to the different segments. We first calculate the posterior probability that a salesperson belongs to segment s as $Prob(\theta_s | y)$, and assign each salesperson to the segment to which they have the highest probability

¹⁷ The effect of tenure on commission sensitivity for this segment is statistically insignificant, but if anything, it is declining over time. This decline in commission sensitivity over time perhaps reflects the increases in base salary with tenure. In Japan, the base salary largely depends on the seniority of workers and less on the work performance.

of belonging. We then run two reduced-form regressions to show how the two segments behave differently in the data. In the first regression, the dependent variable is a salesperson’s aggregate sales in a month divided by the total number of car models that are available in that month.¹⁸ The independent variables include an indicator variable that takes a value of 1 if the salesperson belongs to segment 2 (which is less responsive to commissions), as well as the interactions between the average commission across car models that are available in that month and the indicator variables for the two segments. In the second regression, we further control for month and year fixed effects. The regression results are presented in Table 6 below.

Table 6 Regression results for salespeople sales by segment

	Spec 1	Spec 2
Average Commission*Seg 1	0.905 (0.090)	0.286 (0.145)
Average Commission*Seg 2	0.272 (0.167)	-0.313 (0.197)
Seg 2	0.123 (0.072)	0.110 (0.069)
Month and Year Indicators	Not included	Included

*Commissions are converted to \$1,000 for regressions

The coefficients of the interaction terms capture how changes in commissions affect the sales of salespeople in each of the two segments. The commission effect is positive and significant for segment 1, and insignificant for segment 2 in both regressions. This is consistent with our main model findings that salespeople in segment 1 respond positively to commission incentives but those in segment 2 do not. The fixed effect of segment 2 is positive but insignificant. Ultimately,

¹⁸ Doing so helps minimize the concern that a larger availability of models may increase the demand, as our choice model implies.

considering that the interaction effect between commission and segment 1 is positive and the average commission across car models is at least \$360 in our data, we observe that salespeople in segment 1 generate higher sales than those in segment 2. These results verify our findings from the structural model that salespeople in segment 1 are more effective in generating sales and more responsive to commission incentives than salespeople in segment 2.

3.4 Alternative information assumptions

In our main model, we assume that salespeople obtain perfect information about consumer preferences after talking to the consumers. In this subsection, we briefly discuss two alternative assumptions: 1. Salespeople only know the distribution of consumer preferences but not any individual's preferences, and 2. Salespeople know consumers' preferences but they are uncertain about the effectiveness of the selling effort.

We first consider the case where salespeople only know the distribution of consumer preferences. We also assume that the salesperson only applies effort to one product. Despite its seeming simplicity, this case is computationally difficult to estimate because we can no longer analytically solve for the optimal effort as a function of the other parameters. Because of this computational burden, we only estimate a simpler version without salesperson heterogeneity and compare it to our structural model without salesperson heterogeneity. The results are presented in the Appendix in Table A1. While comparing only the complete information and no information models without salesperson heterogeneity is not the ideal way to ascertain the validity of the complete information model, the estimated parameters of the complete information model, as well as the average commission coefficient, are very close to what we obtained in Table 4 for our complete model. Thus, the comparison between the two models in Table A1 is reasonably

informative, if imperfect. The model estimates in the no-information model remain qualitatively similar to those in the complete information model, except for the coefficients on commissions, which become larger. The reason for the difference in commission sensitivity is that when the salesperson is unable to ascertain the consumer's preferences they do not know if their efforts will be effective. This causes the salesperson to put in less overall effort unless they value commissions more, which is what the model needs to do in order to reflect the fact that commissions affect sales. However, the fit is much worse when we assume that salespeople do not observe individual preferences, as measured by the likelihood or BIC. This implies that salespeople do indeed process private information on the consumers. This result is consistent with the findings of Kim et al. (2018), who detect evidence that salespeople from a microfinance bank process private information about the customers.¹⁹

