Submitted to *Marketing Science*
manuscript (Please, provide the mansucript number!)

Compensation Schemes for Salespersons Who Forecast Demand

(Authors' names blinded for peer review)

We extend the salesforce compensation literature by developing compensation schemes for salespersons who also forecast. We assume that there are different types of salespersons. The types differ in terms of forecasting capability and sales productivity. We show that if potential forecasts of each type belong to a refinement level, then the firm can cause the salesperson to reveal her true type and exert the appropriate level of effort. Unlike the traditional salesforce literature, our analysis incorporates costs of matching supply and demand, which depend on forecast accuracy.

Key words: Salesforce Compensation, Forecasting, Scoring Rules, Principal Agent Models

1. Introduction

In this paper we study compensation schemes for salespersons who are risk averse, rational, and are employed by the firm. The salesforce is in direct contact with the customers and has valuable information about the effort required to generate demand and the level of uncertainty in the market. It is well known that inducing the salesforce to reveal this information truthfully is challenging. As Chen (2005) points out, firms face a conflict between getting the salespersons to exert effort and to reveal market conditions. In the traditional salesforce management literature (Coughlan 1993, Basu et al. 1985, Lal and Staelin 1986, Chen 2005) the firm depends on the salesforce to reveal the first moment of the demand distribution, while all other moments are common knowledge. In general, costs associated with mismatches in supply and demand depend on all the moments of the demand distribution (Ridder et al. 1998, Porteus 2002). In our work we require the salesforce to provide information about all the moments of the demand distribution. Thus we account for situations in which the salesforce has better information about the average demand and the uncertainty in the market. We assume there are different types of salespersons who differ in terms of forecasting and sales capabilities. For these more general settings, we determine conditions under which a firm can induce the salesperson to reveal her true type and exert the targeted effort level. We believe that our work is the first to explicitly model forecast uncertainty while developing optimal salesforce compensation schemes.

Forecast distributions can be partitioned into sub-sets called refinement levels (Dasu and Ochiumi 2005, DeGroot and Fienberg 1983). Refinement levels, summarized in section 3, are used to order forecasting capabilities. If all forecasts generated by a saleperson belong to the same refinement level, then the firm can induce the salesperson to reveal her type and exert the optimal effort level while compensating the salesperson at her reservation level. In general the compensation received by the salesperson under these schemes may not be monotone increasing in the realized demand. If, however, the salespersons are identical in terms of productivity but differ only in their forecasting abilities, then monotone increasing payment schemes can also ensure truth-revealing.

In the next section we review related literature. In Section 3 we review the concept of refinement and introduce our model. The main results are in Section 4.

2. Literature

There is a body of literature that analyzes optimal compensation schemes for managing a salesforce. For a comprehensive review, the reader is referred to Coughlan (1993). Work by Basu et al. (1985) and Lal and Staelin (1986) are representative of research in this field that is related to our work. Basu et al. (1985) assume that the sales of a product depend on the salesperson's costly effort. They assume homogeneity in the salesforce and information symmetry. Lal and Staelin (1986) extend Basu et al. (1985) by permitting heterogeneity and information asymmetry. In their work, demand (x_i) is a random variable that is expressed as a sum of a deterministic term ($\eta_i(t)$) and random shock (ϵ_i). Here i represents the salesperson type, t is the effort level, and $\eta_i(\cdot)$ is a monotone increasing deterministic function. There are two types of salesperson – high type and low type. Lal and Staelin (1986) allow the distribution of the demand, not just the mean, to vary with salesperson type. But they impose a restriction by assuming that for any given effort level t, the realized demand from a high type is stochastically larger than that from a low type.[1] Their focus is on difference in sales capabilities and not forecast precision. They are interested in compensation schemes that optimize the mix of salespersons employed by the firm. They can exclude a low type from participating but otherwise cannot separate the two types. Also, just as with the rest of the salesforce compensation literature, Lal and Staelin (1986) and Coughlan (1993) assume that the firm's payoff depends only on the expected demand. They are not concerned with inventory costs or other costs associated with matching supply and demand. These costs depend on the demand

[1] If ϵ_i are normally distributed, then the standard deviation must be the same for both the high and the low type. This is because they require demand for the high type to be stochastically larger than that for the low type. Consequently, forecast precision is the same for high and low type salespersons.

distribution and not just the expected sales. We account for these costs. Consequently, traditional salesforce literature is a special case of our model.

Given our interest in forecast precision, we assume that there are different types of salesperson and allow forecast precision and sales productivity to vary by type. We let the firm's profits depend not just on the expected demand but on the entire demand distribution. This allows us to account for costs that arise from imbalances in supply and demand. Under these general conditions, we identify conditions that enable the firm to induce the salesperson to truthfully reveal the type and pick the appropriate effort level.

There are a few papers in the operations management literature that address salesforce incentives. An interesting paper that examines the impact of salesforce compensation on inventory and production costs is Chen (2000). He develops incentives that induce the salesforce to smooth demand over time. The paper, however, is not concerned with forecast accuracy, as it assumes common information. Huff et al. (2003) examine the interactions between stocking decisions and salesforce compensation in a newsvendor environment. In their model, demand distribution is stochastically increasing both in the firm's effort and in the salesperson's effort. Their main interest is to address the issue of double-sided moral hazard under various scenarios, and one of their scenarios is similar to our model in Section 4. They find that the firm can delegate stocking decisions to its risk-neutral salesforce without loss of profit. Our model in Section 4 is an extension of their model in several ways. First, their agent is risk-neutral while ours can be risk-averse. Second, Huff et al. (2003) assume that both the principal and the agent have the same information about the future demand whereas we assume information asymmetry. Third, in their model, the agent partially bears the overage/underage costs. In our model, the agent is compensated solely based on submitted forecast distributions and realized demand.

Chen (2005) explicitly incorporates the costs of matching supply and demand. He assumes that a firm employs a single salesperson. This person may be one of two types. He adopts a demand model that is very similar to that of Lal and Staelin (1986), by letting demand: $x = a + \theta_i + \epsilon$. Here a is the agent's selling effort, θ_i is a constant that varies by type, and ϵ is a normal random variable. The error term ϵ does not depend on the type. The payoff to the firm depends not only on the realized demand but also on over-stocking and under-stocking levels, just as in the classic newsvendor problem. For this model, he compares the performance of a menu of linear contracts with that of the schemes proposed by Gonik (1978), whose work is often cited as an example of compensation schemes that have been implemented. Chen (2005) finds that a menu of linear contracts dominates piece-wise linear schemes proposed by Gonik. Our work extends Chen's work by generalizing the cost function. The penalty costs that arise in the newsvendor problems are a special case of our cost function. Although Chen incorporates the cost of matching supply and demand, he is not concerned with forecast precision and forecasting abilities of salespersons. The forecast distribution in his analysis is normally distributed with a fixed standard deviation. We permit the demand distribution to vary by the salesperson's type that is linked to her forecast precision. We study optimal contracts that take into account forecast precision, while he analyzes the performance of linear contracts.

