Information Update and Risk Pooling in a Mixed Distribution Channel

Xiaohui Xu*
California State Polytechnic University, Pomona, CA, U.S.A.

Centralized distribution systems can pool risk across many markets as inventory can shift from low demand markets to high demand areas. On the other hand, a decentralized channel design allows store managers to control inventory according to area specific conditions. In this paper, we study a hybrid system that takes advantage of both types of setups. In addition to the considerations of risk pooling and local differences, we also examine situations where a centralized operation in a hybrid configuration can obtain information about general demand from decentralized local stores. This paper studies the optimal hybrid channel structure of a manufacturer facing a short selling season and long production lead time. The results demonstrate the tradeoff between information update and risk pooling. When the benefits of information flow outweigh the benefits from aggregating regional demands, the manufacturer chooses to have a significant local presence. We also study the manufacturer's decision when only pure strategies are available.

*Corresponding Author. E-mail address: xxu@csupomona.edu

I. INTRODUCTION

Companies that serve multiple markets can choose between centralized vs. decentralized inventory management. In a decentralized setting, local store managers create restocking purchase orders. In a centralized system, headquarters places orders to satisfy regional demands.

When making this key decision, risk pooling is an important factor to consider. Risk pooling effects refer to the fact that aggregate demands are more predictable than individual components. Levi, Kaminsky, and Simchi-Levi (2003) point out that this reduction in variability, measured by either standard deviation or coefficient of variation, allows a decrease in safety stock and therefore a reduction in average inventory carried.

The U.S. Army uses a centrally managed inventory system (Army Regulation 710–1, 2007). On the other hand, some retailers allow local store managers to place their orders based on their specific knowledge of local events and consumers. With the advent of e-commerce, more companies are operating a mixed distribution system. For example, Disney operates gift shops in its theme parks and several shopping malls. Disney also runs a direct online channel to reach markets where there is no Disney stores. Each local store has its own inventory and the store manager determines what to order. So, for local stores, Disney employs a decentralized model. However, for the online store, the company uses a centralized system where orders from across the nation are fulfilled through central warehouses. Sales data from local stores carries location-specific information and preferences that may be relevant to the market as a whole. So orders from Disney's local stores can help the centralized online store to fine tune demand forecasts.

We show in the paper that, for companies like Disney, it is optimal for them to run a hybrid system instead of a pure centralized or decentralized system. Disney's current system is superior to a pure online direct distribution setup or a pure brick-and-mortar, stores-everywhere solution.

Note that local stores often sell only a subset of products. Our study focuses on one particular product and how stocking decisions are made based on demand characteristics. Specifically, how many local stores should carry the product and what characteristics of a product's demand lead to stocking it in more local stores?

This study segments markets by location and customer preferences for buying online or from a local store. The manufacturer can have a local store serving a particular community and a centralized selling channel to directly serve a different set of customers in the same community.

In reality, customers may choose to buy from multiple local stores and an online customer also buys from local stores. We treat the same customer as different ones when s/he buys from different channel and we capture that in the demand function.

The paper is organized as follows:

In section II, we provide a summary of related literature. In section III, we specify our assumptions and set up a model assuming that the manufacturer tries to maximize the total profits by designing the right mixed channel. In this section, analyses are carried out and findings are presented in propositions and lemmas. In the final section, we relate our results to business practices and make recommendations for future research.

II. LITERATURE REVIEW

The role of the advance order and the subsequent information update in the supply chain has been studied. Fisher and Raman (1996) studied the benefits of inducing early customer orders to update the demand information by a retailer. Iyer and Bergen (1997) showed that advance ordering allocates inventory risk along a supply chain. Our model differs from Iyer and Bergen (1997) in that retailers can order only once and the updated information is applied in markets directly served by the manufacturer. Although inventory risk plays a role in our model, it is not the focus of the study. Our model also builds on Donohue (2000) and Tang, Rajaram, Alptekinoglu, and Ou (2004) to capture information updates in the supply chain.

III. Model

3.1. Model Assumptions

We consider a product with a short lifecycle, which is marketed to n markets with correlated demands. Due to a limited selling season and long production lead time, the manufacturer has only one chance to order before the season. We denote the unit cost as c and the selling price as r. See Fig. 1.