We next consider an alternative scenario where the salespeople are able to observe consumers' preferences, but they are uncertain about the effectiveness of their selling effort. In this case, everything is the same as in our main specification, except that the service effort that is received by consumers becomes

$$e_{sijt} = \begin{cases} e_{sijt}^* + \eta_{sijt} & \text{if effort is put towards product } j \\ 0 & \text{otherwise} \end{cases}$$

where e_{sijt}^* is the amount of effort the salesperson exerts, and $\eta_{sijt} \sim \text{Logistic}(0, \sigma)$. This specification means that the effort the salesperson exerts is of uncertain efficacy, but only for the product for which the effort is made. The challenge of incorporating such an error term is that the scale parameter σ is difficult to identify from the data because the salesperson's effort is unobserved. Moreover, the computational burden is dramatically higher, as we have to solve for

¹⁹ We also compute the elasticities from this alternative model, and find that when we assume that salespeople do not have private information, the commission elasticities are much smaller, and thus cannot capture the effect of commission on product sales as well as our proposed model.

the optimal effort numerically in an inner loop of the estimation. Because of the computational burden, we only estimate the model without salesperson heterogeneity, and set σ to 0.5. We compare the estimation results with the complete information model without salesperson heterogeneity, which are presented in Appendix Table A2. The estimates of this model are very similar to those of the complete information model, but the fit as measured by the likelihood and BIC is lower. The results are similar when we vary the value of σ between 0.1 and 1. Thus, we believe that our full information specification is the best specification to use in our market.

Finally we note that the complete information assumption is reasonable in our empirical setting. As many researchers have found, trust is highly valued in Japan. For example, Dyer and Chu (2003) find that trust is significantly higher in Japan than in Korea or the United States between suppliers and automakers. There is greater information sharing among different parties in transactions. In addition, because prices are rarely negotiated in Japan consumers do not have as much of an incentive to hide information from salespeople.

3.5 Counterfactuals

We use the estimated structural model to run two counterfactual policy experiments. First, we study the optimal commission schedule for each car and examine how the optimal commissions depend on the appeal of products for consumers, the substitutability of the products, and profit margins. We then examine whether it is better for the dealer to try to sell more cars through discounted prices or through sales incentives.

3.5.1 Optimal commissions

We first consider how the attractiveness of products, the substitutability between products, and the profit margins of products affect the optimal commissions. We take the car models that are available in the first month in our data and randomly draw the price, cost, car model fixed effect (which is a measure of the attractiveness of the car model to consumers), and the within-category correlations of preferences for all of the car models in our data. The other parameters remain the same as in the proposed full model in Table 4. For each simulation, we calculate the optimal commission for each car model. We repeat this simulation exercise 100 times,²⁰ and then run a linear regression with the optimal commission of each car model as the dependent variable, and use the baseline attractiveness of the focal car (i.e. the difference between the car fixed effect and the average disutility from price across consumers), the average baseline attractiveness of other cars, the profit margin of the focal car, the average profit margin of the other cars, and the correlation between preferences for the cars as covariates.

The optimal commission results are reported in Table 7. We find that the higher the attractiveness of the car, the lower the optimal commission the dealer should set because making the sale does not require much effort from salespeople to sell the product. Second, the higher the attractiveness of substitute products, the lower the optimal commission should be because consumers will still purchase other products if effort is not exerted for the focal product. Third, a higher profit margin implies a higher optimal commission. In contrast, to avoid cannibalization the commission should be lower if other products have higher profit margins. Finally, a high preference correlation across products suggests that the focal product has close substitutes. To avoid cannibalizing the sales of substitute products the dealer should also set a lower commission

²⁰ The ranges we draw from are as follows: price [8.05, 22.39], cost [6.79, 20.73], commission [0.17, 2.55], fixed effect [-2.20, 2.24], correlation [0, 0.91]. These are the ranges from the observed data. We draw from a normal distribution with either the mean of the data or the mean of parameter estimates (for the fixed effects and correlations). Prices are also constrained to be above the cost of the car.

for the focal product. All of these results give useful rules of thumb for firms when deciding how to set commissions.