3. Model Assumptions and the Concept of Refinement

Let x denote the forecasted demand, $f(\cdot)$ the forecaster's true forecast, and $\hat{f}(\cdot)$ the distribution that the forecaster announces to the firm. The events occur in the following sequence.

(i) The firm announces a menu of contracts $\xi(\hat{f}(\cdot), x_0)$ as a function of the announced distribution $\hat{f}(\cdot)$ and the observed demand x_0.

(ii) Each salesperson selects a contract by announcing a distribution $\hat{f}(\cdot)$.

(iii) Each salesperson exerts effort that maximizes her expected utility.

(iv) Once the forecast(s) is received, the firm makes its optimal decision $\Gamma(\hat{f}(\cdot))$. $\Gamma(\hat{f}(\cdot))$ may be the quantity produced or the stock level.

(v) The demand x_0 is observed, and payment ξ is made according to the chosen contract.

Since the forecasted variable is demand, some discussion is in order. Demand is not always observable. In the retail industry where consumers have direct access to inventory, lost sales are difficult to observe. There are many other situations in which the consumer does not have access to inventories. Examples include catalog based sales, Internet based sales, sales from manufacturers to distributors, and sales of industrial supplies and raw materials that are procured on the basis of purchase orders. In all these cases, the firm observes demand. As a concrete example, consider Carhartt, a manufacturer of jackets worn by industrial workers. A supplier to Carhartt that we have studied anticipates orders from Carhartt and produces the fabric in anticipation. Garment production is initiated only after receiving a purchase order. If the order size is larger than anticipated, additional fabric is produced on an expedited basis. From a modeling perspective, we assume, like Chen (2005), that there exists an emergency action such as special expedited order, so that demand is always met.[2]

Demand forecasting is also used in call centers to aid staff planning. Most call centers can observe lost calls or dropped calls and therefore can observe demand.

Some salespersons are better at forecasting than others. The productivity of effort may differ among salespersons. Furthermore, salespersons may differ in their disutility of effort.

[2] In operations management, particularly in inventory theory, there is a long tradition of assuming that lost sales are back-logged.

Accordingly, we assume that there are multiple salesperson types with different forecasting capabilities, different productivity functions, and different disutility functions. Associated with each type are a base forecast distribution, a sales productivity function, and a disutility function. For a type i salesperson, the base density is $f_i(x)$. If a type i salesperson exerts t units of effort, demand is distributed as $f_i(x|t)$, and the salesperson incurs disutility of $V_i(t)$. The firm is aware of all the possible $f_i(\cdot)$ and V_i but does not know which agent is of which type. The firm, however, has a prior distribution about the likelihood of a salesperson's type.

Thus far, we have not specified the firm's objective function . We show below how our model incorporates operational costs that depend on forecast uncertainty. Traditional salesforce compensation literature has not taken these costs into account. Recall that, given a forecast $f(\cdot)$, the firm's best decision is $\Gamma(f(\cdot))$. While making this decision the compensation paid to the salesperson is treated as a sunk cost. For simplicity let us also assume that this decision is measured in the same unit as the forecasted variable. For example, $\Gamma(f(\cdot))$ may the optimal stocking level. The objective function ($\Pi(\cdot)$) of decision problems encountered when there is supply-demand mismatch cost can frequently be parsed into two components. One component depends on the realization of the forecasted variable and the second component is a penalty for not correctly anticipating the state of nature.

$$\Pi(x_0,\Gamma(f(\cdot))) = \pi(x_0) - \kappa_e([\Gamma(f(\cdot)) - x_0]^+) - \kappa_s([x_0 - \Gamma(f(\cdot))]^+), \tag{1}$$

where $\kappa_e(\cdot)$ and $\kappa_s(\cdot)$ are monotone non-decreasing positive valued functions; and $\pi(\cdot)$ and $\Pi(\cdot)$ are real valued functions. The set of problems where this decomposition applies includes inventory problems, capacity planning, manpower scheduling, and production planning. We formulate the salesforce incentive problem using this objective. This is done purely to illustrate the link to operational problems and the need for considering higher moments of forecast distributions.

3.1. Refinement

Dasu and Ochiumi (2005), building on the work by DeGroot and Fienberg (1983), partition the space of forecast distributions into sets that correspond to different levels of refinement. The possible set of forecasts is given by $\mathcal{F}$. $\mathcal{F}$ is such that if $f(x) \in \mathcal{F}$, then $f(x) \geq 0$, $\forall x$, and $\int f(x)dx = 1$. The set $\mathcal{F}$ is partitioned into *refinement sets* $\chi(i) = \{f_{i,s}(\cdot)\}$, where i is a refinement level, and s and r are indices, such that:

(a): All the members of the set $\chi(i)$ are linearly independent.

(b): If $H > L$, then for any $f_{L,r}(x) \in \chi(L)$, there exists a density $\alpha_r(s)$ such that $f_{L,r}(x) = \int f_{H,s}(x)\alpha_r(s)ds\ \forall d$, where $f_{H,s} \in \chi(H)$.[3]

[3] We use the phrase "$f_{L,r}$ is a convex combination of $f_{H,s}$," although convex combination normally implies convex combination of a finite number of elements.

Note $\chi(i)\cap\chi(j)=\emptyset$ if $i\neq j$. For a given partition, a salesperson is considered to be at refinement level H if all her true forecasts belong only to the set $\chi(H)$, and forecasts are well calibrated in the sense defined in Dawid (1982) and DeGroot and Fienberg (1983). In summary, a salesperson is well calibrated if the long run relative frequency of a measurable outcome is the same as the probability assigned to that outcome, holding the information available constant.

The set of probability distribution functions can be partitioned into refinements levels in a number of different ways. Given a partitioning scheme, we can always determine the refinement level of any given distribution, but there may not be any parameters of the density function that allow us to easily do so. There are, however, a few families of distributions in which one or two parameters can be used to define the level of refinement.

(a): If all the forecast distributions are uniform or normal then the standard deviation of the forecast determines the refinement levels. Distributions with smaller standard deviations belong to higher refinement levels.

(b): Given any distribution $f(\cdot)$ let $S[f]$ be the collection of distributions f^b such that $f^b(x)=f(x-b)$ for some $b\in R$. Let $f_a(x)=f(ax)$ where $a\in(0,1)$, then forecasts in $S[f]$ are of higher refinement level than those in $S[f_a]$.