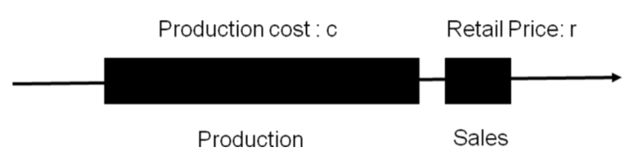

Figure 1. Model Assumptions.

We also assume that area demands are unique, but may be positively correlated. For each market, the manufacturer must determine whether to serve that market via centralized distribution or through a local store through which the manufacturer may acquire demand information hard to obtain in other ways. Our assumption is that the manufacturer will be able to observe local information only if he maintains a physical presence in the market by operating a retail facility that is responsible for holding its own inventory. This is reasonable, since holding the local manager accountable for her own inventory provides strong incentives for her to put in the effort necessary to observe the demand accurately and make appropriate stocking decisions.

Since the selling season is short relative to the production lead time, all production must occur prior to the realization of demand. Prior to production, the manufacturer determines which markets to serve directly and which to serve through local stores. In our model, this decision reduces to the determination of a number, denoted by t, of markets to serve through physical stores. See Fig 2: Channel Structure.

After the manufacturer determines the number of markets, t, to serve through local retail outlets, each local store observes its own local information, represented by y_i, and places an order, denoted by $q_i(y_i)$, with the manufacturer. The manufacturer responds by producing enough product to satisfy each local store's order plus an additional amount that will be used to satisfy demand in the $n - t$ markets that he serves directly. We represent the additional amount that the manufacturer produces for the $n - t$ markets that he serves directly by $Qm\ (t, y)$, where $y = [y_1, ..., y_t]$.

Finally, after all physical stores receive the quantities that they have ordered, demand is realized. We do not allow for any excess inventory held by the manufacturer to be used to satisfy demand at any local store, nor do we allow for any inventory held by a physical store to be used to satisfy demand in any market other than its own. This "no-sharing" policy is observed at Target, Wal-Mart, Victoria's Secret, and many other retailers operating a mixed channel.

The demand for market i is represented by $d_i = y_i + \varepsilon_i$. The random variable, y_i, which is normally distributed with the mean θ and standard deviation σ, represents information that is available prior to the point at which production must occur, but is observable only to a local store because of its physical presence in the market i.

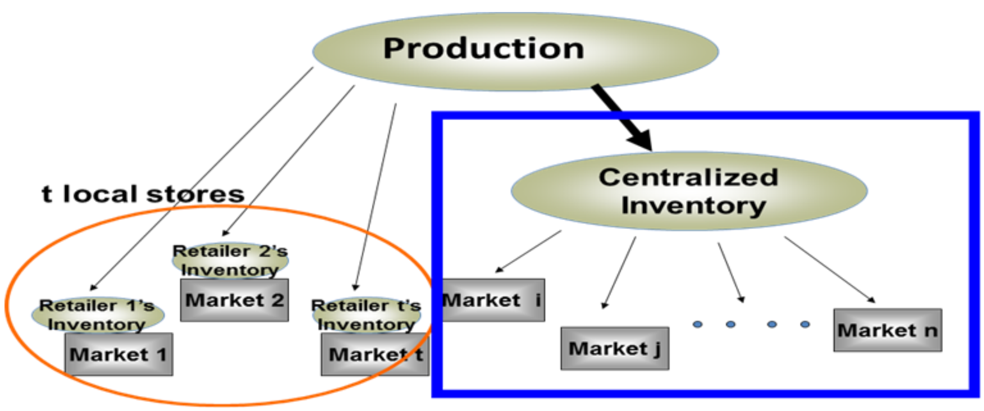

Figure 2. Channel Structure.

The manufacturer has a prior distribution for the mean θ of the distribution for y_i, specifically θ follows $N\ (\mu,\ v_2)$. Note that, although the conditional distributions for the y_is are i.i.d., the d_is are not. However, because of their dependence on θ, the demands in different markets are correlated. Note that we assume that the y_is are i.i.d., which implies that the manufacturer is unable to discern any differences among various markets unless he has access to the local information contained in the realization of y_i. In reality, a manufacturer is likely to have at least some information about differences among individual markets. For example, a manufacturer could reasonably expect that sales will be stochastically larger in Chicago than in Debuke. We assume that the y_is are i.i.d. for ease of exposition. However, by including a market specific shift parameter in the conditional distribution for each y_i, we could easily account for the manufacturer to be partially aware of the differences among markets.