Table 7. Results from the optimal commission regression

Intercept	-0.74 (0.27)
Baseline attractiveness of focal car	-0.08 (0.01)
Average baseline attractiveness of other cars	-0.12 (0.07)
Profit margin of focal car	0.61 (0.01)
Profit margin of other cars	-0.10 (0.03)
Preference correlation across cars	-0.04 (0.02)
R^2	0.92

3.5.2 The profit impact of discounting prices and increasing commissions

In marketing, the use of price promotions targeting end consumers is a popular “pull” strategy to stimulate sales. Another strategy is to use commissions to incentivize salespeople to invest service efforts, which increase sales. This can be viewed as a “push” strategy since it targets the middleman in a transaction process. In particular, end consumers do not directly benefit from commissions, nor are they usually aware of how much commissions a salesperson earns. Our model allows firms to evaluate and compare the effectiveness of the “pull” and “push” strategies within a unified framework.

We use a counterfactual exercise as an illustration. Suppose that the dealer’s goal is to increase the sales of car model 6 in the MPV class by 50% in the first month of the data. This could arise because of the need to clear the excess inventory. The dealer can either discount the selling

price of the car model or increase the commission for salespeople who successfully sell a car.²¹ We calculate the level of price discount or the commission increase the dealer has to offer in order to achieve their goal. We also calculate the expected unit sales of the other car models in order to calculate the total profit impact for the dealer. Results are reported in Table 8.

Table 8. Results from a price discount and a commission increase

	Current setting	Price discount	Commission increase
\$ Amount		\$816	\$205
Car model 6 unit sales	144	216	216
Seg 1	125	186	197
Seg 2	20	31	20
Other cars unit sales	647	638	636
Seg 1	549	540	538
Seg 2	98	98	98
Total net profit	\$770,683	\$722,492	\$853,584
Seg 1	\$666,085	\$622,990	\$753,021
Seg 2	\$104,598	\$99,502	\$100,564

Table 5 indicates that the magnitude of the own-price elasticities is much larger than the own-elasticity for commissions for all three car classes, which would seem to imply that price discounts are more effective than commission increases. Table 8, however, shows that this intuition is incorrect. To increase the sales of car model 6 by 50% (i.e. from 144 units to 216 units), the dealer has to offer consumers an \$816 price discount, but it only needs to offer salespeople a \$205 commission increase. The commission increase is only one-fourth of the price discount required to achieve the sales growth. The reason is that because the average selling price of model 6 is

²¹ When the selling price drops, the commission will also decrease based on the way it is calculated. To separate the effects of the “pull” from the “push” strategy, we assume that the commission rate will be adjusted upward to make the commission remain unchanged in the first case. We also assume the selling price remains unchanged when the commission increases in the second case.

\$14,440 (see Table 2), \$816 is merely a 5.7% price discount. The \$205 commission increase, however, represents a 56.7% increase from the original \$360 commission level. In both cases, the unit sales of other car models will only slightly decrease. Table 8 also shows that, compared with the current setting, the dealer's total net profit will decrease by 6.25% if they use the price discount, while it will increase by 10.75% by raising the commission.

The strategies work differently for the two segments of salespeople. While the price discount increases sales from both segments, the commission increase only motivates segment 1, the segment that has a higher selling ability and is more responsive to commission changes, to exert more effort. Consequently, the commission increase will increase the dealer's net profit coming from the salespeople in segment 1, yet it reduces the profit from the salespeople in segment 2. The price discount will decrease the dealer's profit from both segments.

Although the above counterfactual suggests that commission increases are more profitable than price discounts, this result crucially relies on how consumers respond to price changes and salespeople respond to commission changes. We further investigate how a firm's optimal strategy may change under different commission coefficients. In the exercise, we lower the commission coefficient (α) from the average of 2.58 (see the last column of Table 4) to 1, 0.7 or 0.5, respectively, while keeping other model parameters unchanged. We again assume that the dealer's goal is to increase the sales of car model 6 by 50%, and calculate the required price discount and commission increase for the dealer. The results are reported in Table 9. When the commission coefficient is still high enough (i.e. $\alpha = 1$), we find that it is still more profitable for the dealer to increase the commission and not to cut the price. When the commission coefficient is much lower (i.e. $\alpha = 0.5$), however, it is more profitable for the dealer to only offer price discount to consumers. This demonstrates the importance of quantifying the sensitivity of salespeople in