4. Optimal Compensation Schemes for Salespersons with Different Forecasting Capabilities

The firm's problem is to design a set of profit maximizing compensation schemes $\xi_i(\cdot)$ that ensure truth-telling for each type of salesperson. For ease of exposition, we present our analysis assuming two types, $i=1$ or 2, and based on firm's prior knowledge a salesperson is of Type 1 with probability α and Type 2 with probability $1-\alpha$. Extension of the analysis to N types, $N\geq 3$, is relatively straight-forward. We formulate the firm's problem as a two-stage optimization problem:

$$\left(P^I\right):\quad \max_{t_1,t_2}\quad \alpha\int \Pi(x,\Gamma(f_1(\cdot|t_1)))f_1(x|t_1)dx+(1-\alpha)\int \Pi(x,\Gamma(f_2(\cdot|t_2)))f_2(x|t_2)dx-T(t_1,t_2)$$

$$\left(P^{II}\right):\quad T(t_1,t_2)=\min_{\xi_1(\cdot),\xi_2(\cdot)}\quad \alpha\left[\int \xi_1(x)f_1(x|t_1)dx\right]+(1-\alpha)\left[\int \xi_2(x)f_2(x|t_2)dx\right]$$

Subject to:

$$\int U(\xi_1(x))f_1(x|t_1)dx\geq V_1(t_1), \tag{2}$$

$$\int U(\xi_2(x))f_2(x|t_2)dx\geq V_2(t_2), \tag{3}$$

$$\int U(\xi_1(x))f_1(x|t)-V_1(t)\leq \int U(\xi_1(x))f_1(x|t_1)dx-V_1(t_1)\qquad \forall t, \tag{4}$$

$$\int U(\xi_2(x))f_2(x|t)dx-V_2(t)\leq \int U(\xi_2(x))f_2(x|t_2)-V_2(t_2)\qquad \forall t, \tag{5}$$

$$\int U(\xi_2(x))f_1(x|t)dx - V_1(t) \leq \int U(\xi_1(x))f_1(x|t_1)dx - V_1(t_1) \qquad \forall t, \tag{6}$$

$$\int U(\xi_1(x))f_2(x|t) - V_2(t) \leq \int U(\xi_2(x))f_2(x|t_2) - V_2(t_2) \qquad \forall t. \tag{7}$$

The first two constraints ensure the agent's participation. Constraints (4) and (5) place the salesperson in a worse position if she does not make the target effort. Due to the last two constraints, (6) and (7), the salesperson is in a worse position if she does not announce the true type. The utility functions for a salesperson is given by $U(\cdot)$.

Given a pair of effort levels t_1 and t_2, P^{II} determines the lowest cost at which the firm can influence the salesperson to make the effort level t_i when the salesperson is of type i. The key here is to determine whether P^{II} is feasible for all possible (t_1, t_2) combinations. We discuss the second stage optimization problem before we discuss the overall optimal compensation schemes.

4.1. The Second Stage Problem

If P^{II} is always feasible, then the firm can pick any combination of effort levels it wishes and incent the salesperson's to exert that effort level. This in turn will enable the firm to find the overall optimal solution. Unfortunately, it is not always possible for the firm to provide incentives that will cause the salesperson to make the designated effort levels.[4]

Although we cannot guarantee truth-telling in general, in the following proposition we show that if all forecast distributions for each type belong to the same refinement level, then the salesperson can be incented to reveal the true type and select any given effort level.

Theorem 1. *If $f_i(\cdot|t)$ for all t belongs to refinement set $\chi(i)$, $i = 1, 2$, and $\chi(1) \neq \chi(2)$, then there exists a set of truth-revealing compensation schemes that ensures that the salesperson reveals her true type and selects any given target effort level t_i, $i = 1, 2$.*

Theorem 1 shows that if salespersons differ in refinement levels, then firms can ensure that each salesperson truthfully reports her type. Notice that the only requirement is that distributions of a type belong to the same refinement level. We do not place any restrictions on the effort function. The effect of effort can be additive, multiplicative, or otherwise. The theorem also accommodates situations in which the shape of the distribution changes as the salesperson's effort increases. The base distributions of different types of forecasters, $f_i(\cdot|t = 0)$, $i = 1, 2$, do not have to have the same mean. Additionally, we do not place any restrictions on the disutilities associated with effort ($V_i(t)$). Technically, V_i does not even have to be non-decreasing. Finally, let us note that salespersons with more refined distributions will receive higher payments in expectation.

[4] Such an example is provided in the Appendix.