An important feature of our model of demand is that local information about one market conveys something about general market conditions. Specifically, the observation of local information y_i in market i signals something about the mean, θ, of the distribution for demand in all markets $j = 1, ..., n$. See Fig. 3 for the information update process.

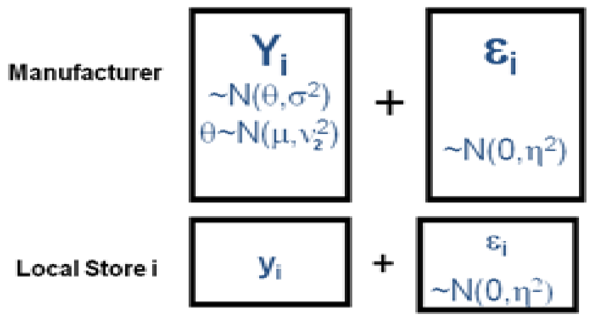

Figure 3. Prior Market Demand Information.

3.2. Model Analysis

As described above, each local store observes the realization of y_i prior to placing her order with the manufacturer. Note that, throughout the paper, we use $\emptyset(\cdot)$ and $\Phi(\cdot)$ for probability density function and cumulative distribution function for standard normal distribution, respectively. Based on the normal distribution's property and analysis, it is easy to see that, after observing y_i,, the conditional distribution of local store i's demand will be normal with mean y_i and standard deviation η. It follows that retailer i's optimal order quantity will be

$$q_i(y_i) = y_i + \eta\Phi^{-1}((r-c)/r). \qquad (1)$$

In our model, since the manufacturer has the same information as local store i about everything except y_i, he is able to infer y_i from each market in which a local store operates. In reality, uncertainty about local demands might prevent the manufacturer from perfectly inferring this value. However, although this would tend to diminish the information value of the local store's order, it would not eliminate it. The local information can be used to the manufacturer's advantage in serving the $n - t$ markets that he serves directly. Recall that the demands in the various markets are correlated through the mean θ of the distribution for y_i. Since each observed value of y_i can be used to update the manufacturer's prior for θ, it follows that, as the manufacturer obtains more orders from local stores, he will have better information about the demand in the remaining markets that he serves directly.

As discussed above, the manufacturer can infer the information about market i from the order placed from a local store in that market. From (1), it is easy to see that the manufacturer can infer the value of y_i as follows:

$$y_i = q_i(y_i) - \eta\Phi^{-1}((r\text{-}c)/r).$$

Recall that, before receiving orders from the local stores, the manufacturer's prior distribution for the mean θ of the market information y_i is $N(\mu, v_2)$. After obtaining orders from local stores operating in t of the markets, the posterior distribution of θ is a normal distribution for which the mean μ_1 and standard deviation v_1 are as follows:

$$\mu_1 = (\sigma^2\mu + v^2(y_1+y_2+...+y_t) / (\sigma^2 + tv^2), \qquad (2)$$

$$v_1^2 = (\sigma^2 v^2) / (\sigma^2 + tv^2) = (1/v^2+t/\sigma^2)^{-1}. \qquad (3)$$

Let D_{n-t} be the random variable corresponding to the total demand that the manufacturer will experience from the $n - t$ markets he serves directly.

$$D_{n-t} = \sum_{t+1}^{n}(y_i + \varepsilon_i). \qquad (4)$$

By assumption, the ε_is are i.i.d. random variables that are independent of the y_is. It follows that:

$$E[D_{n-t}] = (n-t)\mu_1, \qquad (5)$$

$$\mathrm{Var}\,[D_{n-t}] = \mathrm{Var}\left[\sum_{t+1}^{n} y_i\right] + \mathrm{Var}\,[\sum_{t+1}^{n}(\varepsilon_i)] = (n-t)(\sigma^2+\eta^2) + (n-t)^2 v_1^2 \quad (6)$$

Note that both expressions are conditional on the realization of $y = [y_1, ..., y_t]$. However, to avoid overly cumbersome notation, we have omitted this conditioning from the notation.

