Multilocation Combined Pricing and Inventory Control

Awi Federgruen • Aliza Heching
Graduate School of Business, Columbia University, New York, New York 10027
IBM T.J. Watson Research Center, Yorktown Heights, New York 10598
af7@columbia.edu • ahechi@us.ibm.com

We consider the problem of managing inventories and dynamically adjusting retailer prices in distribution systems with geographically dispersed retailers. More specifically, we analyze the following single item, periodic review model. The distribution of demand in each period, at a given retailer, depends on the item's price according to a stochastic demand function. These stochastic demand functions may vary by retailer and by period. The replenishment process consists of two phases: In some or all periods, a distribution center may place an order with an outside supplier. This order arrives at the distribution center after an ''order leadtime'' and is then, in the second phase, allocated to the retailers. Allocations arrive after a second ''allocation leadtime.''

We develop an approximate model that is tractable and in which an optimal policy of simple structure exists. The approximate model thus provides analytically computable approximations for systemwide profits and other performance measures. Moreover, the approximate model allows us to prove how various components of the optimal strategy (i.e., prices and order-up-to levels) respond to shifts in the model parameters, e.g., to shifts in the retailers' demand functions. In addition, we develop combined pricing, ordering, and allocation strategies and show that the system's performance under these strategies is well gauged by the above approximations. We use this model to assess the impact of different types of geographic dispersion on systems with dynamically varying prices and how different system parameters (e.g., leadtimes, coefficients of variation of individual retailers' demand, price elasticities) contribute to this impact. Similarly, we use the model to gauge the benefits of coordinated replenishments under dynamic pricing, and how these benefits increase as the allocation decisions of the systemwide orders to individual retailers are postponed to a later point in the overall replenishment leadtime.

We report on a comprehensive numerical study based on data obtained from a nationwide department store chain.
(*Inventory; Pricing Strategies; Two-Echelon System*)

1. Introduction

We consider the problem of managing inventories and dynamically adjusting retailer prices in distribution systems with geographically dispersed sales locations (''retailers''). Inventories can be controlled via the supply process, the demand process, or a combination thereof. The *supply side* can be managed by an effective replenishment strategy, placing systemwide orders and determining retailer allocations in proper amounts and at the proper points in time. The *demand* process can be controlled by stimulating (dampening) demand via price decreases (increases) or by the use of advertising and coupon promotions.

Simultaneous and dynamic control of the demand and the supply side, i.e., the integration of inventory control, in the classical sense of the word, with that

1523-4614/02/0404/0275$05.00
1526-5498 electronic ISSN

of revenue management, has been addressed by few, and only in single-location, single-item settings. Yet, Federgruen and Heching (1999) show in this single-location context that the optimal integrated price and inventory control policy can result in major profit increases—by 6.5% for a specialty retailer of high-end (private label) women's apparel—as compared to a two-stage procedure in which a price trajectory is determined *first,* followed by a dynamic inventory policy to optimally respond to the resulting sequence of demand distributions. These differences relate to instances with weekly replenishments, while even larger profit increases are observed when only a few replenishments are permitted over the course of the sales season. In the retail sector such differences may have large impacts on the bottom line.[1] Integrated price and inventory control allows for *contingent* pricing, i.e., the price may be varied as a function of the prevailing inventory (or inventories) in the system. Thus, prices may be reduced to stimulate demand when the firm is faced with large inventories, or they may be increased to dampen demand when faced with low inventories or backlogs.

The dichotomy between supply- and demand-side analytical models reflects traditional organizational structures, where supply-side decisions are relegated to production or operations units and demand-side decisions to marketing departments. Decisions in both are made in isolation or, at best, in a *sequential* manner, as when prices are selected by the marketing unit (without proper recognition of their operational consequences), and subsequently transmitted to the operational unit. Even when a single unit or individual makes both the pricing and operational decisions, the trade-offs between these two areas are, as of yet, poorly understood.