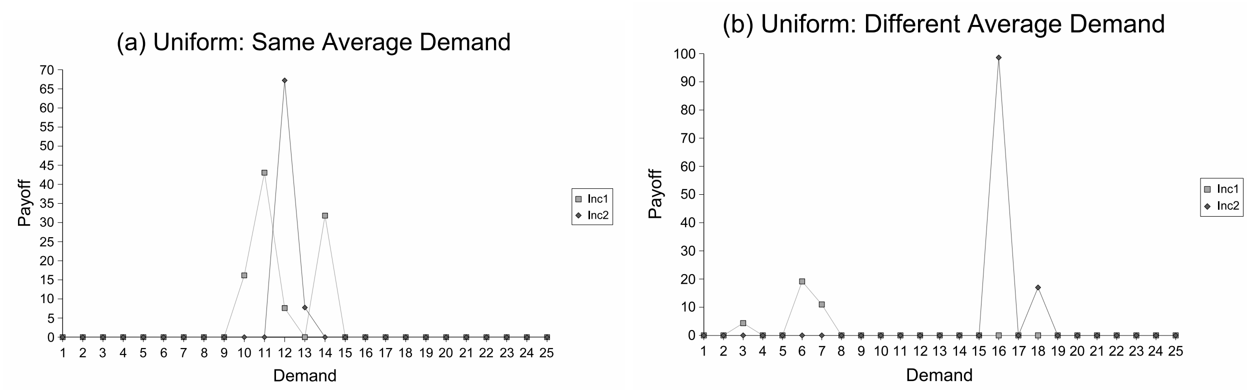

Figure 1 Uniformly distributed forecasts

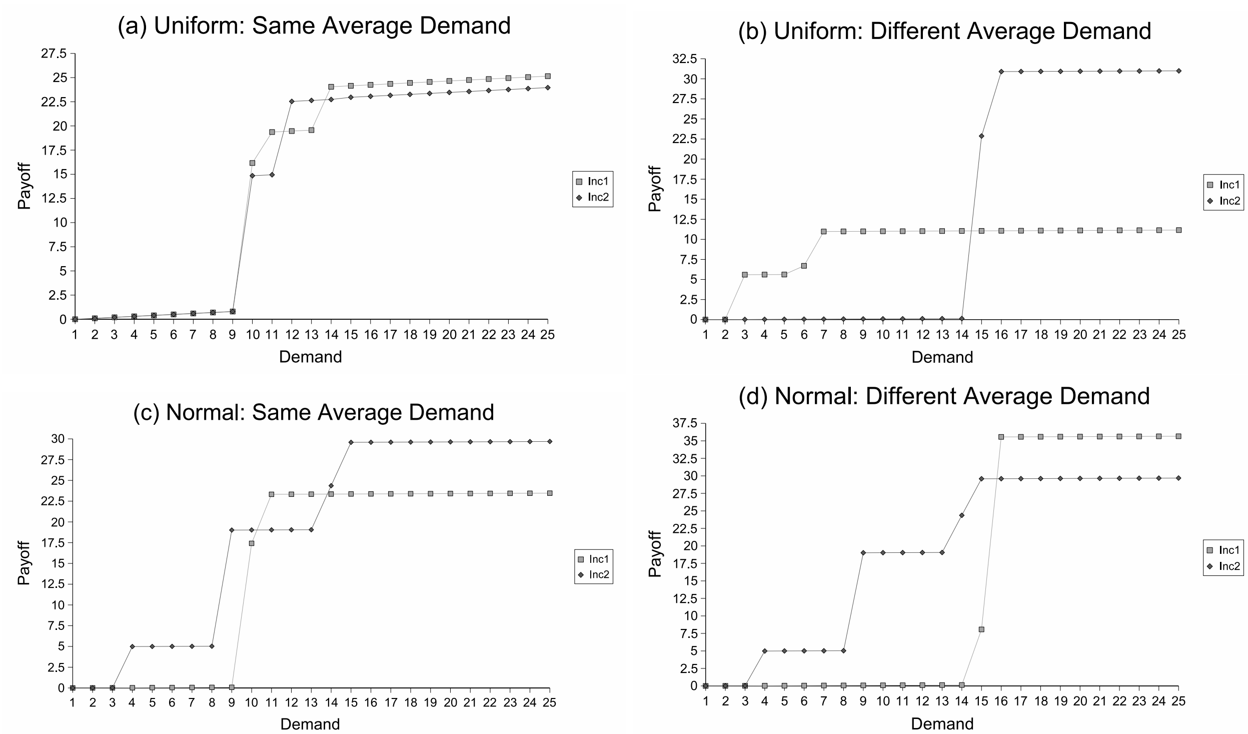

Figure 2 Monotone compensations with uniformly or normally distributed demand.

Figures 1, 2, and 3 illustrate solutions to P^{II} when the forecast distributions are either uniform or normally distributed.[5] In these figures, the X-axis is the realized demand and the Y-axis is the payment to the forecaster. In Figure 1, there are two types of salespersons. Forecasts of both types are uniform distributions but with different spreads. Potential distributions of the type 1 salesperson lie in the convex hull of the distributions of the type 2 salesperson. Effort disutility functions are the same for both types. In Figure 1-(a), the target effort levels are the same for both

[5] The details are provided in the Appendix.

types, that is, $t_1 = t_2$. Consequently, the expected demand would be the same. In Figure 1-(b), the target effort levels are set such that the type 2 agent expends greater effort than the type 1, that is, $t_1 < t_2$.

Notice the compensation schemes are not monotone non-decreasing functions of demand. In the salesforce literature, compensation schemes often are non-decreasing functions (Basu et al. 1985, Lal and Staelin 1986, Chen 2005). Although Lal and Staelin (1986) and Chen (2005) employ non-decreasing schemes, they do not establish the overall optimality of their contracts. Lal and Staelin (1986) only consider non-decreasing compensation functions. Chen (2005) finds that the optimal menu of linear contracts outperforms the optimal menu of piece-wise linear functions proposed by Gonik. Neither of these papers claims that their solutions are optimal among all the possible compensation functions.

Basu et al. (1985) assume $\frac{df(x|t)}{dt}/f(x|t)$ is strictly increasing, to show that optimal compensations are non-decreasing. They use gamma and binomial distributions. For gamma and binomial distributions, the mean and the standard deviation are determined by the effort level. In other words, precision and location of the distribution are connected via effort level. We do not assume any relationships between the location and the shape of the distribution. Also, we do not make any assumption about how the distributions shift as the agent changes her effort level. Therefore, in our models $\frac{df(x|t)}{dt}/f(x|t)$ need not be increasing, and the compensation schemes we present may not be monotone.

We have implicitly assumed that the agent does not observe sales until she finishes exerting all the intended effort. In some situations, the outcome of the sales effort can be partially observed prior to the end of the period. If we incorporate this intertemporal aspect, non-decreasing compensations may be optimal.[6]

In Figure 2, we show the result of adding a constraint to P^{II} that forces the payoffs to be monotone increasing,[7] Figures 2-(a) and 2-(b) show the results of adding this constraint to the problem depicted in Figure 1. In Figures 2-(c) and 2-(d), the underlying distributions are normal. The standard deviation of a type 2 salesperson is twice that of type 1. Figure 3 shows the case in which there are three types of salespersons. In all of the examples, the cost of effort is modeled as μ^{k_i}, where μ is the mean of the distribution and i is the type. We find that as the difference among the k_i's increases, monotone increasing policies cease to be feasible.

Although monotone increasing compensation schemes are desirable, we cannot always achieve truth-telling if we restrict the policies to be monotone increasing. Our numerical experiments

[6] Lal and Srinivasan (1993) study such intertemporal behavior of a salesperson.

[7] The details are provided in the Appendix

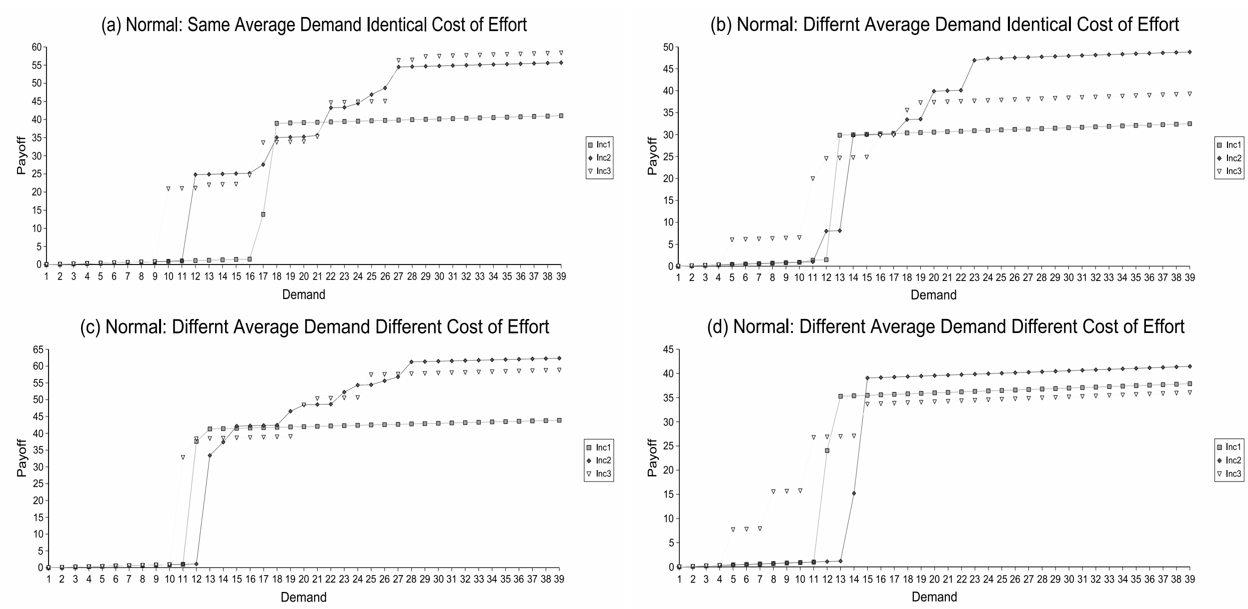

Figure 3 Monotone compensations with uniformly or normally distributed demand.

suggest that the feasibility depends on the effort functions. In the proposition below, we identify sufficient conditions under which monotone policies will be truth-revealing. We now assume that the base demand distribution of each type, $f_i(\cdot|t=0)$, has the same mean, that the effort effect is additive, i.e., $f_i(x|t)=f_i(x-\eta_i(t))$, where $\eta_i(\cdot)$ is the effort productivity function of Type i, that $\eta_1=\eta_2$, and that the sales disutility function remains the same: $V_1=V_2$.