Given that the local store operating in market i orders the optimal quantity in (1), the conditional expected sales in this market will be

$$s_{Li}(y_i) = y_i + L(z_m)\,\eta. \quad (7)$$

where $L()$ is the expected unit normal loss function and $z_m = \Phi^{-1}[(r-c)/r]$. Recall that, at the time that the manufacturer decides to operate a local retail operation in market i, he does not know the realization of y_i. Taking the expectation of (7), we have $E[s_{Li}(y_i)] = \mu + L(zm)\ \eta$. It follows that the expected profits that the manufacturer will earn from operating a physical store in a local market are

$$E_y[\pi_{mL}] = E_y[-cq_i(y_i) + rs_{Li}(y_i)] = (r-c)\mu - \eta(rL(z_m) + cz_m). \quad (8)$$

The manufacturer will produce enough to fill all orders from the t markets in which he has decentralized operations, plus an additional amount to fill the demand from markets that he intends to serve through centralized direct sales.

Since local retail operations provide information to the manufacturer about the value of θ, the conditional distribution of the aggregate demand in the markets served centralized will be $N\ (E[D_{n-t}\ ],\ Var[D_{n-t}\ ])$ (see (5) and (6) respectively).

It follows that the optimal quantity for the manufacturer to produce for the $n - t$ markets that he will serve directly is:

$$Qm(y,t) = (n-t)\mu_1 + z_m\sqrt{(n-t)(\sigma^2+\eta^2) + (n-t)^2 v_1^2} \quad (9)$$

Let $S_C(y, t)$ be the total expected sales across all $n-t$ centrally served markets when the manufacturer produces $Qm(y, t)$ as shown above. Then we have:

$$Sc(y,t) = (n-t)\mu_1 + L(z_m)\sqrt{(n-t)(\sigma^2+\eta^2) + (n-t)^2 v_1^2} \quad (10)$$

where $L()$ is the expected unit-normal loss function.

In our analysis, it will occasionally be useful to refer to the average profit per market served from the central sales operation, which we will denote by lower case $s_C(y, t)$. By dividing (10) by $n - t$, it is easy to see that:

$$s_c(y,t) = \mu_1 - L(z_m)\sqrt{\frac{(\sigma^2+\eta^2)}{n-t} + v_1^2}$$

Thus, given that the manufacturer has access to the realization of the local information, $y = [y_1, y_2, ...y_t]$, in markets $1, 2, ..., t$, his optimal total profit from centralized sales to the remaining $n - t$ markets can be expressed as follows:

$$\Pi_{mC}(y,t) = rS_c(y,t) - cQ_m(y,t) \quad (11)$$

We can now substitute (9) and (10) into (11). After using (5) and (6), and taking expectation with respect to $y = [y_1, \cdots, y_t]$, we have:

$$E_y[\Pi_{mC}(y,t)] = (n-t)[(r-c)\mu - (rL(z_m) + cz_m)\sqrt{\frac{(\sigma^2+\eta^2)}{n-t} + \frac{\sigma^2 v^2}{\sigma^2 + tv^2}}] \quad (12)$$

where we have used the fact that $E_y[\mu_1] = \mu$. Let us use lower case $\pi_{mC}(y, t)$ to represent the average profit earned by the manufacturer in the markets that he serves directly. By dividing (12) by $n - t$, it is easy to see that:

$$E_y[\pi_{mC}(y,t)] = (r - c)\mu - (rL(z_m) + cz_m)\sqrt{\frac{(\sigma^2 + \eta^2)}{n - t} + \frac{\sigma^2 v^2}{\sigma^2 + t v^2}} \quad (13)$$

If the manufacturer must choose between selling entirely through a centralized operation vs. entirely through local stores, the following proposition shows the optimal strategy. Note that no information update is involved in any pure strategy.

Proposition 1 If and only if $(n-1)\eta^2 > \sigma^2$, the manufacturer will prefer to use pure centralized system, i.e. t = 0, over pure local retail sales, t = n. Otherwise, if this inequality is reversed, then the manufacturer will prefer pure local retail sales over pure centralized distribution.