response to commission changes for the optimal promotion strategy. Interestingly, when the commission coefficient is at the medium level (i.e. $\alpha = 0.7$), it is better for the dealer to combine both “pull” and “push” strategies, implying that price discounts and commission increases can complement the effectiveness of each other. When there is a price discount, salespeople find it easier to sell the car to wider audience of consumers; when the price discount is accompanied by a commission increase, the salespeople have a greater incentive back up the sale with greater sale effort. These results are also consistent with car companies using both dealer cash and customer cash, even though dealers only pass through a small fraction of the dealer cash (Busse et al. 2006).²²

Table 9. Optimal promotional strategies under different commission sensitivities

	Price discount	Commission increase
$\alpha = 1$	\$0	\$559
$\alpha = 0.7$	\$456	\$311
$\alpha = 0.5$	\$770	\$0

4. Conclusion

We develop a salesforce-driven model of consumer choice to study how performance-based commissions incentivize a salesperson’s service effort toward heterogeneous, substitutable products. This study bridges the gap between the consumer-choice literature and the salesforce-management literature. In particular, the latter literature generally focuses only on how commissions affect the performance of salespeople in terms of aggregate sales. In contrast, we model the selling process as a joint outcome of decisions by a salesperson and a consumer. The

²² Whether it is optimal to cut prices or increase sales incentives also depends on the baseline prices and commission rates. Our key takeaway is to help lay out when one tool is likely to outweigh the other.

likelihood of selling a product is influenced not only by the salesperson's efforts induced by commissions but also by the consumer's innate product preferences. Our model thus allows the salesperson's efforts to vary across different transactions depending on the unique product preferences of each consumer.

We estimate the model using data from a car dealership in Japan. The results show that not only do consumers have heterogeneous product preferences but also salespeople have heterogeneous sensitivity towards commissions. We then use counterfactuals to compare the effectiveness of using price promotions vs. commission incentives to increase sales. We also illustrate how product-specific commissions should be set differently depending on the popularity and substitutability of products.

Although our model makes significant progress toward accounting for salesforce effort in modeling choices, it also has limitations. First, we assume that salespeople have perfect knowledge regarding consumers' product preferences. Future research may investigate how, when a salesperson is uncertain of the consumer's type, they will allocate different levels of service efforts across differentiated products and how such uncertainty will impact the effectiveness of using commissions as incentives. Second, it would be interesting to see if consumers differ in their responsiveness to salesperson effort. However, our data do not allow us to study this question because we do not have any customer level information – each customer only shows up once in our data, and we do not know any demographic data about the customer, such as their gender or age. We could study this type of heterogeneous responsiveness if we had data about consumers, or even better if we had a consumer panel with repeat purchases. Third, if salespeople push the wrong products to consumers due to high commissions, this can lead to increased dissatisfaction and reduced trust among consumers; therefore, a short-term increase in profits may lead to a long-

run cost for the dealer. In contrast, price promotions directed to consumers will not have this issue. Such long term consequences are not considered in this paper. Similarly, our model also abstracts from how a price promotion may act as an advertising tool that can attract more store visits from customers. Finally, we only study consumer and salesforce decisions in a static framework, ignoring any dynamics that may arise if quotas or ratcheting exist. Current performance might also affect a salesperson's future base salary and career movement. Therefore, the salesperson may face a dynamic optimization problem, which is beyond the scope of the current study. We view this study as the first step in the literature that simultaneously models the two-sided decisions in a market with differentiated products. We hope our study will lead to more research in the future that can address the above issues.

References

Albuquerque, P., & Bronnenberg, B. J. (2012). Measuring the impact of negative demand shocks on car dealer networks. *Marketing Science*, 31(1), 4-23.

Basu, A. K., Lal, R., Srinivasan, V., & Staelin, R. (1985). Salesforce compensation plans: An agency theoretic perspective. *Marketing science*, 4(4), 267-291.

Busse, Meghan, Jorge Silva-Risso and Florian Zettelmeyer (2006), "\$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions," *American Economic Review*, 96(4), 1253-1270.