Lemma 1. *If the cost of effort does not depend on the type, and if the mean of the distribution given an effort level is the same for two types, then we can achieve truth-telling even if we restrict the compensation schemes to be monotone increasing. Further, the salesperson will be paid in expectation an amount exactly equal to the cost of her effort.*

Thus, if the only difference between the types is the refinement level of their forecasts, then we can restrict the incentive schemes to be monotone increasing. In this case we can also show that the firm pays only the reservation utility level, regardless of the choice of effort levels for each type.

4.2. Overall Optimal Compensation Schemes

We know that so long as all forecasts belonging to a type belong to the same refinement level, the firm can induce the salesperson to reveal her true type and pick the designated effort level. This in turn enables the firm to pick the effort levels that optimizes P^I. Figure 4-(a) shows an overall optimal compensation schemes for two types of salesperson when we don't enforce monotone non-decreasing compensation schemes. Here Type 1 salesperson has normally distributed demand estimations with a standard deviation of 5, while Type 2 salesperson's standard deviation is 10. As

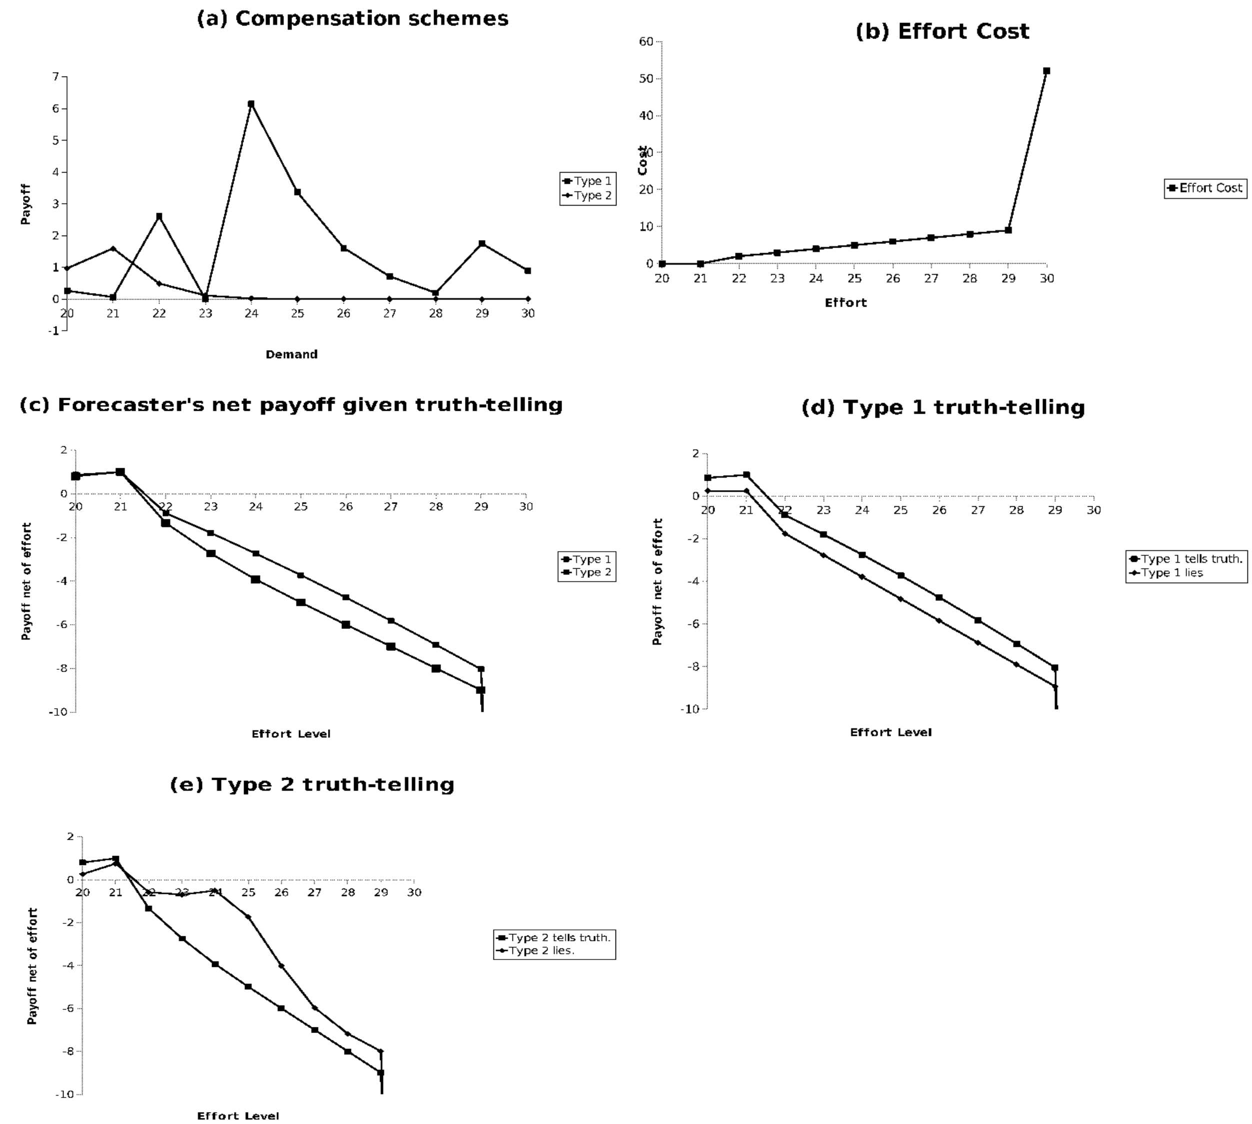

Figure 4 **Overall Optimal Solution**

each salesperson makes more effort, demand distribution shifts to the right, and the salesperson's cost of effort increases accordingly, as in Figure 4-(b). This effort cost function is the same for both types. Figure 4-(c) shows the expected payoff of each salesperson, net of effort cost, given that each salesperson truthfully reports her type. In this case each salesperson will make 21 units of effort. Finally, we can tell from Figures 4-(c) and 4-(d) that each salesperson is better of revealing her true type.

If non-decreasing compensation functions are desired, we can add appropriate constraints. After adding non-decreasing constraints, we get the results in Figure 5. Although it is not evident, in figures 5-(c) and 5-(d) the salesperson's payoffs are lower if they fail to reveal their true type.