The above proposition shows that the manufacturer will prefer direct centralized sales when the amount of residual uncertainty (η) after observing local information is high. In this case, risk pooling is particularly beneficial. On the other hand, when a large portion of uncertainty can be resolved by observing y_i (σ is relatively large), the local information dominates risk pooling.

Proposition 1 also indicates that, when the number of markets is large, risk pooling effects are most significant and centralized direct sales exploit these effects.

Proposition 2 *a)* There is a threshold value t**< n, such that the expected profit per market in markets served through centralized direct sales is increasing in t when t < t** and is decreasing when t > t**.
b) t** > 0 if and only if $n > \frac{\sigma}{v^2}\sqrt{\sigma^2 + \eta^2}$

This proposition illustrates the tradeoff between the benefit of early information and risk-pooling effects. The marginal benefit from more accurate demand information decreases with *t*. On the other hand, the loss due to fewer and fewer markets in the centralized mode pool gets larger and larger.

If *n* is relatively large, such that at the beginning, the loss from risk pooling is marginal, but the benefit of the information is large, we have incentives to set up local stores in at least some markets and $t^{**} = \frac{n\sigma v^2 - \sigma^2\sqrt{\sigma^2+\eta^2}}{v^2(\sigma+\sqrt{\sigma^2+\eta^2})} > 0$. If *n* is small, the loss from risk pooling dominates at the beginning, and *t** < 0.* The manufacturer sells centralized to all markets in this context.

We also observe this in reality. When a business starts and serves a limited number of markets, it usually chooses to sell through the Internet using centralized inventory. Later, as the business builds up brand reorganization and has a larger *n*, the online store will serve more markets and the manufacturer has incentives to set up local presence(s).

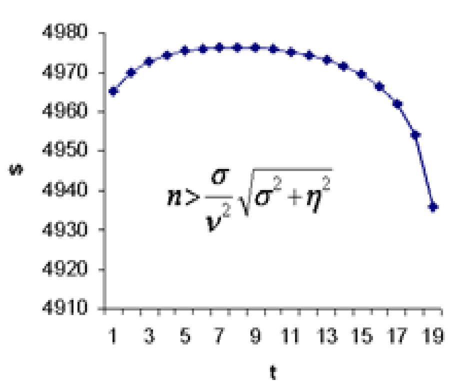

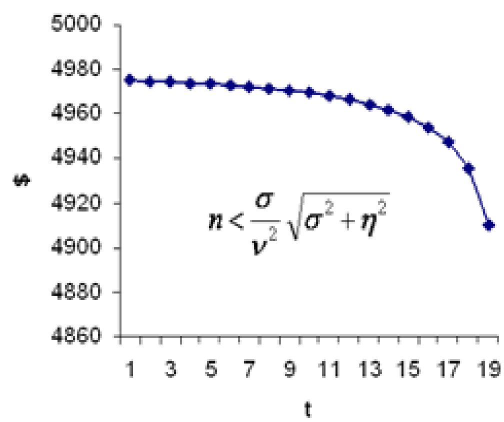

Figure 4. Per Market Profit in Centralized Channel.

Finally, note that t^{**} is always less than n. Intuitively, the early demand information from local stores benefits the manufacturer only when the manufacturer needs to make quantity decisions himself for the centralized areas. If $t = n$, the quantity decisions in all markets are solely the responsibility of the retailers and the manufacturer never needs to know anything about demand.

The above proposition shows the per market profits for centralized part of the channel and t^{**} is the number that optimizes the direct operations. We must look at overall profits to derive the global optimization. We denote the optimal number of local stores in the mixed channel by t^*, which leads to maximum total profits.

Proposition 3. The optimal number of local stores t*<n solves the equation

$$\left(rL(z^0)+cz^0\right)\left(-\frac{d}{dt}\left(\sqrt{(n-t)(\sigma^2+\eta^2)+(n-t)^2v_1^2}\right)-\eta\right)=0$$

In general, if

$$\sqrt{n(\sigma^2+\eta^2)+n^2v^2}-\eta-\sqrt{(n-1)(\sigma^2+\eta^2)+(n-1)^2\frac{\sigma^2v^2}{\sigma^2+v^2}}>0$$

then $t^*>0$.