Chung, D. J., Steenburgh, T., & Sudhir, K. (2013). Do bonuses enhance sales productivity? A dynamic structural analysis of bonus-based compensation plans. *Marketing Science*, 33(2), 165-187.

Copeland, A. M., & Monnet, C. (2003). The welfare effects of incentive schemes. *Working paper*.

Daljord, Ø., Misra, S., & Nair, H. S. (2016). Homogeneous Contracts for Heterogeneous Agents: Aligning Sales Force Composition and Compensation. *Journal of Marketing Research*, 53(2), 161-182.

Farley, J. U. (1964). An optimal plan for salesmen's compensation. *Journal of Marketing Research*, 39-43.

Kim, T. T. (2014). When Franchisee Service Effort Affects Demand: An Application to the Car Radiator Market and Resale Price Ceiling. *Working Paper*.

Kim, M., Sudhir, K., Uetake, K, Canales, R. (2017). Multidimensional Sales Incentives in CRM Settings: Customer Adverse Selection and Moral Hazard. *Working Paper*.

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.

Mantrala, M. K., Albers, S., Caldieraro, F., Jensen, O., Joseph, K., Krafft, M., ... & Lodish, L. (2010). Sales force modeling: State of the field and research agenda. *Marketing Letters*, 21(3), 255-272.

Misra, S. and Nair, H.S. (2011). A structural model of sales-force compensation dynamics: Estimation and field implementation. *Quantitative Marketing and Economics*, 9(3), 211-257.

Petrin, A., & Train, K. (2010). A control function approach to endogeneity in consumer choice models. *Journal of marketing research*, 47(1), 3-13.

Rao, R. C. (1990). Compensating heterogeneous salesforces: Some explicit solutions. *Marketing Science*, 9(4), 319-341.

Srinivasan, V. (1981). An investigation of the equal commission rate policy for a multi-product salesforce. *Management Science*, 27(7), 731-756.

Weinberg, C. B. (1975). An optimal commission plan for salesmen's control over price. *Management Science*, 21(8), 937-943.

Yang, B., Chan, T.Y., Owan, H., Tsuru, T (2017). Incentives from compensation and career movements. *Working Paper*.

Zoltners, A., Sinha, P., & Lorimer, S. (2008). Sales force effectiveness: A framework for researchers and practitioners. *Journal of Personal Selling and Sales Management* 28, 115-131.

Appendix

Table A1: Comparison of the proposed model and the no information model

	Complete information model	No information model
Consumers		
Price coefficient: mean	-0.244 (0.014)	-0.163 (1.140)
Price coefficient: s.d.	0.006 (0.013)	0.028 (0.046)
Correlation:MPV	0.929 (0.055)	0.576 (0.005)
Correlation:Sedan	0.011 (0.414)	1e-4 (0.061)
Correlation: Mini-SUV	0.690 (0.061)	0.737 (0.004)
Price control	0.120 (0.008)	0.041 (0.292)
Salespeople		
Commission Coefficient	3.021 (0.336)	5.400 (0.001)
Car model indicators	Included	Included
Month indicators	Included	Included
Loglikelihood	-131426.257	-131931.000
BIC	263404.641	264426.398

Table A2: Comparison of the proposed model and the incomplete information model with $\sigma=0.5$

	Complete information model	Incomplete information with $\sigma =0.5$
Consumers		
Price coefficient: mean	-0.244 (0.014)	-0.187 (0.012)
Price coefficient: s.d.	0.006 (0.013)	0.006 (0.001)
Correlation:MPV	0.929 (0.055)	0.927 (0.023)
Correlation:Sedan	0.011 (0.414)	0.015 (1.33e-4)
Correlation: Mini-SUV	0.690 (0.061)	0.685 (0.008)
Price control	0.120 (0.008)	0.144 (0.338)
Salespeople		
Commission Coefficient	3.021 (0.336)	3.214 (1.70e-4)
Car model indicators	Included	Included
Month indicators	Included	Included
Loglikelihood	-131426.257	-131607.099
BIC	263404.641	263778.596

*Prices and commissions are converted in \$1,000 for estimation