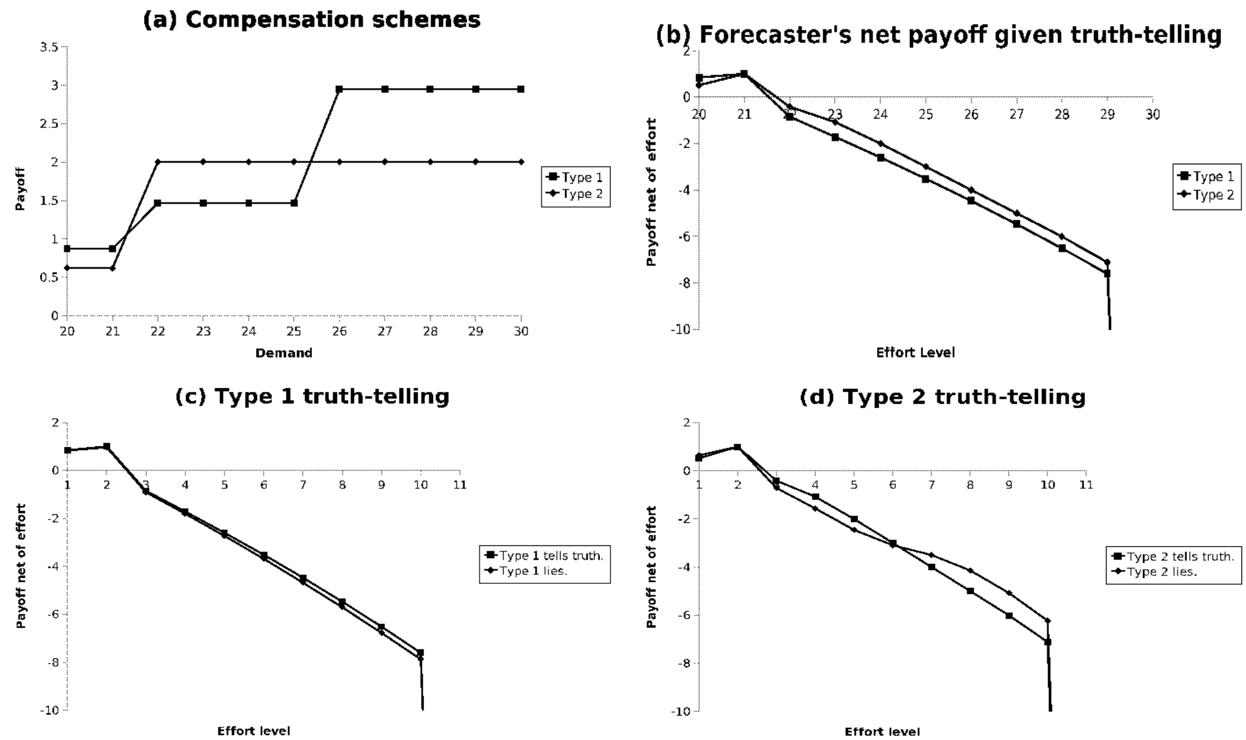

Figure 5 **Overall Optimal Solution with Monotone Compensation.**

5. Conclusions

In practice, salesforce compensation is a combination of a salary and a commission that depends on sales (Basu et al. 1985). Since salespeople are typically compensated on the basis of sales and are not responsible for any excess inventory or shortages, they have an incentive to inflate demand estimates. It is rarely the case that salespeople are compensated based on the accuracy of their demand estimates; i.e. their forecasting capabilities. However, a few attempts have been made to alleviate this problem. Gonik (1978) describes one such example that was implemented in IBM Brazil. By offering a menu of piece-wise linear contracts, IBM Brazil successfully extracted important information about the sales prospect of each sales territory and used the information for their production and logistics planning. Gonik's schemes are by no means optimal. His focus is on extracting crucial information about the local sales environment from the salesforce, not on designing optimal contracts. Our compensation schemes are more involved than those proposed by Gonik, but given Gonik's scheme's success, our compensation schemes have a potential that allows a firm to understand the market better and become more profitable. Even if the salespersons differ on their ability to forecast, we can induce them to reveal their true type and exert the given effort level. If all forecasts have normal distributions then the standard deviation defines the refinement level. In this case, our computational experiments suggest that incentive payments that are piece-wise linear may be truth revealing and near optimal for the firm.

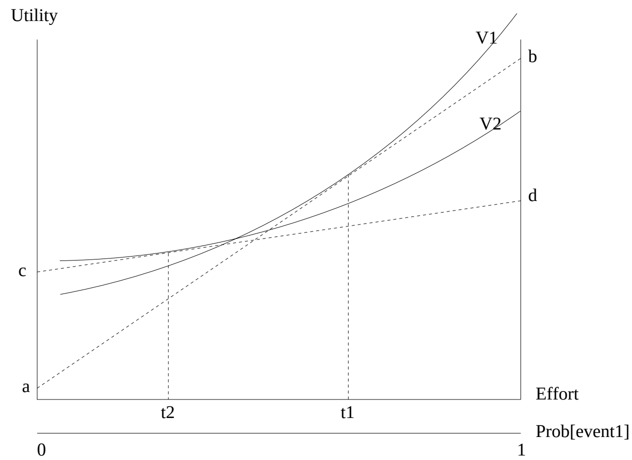

Figure 6 **Separation is impossible**

Appendix A: Example: Truth-revealing is not always possible.

We present this example to show that truth-revealing is not always possible. There are two potential outcomes (0 and 1) and two types of salespersons, Type 1 and Type 2. The salesperson's disutility function is $V_i(\cdot)$ if the salesperson is of Type i. The salesperson's effort increases the probability of outcome 1. Figure 6 depicts the effort curves and the effort levels t_1 and t_2 that the firm would like the salespersons to exert, or equivalently the probability levels that the firm desires for Types 1 and 2, respectively. This combination cannot be induced. By Hendrickson and Buehler (1971) and Dasu and Ochiumi (2005), truth revealing compensations schemes must be based on supporting hyperplanes to the effort curves. They must be straight lines with slopes $V_1'(t_1)$ and $V_2'(t_2)$. Lines (a,b) and (c,d) represent one possible set of payoff functions. If these payoff functions are employed, then salesperson of type 1 will receive a payment of \$$a$ if the outcome is 0 and a payment of \$$b$ if the outcome is 1. The corresponding payments for a salesperson of type 2 are \$$d$ and \$$c$, respectively. But under this scheme the salesperson will not reveal her true type. Any other truth revealing scheme must be lines that are parallel to (a,b) and (c,d). It is clear that shifting these lines cannot induce the salesperson to make the desired level of effort and reveal her true type.[8]

Appendix B: Example: Truth-revealing Compensation Schemes in Section 4.1

In all the examples the cost of attaining a forecast distribution equals (mean of distribution)^{k_i}, where k_i is a constant for a Type i salesperson. The firm specifies the average demand to be achieved for each type.