Lemma 1. If the manufacturer already chooses to serve n−1 markets through local stores, he chooses to do the same for the last market.
This is because the benefit of risk pooling is limited when there is only one regional market left.
Lemma *2. t*>0* when
1). n is large, or
2). v^2 *is large, or*
3). n and v^2 *is at least moderate* ($\sqrt{n\eta^2+n^2}-\eta-\eta\sqrt{(n-1)}>0$), *and* σ *is small.*

From lemma 2, when n is large, the marginal loss from risk pooling is small and the benefits associated with knowing more about demand dominate the trade-off.

Lemma 2 also shows that when v_2 is large, the benefit of early information is evident and $t^*>0$. If v_2 is small, we don't get much value from the local information and we do not maintain decentralized channels. We can also observe the following properties of the optimal solution by numerical study.

Observation 1: The optimal number of local stores decreases in the residual unresolved uncertainty (η^2). See Fig 5.

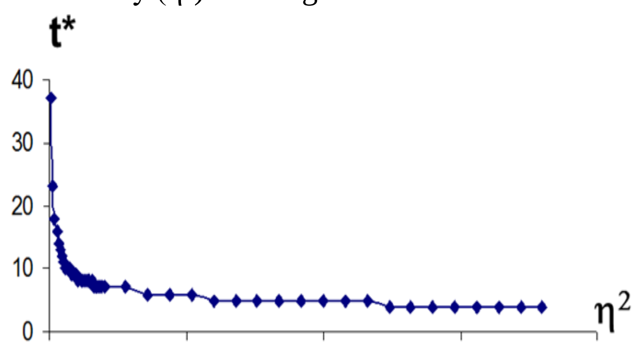

Figure 5. t* Decreases in η^2

Observation 2: The optimal number of indirect markets (t*) is non-decreasing in the variance of the mean of the demand distribution (v^2), which the manufacturer updates after obtaining local demand information. See Fig. 6.

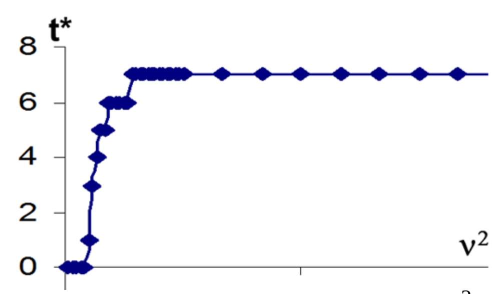

Figure 6. t* Non-Decreasing in v^2

Observation 3: The optimal number of local stores (t*) is increasing in the dispersion of local market demands *(σ^2)* See Fig. 7.

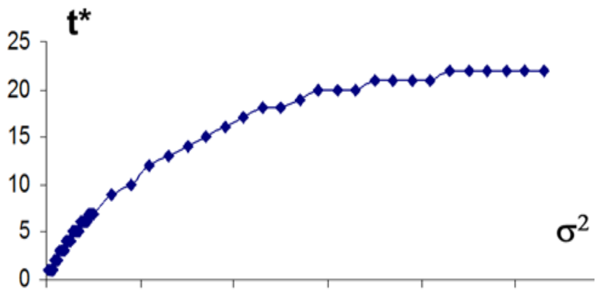

Figure 7: t* Increasing in σ^2

IV. CONCLUDING REMARKS

Our model studies the tradeoff between information updates and risk pooling in a hybrid channel. We found that the more markets a manufacturer serves, the more likely he will use a hybrid channel. The more value the information has, the more markets the manufacturer may serve through local stores. The number of local stores decreases if the residual uncertainty is large, i.e. the information update can only remove a small portion of uncertainty. Also the number of local stores increases if the dispersion of local markets is large, i.e., markets are quite different in nature. This is because the manufacturer is better off allowing local managers with regional specific information to determine stocking quantities. As more companies take advantage of the Internet and new media to operate a hybrid channel, our research can be important to understand channel decisions.

Note that we do not assume any costs to set up local operations. This is because many businesses are moving from pure brick-and-mortar to hybrid channels. These businesses have physical stores already and do not need to invest to have a local presence. Furthermore, the focus of the paper is to study the tradeoffs between information updates and risk-pooling effects. Since setup costs to adjust a physical-store network can be substantial, we recommend incorporating the setup costs in future research. Also, we can expand the research by assuming that customers can buy from either channel. Finally, empirical research is highly recommended to understand the information-update and risk-pooling tradeoffs in the business world.