Figures 1-(a) and 1-(b) of the paper assume uniform distributions (discrete outcomes) and two types of salesperson with different levels of refinement. If Type 1, the probabilities are (0.2, 0.2, 0.2, 0.2, 0.2). If Type 2, the probabilities are (0.25, 0.25, 0.25, 0.25).

In Figures 2-(c) and 2-(d), the underlying distributions are normal. Standard deviation of the distribution is 2 for Type 1 salespersons whereas 4 for Type 2. We discretize the normal distribution to solve for the compensation schemes.

[8] One can easily construct other examples in which the two disutility curves do not cross.

In Figure 3, there are three types of salespersons. The underlying distributions are normal with standard deviations of 2, 4, and 5, respectively. The following table gives the remaining parameters for each of the figures.

Figure#	Average Demand 1	Average Demand 2	Average Demand 3	k_1	k_2	k_3
2(a)	12	12		1.2	1.2	
2(b)	5	15		1.2	1.2	
3(a)	12	12		1.2	1.2	
3(b)	5	15		1.2	1.2	
3(c)	12	12		1.2	1.2	
3(d)	18	12		1.2	1.2	
4(a)	20	20	20	1.2	1.2	1.2
4(b)	15	20	15	1.2	1.2	1.2
4(c)	15	20	15	1.1	1.2	1.3
4(d)	15	20	15	1.3	1.1	1.1

Appendix C: Example: Truth-revealing Compensation Schemes in Section 4.2

The table below shows the cost of attaining a mean demand, which is monotone increasing. This function is the same for both types of salespersons. Salespersons' demand is normally distributed with different standard deviations. Type 1 salesperson has a standard deviation of 5 while type 2 salesperson's standard deviation is 10. For simplicity, effort level and demand are in the same unit. For example, if a salesperson makes, say, 25 units of effort, then the mean of demand will also be 25. After solving the overall compensation problem, we get Figure 4-(a). Notice they are not monotone. Figures 4-(c), 4-(d), 4-(e) show that each constraint in the optimization problem is satisfied.

After adding monotone non-decreasing constraints, we get Figure 5.

Effort	Demand	Cost
20	20	0
21	21	0.00001
22	22	2
23	23	3
24	24	4
25	25	5
26	26	6
27	27	7
28	28	8
29	29	9
30	30	52

Appendix D: Proof of Theorem 1

Note due to condition (b) in section 3.1, forecasts of a lower refinement level lie in the interior of the convex hull generated by forecasts of a higher refinement level. Assume that forecasts of Type 1 lie in the convex hull of the forecasts of Type 2. We need two convex functionals, $G_1(\cdot)$ and $G_2(\cdot)$ that serve as the basis for the incentive schemes for the two types, respectively. Functions $G_i(\cdot)$ map from the space of probability distributions to R. The compensation schemes will be hyperplanes to these convex functionals (Hendrickson and Buehler 1971). Hendrickson and Buehler (1971) show that if a salesperson declares she is of type i then she will reveal her true forecast and will receive payment in expected utility equal to $G_i(f(\cdot))$. We have to ensure that she declares her true type and picks the right effort level. Consequently we need:

(i) For all $f_{i,s}(\cdot) \in \chi(i)$, $G_1(f_{1,s}) > G_2(f_{1,s})$ and $G_2(f_{2,s}) > G_1(f_{2,s})$.

To ensure that each type picks the assigned effort level, given any $f_{i,s}$, we need

(ii) $G_i(f_{i,s}) > G_i(f_{i,r})$ for every $r \neq s$ and $f_{i,r} \in \chi(i)$.

To show that we can construct functions $G_1(\cdot)$ and $G_2(\cdot)$ that meet conditions (i) and (ii), we need to establish two properties.

Property 1: We can construct convex functionals $G_i(\cdot)$ such that $G_i(f_{i,s}) = k_{f_{i,s}}$, where $k_{f_{i,s}}$ are an arbitrary set of scalars.

This will ensure that condition (ii) above is met. We select the $k_{f_{i,s}}$ such that (ii) holds.

Property 2: Given any convex functional $G_2^*(\cdot)$ we can construct another convex functional $G_2(\cdot)$ such that $G_2(f_{2,s}) = G_2^*(f_{2,s})$ and $G_2(f^*) < G_2^*(f^*)$ for any f^* that is in the convex hull generated by $\chi(2)$.

Armed with these two properties we can construct convex functionals that satisfy conditions (i) and (ii). Using Property 1, we generate a convex functional $G_1(\cdot)$ so that condition (ii) is met for Type 2 salesperson. Next, using Property 1, we generate a convex functional $G_2^*(\cdot)$ such that $G_2^*(f_{2,r}) > G_1(f_{2_r})$ and $G_2^*(f_{2,s}) > G_2^*(f_{2,r})$ for every $r \neq s$. Finally, using Property 2, we construct $G_2(\cdot)$ such $G_2(f_{2,s}) = G_2^*(f_{2,s})$ and $G_2(f_{1,s}) < G_1(f_{1,s})$. Recall that forecasts $f_{1,s}$ lie in the convex hull generated by $\chi(2)$. Hence we can ensure that conditions (i) and (ii) are met.

In the following we establish the two properties.

Property 1: Let $C_1 \subset \mathbb{H}$ be a set of density functions such that if $f \in C_1$, then f cannot be expressed as a convex combination of other members of the set C_1. Given a set of scalars $k_f \in \mathbb{R}$, $f \in C_1$, there is a convex functional $G : \mathbb{H} \to \mathbb{R}$, such that $G(f) = k_f$.

Proof: For each $g \in C_1$, there exists a vector a_g and $s \in \mathbb{R}$, such that:

$$
\begin{aligned}
a_g \cdot g + s =& k_g \\
a_g \cdot f + s \leq& k_f, \qquad \text{for} f \in C_1 \backslash g.
\end{aligned}
$$

By the theorem of the alternative (Corollary 5.84 of Aliprantis and Border 2007), if there is no vector a_g satisfying these constraints, then there exists $W_f \geq 0$ and W_g in $\mathbb{R}$, such that:

$$
\begin{aligned}
\sum W_h k_h <& 0 \\
\sum W_h h(x) =& 0, \qquad \forall x \\
\sum W_h =& 0.
\end{aligned}
$$

This would mean that g can be expressed as a convex combination of other members of C_1. The functional $G(f) = \sup_h a_h \cdot f$ is a convex functional with $G(f) = k_f$.

Property 2: Let $G : \mathbb{H} \to \mathbb{R}$ be a proper convex functional with linear supports on a closed convex set $C \subset \mathbb{F} \subset \mathbb{H}$. Let ∂C be the boundary of C and C^0 its interior. There exists another convex functional $\tilde{G}(\cdot)$ such that $G(f) = \tilde{G}(f)$ for $f \in \partial C$, and $\tilde{G}(f) < G(f)$ for $f \in C^0$.

Proof: C is a convex set in a Hilbert space with an interior C^0. The convex set C is properly supported by some linear operator L_f at each point f on the boundary of C. $L_f(\cdot)$ will be such that $L_f \cdot f < L_f \cdot g$ for $g \in C \backslash f$, where $\cdot$ denotes the inner product (Lemma 7.37 of Aliprantis and Border 2007). Let S_f be the supporting hyperplane for the convex functional $G(\cdot)$ at $f \in \partial C$. Let $\tilde{G}(g) = \sup_{f \in \partial C} \{S_f \cdot g - k[L_f \cdot g - L_f \cdot f]\}$

for all $g \in C$, where $k \geq 1$ is a scalar. $\tilde{G}(\cdot)$ is a convex functional because it is the *sup* of linear functionals. Convexity of $G(\cdot)$ ensures that $\tilde{G}(f) = G(f)$ for $f \in \partial C$. For $g \in C^0$ we have $G(g) > sup_{f \in \partial C}\{S_f \cdot g\} > \tilde{G}(g)$. The first inequality is true because $G(\cdot)$ is a strictly convex functional and the second inequality is true because $L_f \cdot g - L_f \cdot f > 0$. $\square$

Appendix E: Proof of Lemma 1

Let $f_i(x|t) \in \chi(i)$ denote distributions that are realized by a Type i salesperson exerting level $t \in T$ effort, $i = 1, 2$. Without loss of generality, assume

$$\int x f_1(x|t)dx = \int x f_2(x|t)dx = t,$$

and $V(t)$ is strictly convex increasing in t. We assume that forecasts of a Type 2 salesperson are refinements of the forecasts of Type 1 and there exists a density $\alpha(s)$ such that:

$$\int \alpha(s) f_2(x - s)ds = f_1(x).$$

Because $f_i(x|t) = f_i(x - t)$,

$$\int \alpha(s) f_2(x - s|t)ds = \int \alpha(s) f_2(x - s - t)ds = f_1(x - t) = f_1(x|t) \quad \forall t, x.$$

Let $\xi_2(x)$ be a strictly concave increasing function such that

$$\int f_2(x - t_2)U(\xi_2(x))dx = V(t_2), \tag{8}$$

and for $t \neq t_2$

$$\int f_2(x - t)U(\xi_2(x))dx < V(t). \tag{9}$$

Under this scheme a Type 2 salesperson will pick effort level t_2 if she truthfully declares her type. Next we show that a Type 1 salesperson who falsely declares she is of Type 2 will not recover her cost of effort for any effort level.

If a Type 1 salesperson picks effort level t and declares she is a Type 2 salesperson, her expected payoff is: $\int f_1(x - t)U(\xi_2(x))dx = \int [\int \alpha(s) f_2(x - t - s)]U(\xi_2(x))dx$

$$= \int \alpha(s)[\int f_2(x - t)U(\xi(x + s))dx]ds. \tag{10}$$

$U(\xi_2(x))$ is strictly concave increasing in x because $U(\cdot)$ and $\xi_2(\cdot)$ are monotone increasing and strictly concave. Therefore:

$$\int \alpha(s) \int f_2(x - t)U(\xi(x + s))dxds < \int f_2(x - t)U(\xi(\int \alpha(s)(x + s)ds))dx. \tag{11}$$

$\int s\alpha(s)ds = 0$ because:

$$\int s\alpha(s)ds = \int \alpha(s)[\int x f_2(x - s)dx]ds$$

$$= \int x[\int \alpha(s) f_2(x - s)ds]ds = \int x f_1(x)dx = 0.$$

Combining equations (10) and (11) we get:

$$\int f_1(x-t)U(\xi_2(x))dx < \int f_2(x-t)U(\xi_2(x))dx \leq V(t) \quad \forall t. \tag{12}$$

Let $\xi_1(x)$ be the payment scheme when a salesperson declares her type to be Type 1 and the demand is x. We select a function that has the following properties. $\xi_1(x)$ is a monotone increasing function, $U(\xi_1(x))$ is convex increasing in x, and

$$\int f_1(x-t_1)U(\xi_1(x))dx = V(t_1), \tag{13}$$

and for $t \neq t_1$

$$\int f_1(x-t)U(\xi_1(x))dx < V(t), \tag{14}$$

These conditions and the properties of $\xi_2(\cdot)$ ensure that a Type 1 salesperson will declare her true type and pick effort level t_1. We are assured of the existence of such a function because $V(t)$ is a strictly increasing convex function and $f_i(x|t) = f_i(x-t)$. All that is left to show is that the Type 2 salesperson will prefer ξ_2 to ξ_1. By construction $U(\xi_1(x))$ is strictly convex. Hence, employing the approach used to show that a Type 1 salesperson will never recover her cost of effort we can show the following:

$$\int f_2(x-t)U(\xi_1(x))dx < V(t) \quad \forall t.$$

The direction of the inequality is reversed because we are replacing a concave function by a convex function. This completes the proof. □

References

Aliprantis, C. D., K. C. Border. 2007. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. 3rd ed. Springer, New York.

Basu, A. K., R. Lal, V. Srinivasan, R. Staelin. 1985. Salesforce compensation plans: An agency theoretic perspective. *Marketing Sci.* **4** 267–291.

Chen, F. 2000. Sales-force incentives and inventory management. *Manufacturing Service Oper. Management* **2**(2) 186–202.

Chen, F. 2005. Salesforce incentives, market information, and production/inventory planning. *Management Sci.* **51**(1) 60–75.

Coughlan, A. 1993. Salesforce compensation: A review of MS/OR advances. J. Eliashberg, G. Lilien, eds., *Handbooks in Operations Research and Management Science: Marketing*, vol. 5. North-Holland, Amsterdam, The Netherlands.

Dasu, S., H. Ochiumi. 2005. Incentives for forecasters. *University of Southern California* .

Dawid, A. P. 1982. The well-calibrated bayesian. *Journal of the American Statistical Association* **77**(379) 605–610.

DeGroot, M. H., S. E. Fienberg. 1983. The comparison and evaluation of forecasters. *The Statistician* **32**(1/2) 12–22.

Gonik, J. 1978. Tie salesmen's bonuses to their forecasts. *Harvard Business Review* .

Hendrickson, A. D., R. J. Buehler. 1971. Proper scores for probability forecasters. *The Annals of Mathematical Statistics* **42**(6).

Huff, D., P. J. Lederer, S. Misra, N. Rudi. 2003. Interdependencies between supply level choice and sales-force incentives. *New York University and the University of Rochester* .

Lal, R., S. Srinivasan. 1993. Compensation Plans for Single- and Multi-product Salesforces: An Application of the Holmstrom-Milgrom Model. *Management Sci.* **39**(7) 777–793.

Lal, R., R. Staelin. 1986. Salesforce compensation plans in environments with asymmetric information. *Marketing Sci.* **5** 179–198.

Porteus, E. 2002. *Foundations of Stochastic Inventory Theory*. 1st ed. Stanford Business Books.

Ridder, A., E. van der Laan, M. Salomon. 1998. How larger demand variability may lead to lower costs in the newsvendor problem. *Oper. Res.* **46**(6) 934–936.