V. REFERENCES

Army Regulation 710–1, Centralized Inventory Management of the Army Supply System, Headquarters, Department of the Army,Washington, DC, 20 September 2007.

Donohue, K. L., "Efficient Supply Contracts for Fashion goods with forecast updating and Two production modes," *Management Science*, 46 (11), 2000, 1397-1411.

Fisher, M. and Raman, A., "Reducing the Cost of Demand Uncertainty through Accurate Response to Early Sales," *Operations Research*, 44(1), 1996, 87-99.

Iyer, A. and Bergen, M., "Quick Response in Manufacturer-Retailer Channels," *Management Science*, 43(4), 1997, 559-570.

Levi, D. S., Kaminsky, P., and Simchi-Levi, E., *Designing & Managing the Supply Chain* (2nd ed.), The McGraw-Hill Companies, New York, 2003.

Tang, C. S., Rajaram, K., Alptekinoglu, A and Ou, J., "The Benefits of Advance Booking Discount Programs: Model and Analysis," *Management Science*, 50(4), 2004, 465-478.

APPENDIX. Proofs for propositions and lemmas

Proof for Proposition 1:

Proposition 1 is about pure strategies.

When the manufacturer uses a pure decentralized strategy, a local store's expected profit is shown in equation (8).

When the manufacturer deploys a pure centralized strategy, the per-market expected profit is shown in equation (13).

By comparing (8) and (13), we get the result in proposition 1.

Proof for Proposition 2:

a) The first-order derivative of the manufacturer's expected profit per market served directly is:

$$-\left(\frac{rL(z_m)+c\cdot z_m}{2}\right)\left(\frac{\frac{\sigma^2+\eta^2}{(n-t)^2}-\frac{\sigma^2 v^2}{\left(\sigma^2+t\cdot v^2\right)^2}}{\sqrt{\frac{\sigma^2+\eta^2}{n-t}+\frac{\sigma^2 v^2}{\sigma^2+t\cdot v^2}}}\right) \tag{14}$$

Note that the first term is always positive and that, once the latter term turns positive, it remains positive for all larger values of *t*. To see that the first term is positive, note that it is obviously positive for $z_m \geq 0$ since the unit normal loss function, *L()* is non-negative by definition. For $z_m < 0$, we can use the definition of the expected unit normal loss function, to obtain:

$$\begin{aligned} L(z_m)+z_m &= \int_{z_m}^{\infty} r(x-z_m)d\Phi(x)+cz_m \geq \int_{z_m}^{\infty} rxd\Phi(x)+\int_{-\infty}^{z_m} rz_m d\Phi(x) \\ &= \int_{-\infty}^{z_m} -rxd\Phi(x)+\int_{-\infty}^{z_m} rz_m d\Phi(x)=\int_{-\infty}^{z_m} r(z_m-x)d\Phi(x)>0 \end{aligned}$$

Where the inequality follows from $z_m < 0$, and the assumption that $r > c$. It follows that (14) has the same sign as $\sigma^2 v^4 (n-t)^2 - (\sigma^2 + \eta^2)(\sigma^2 + t \cdot v^2)^2$, which is decreasing in *t* for $t \leq n$.

b). As shown in part a), the sign of (14) has the same sign as $\sigma^2 v^4 (n-t)^2 - (\sigma^2 + \eta^2)(\sigma^2 + t \cdot v^2)^2$. It is easy to confirm that this will be positive for *t=0* if and only if $n > \frac{\sigma}{v}\sqrt{\sigma^2 + \eta^2}$.

Proof for Proposition 3:

Proposition 3 is simply the first order condition. Only the sufficient condition for $t^* > 0$ is specified in the proposition. The proof is straightforward and omitted here.

Proof for Lemma 1:

Proof outline:

Compare total profits when *t=n* and when *t=n-1*. The difference is

$$\left(rL(z^0) + cz^0\right)\left(\eta - \sqrt{\sigma^2 + \eta^2 + v_1^2}\right) < 0.$$

Proof for Lemma 2:

Lemma 2 is directly derived from part b) of proposition 3.