We analyze the following single item, periodic review model. The distribution of demand in each period, at a given retailer, depends on the item's price according to a retailer- and period-specific stochastic demand function. The demands faced by different retailers in a given period may be correlated, but demands are independent across time. (Our numerical studies confine themselves to the case of independent retailer demands.) The price may be specified dynamically as a function of the state of the system. The company thus acts in a market with imperfect competition; for example, the company may be a monopolist or the market may allow for product differentiation. The analysis lays the groundwork for more complex game-theoretical models that can be used to address markets with perfect or limited competition.

An important assumption is that the retailers adopt a *common* price at any given point in time. This is the practice of most retail chains over large geographic regions (e.g., the Northeast, the West Coast, etc.); cf. Kane (2000), including the nationwide department store chain on which the numerical studies in this paper are based. Common pricing is often driven by the desire to conduct national or regional advertising campaigns or to distribute national or regional catalogs. Another factor is that customers tend to react most adversely to price differences between sales outlets in the same region. In a recent experiment, Amazon.com charged different prices for different customers for some of its DVD movies; the price charged depended on which browser the customer used. Several customers, upset by what they felt was unfair treatment, filed complaints with the FTC and the Better Business Bureau, accusing the company of ''dishonesty'' and ''sneakiness'' (see Wolverton 2000 and Newman 2000). Amazon.com stopped the practice in response to customer complaints.[2] Another factor cited by executives is the difficulty to maintain location-dependent prices in their accounting systems, in particular when purchases may be returned to any of the sales locations. Finally, various papers, e.g., Thisse and Vives (1988), have shown that when multiple retail chains compete, common pricing often arises as a Nash equilibrium. In spite of these considerations,

[1]For a particular ski-wear manufacturer Fisher and Raman (1996) report a cost reduction in the amount of 1% of sales results in a 60% increase in bottom-line profits. In §3, we show similar potential for profit increases in multilocation settings.

[2]May, a Jupiter analyst, states: ''This was a major moment in e-commerce, not only because it was the *first* widespread test of a powerful pricing tool, but because Amazon backed off so quickly and so definitively'' (see *Washington Post* 2000). In the words of the *Washington Post* (2000): ''Few things stir up a consumer revolt quicker than the notion that someone else is getting a better deal.''

there are nevertheless cases where price *differentiation* may be acceptable, allowing for even higher profits. The associated dynamic control problem is significantly more complex than the one addressed in this paper. It is unknown, to date, via what type of strategies the potential for price differentiation can be exploited in an effective manner.

The replenishment process consists of two phases. At the beginning of some or all of the periods the distribution center may place an order, of possibly limited capacity, with an outside supplier. (Our numerical studies confine themselves to the case of uncapacitated orders.) The order arrives at the distribution center after a first ''*order leadtime*'' and is allocated to the retailers, requiring a second retailer-specific ''*allocation leadtime.*'' The distribution center thus acts as a transshipment or cross-docking point only, where no inventory is kept, or it may represent a coordination function only, without any specific physical location. Many distribution systems employ cross-docking to reduce the total replenishment leadtime and to avoid costly, time consuming, and error-prone storage and retrieval activities in intermediate warehouses.[3] The cross-docking practice has also been implemented by the nationwide department store chain whose data are used in our case studies. (Goods reside no more than 48 hours in its distribution centers.) On the other hand, where applicable, our analysis can be extended along the lines of that in Federgruen and Zipkin (1984b) and Aviv and Federgruen (1999) to allow for central inventories. Where storage and retrieval in the distribution center can be undertaken rapidly and inexpensively, the storage option further strengthens the profit potential of coordinated replenishments. The literature, admittedly confined to settings with prespecified price and demand distributions, however, suggests that the benefits of central inventories are relatively modest. (This may explain why cross-docking is so prevalent.) See, in particular, Svoronos and Zipkin (1988) and Axsater (1993) who deal with a one-warehouse, multiretailer system similar to ours, and Gallego and Zipkin's (1999) who address serial systems.

[3]Walmart's developing into the world's largest and most profitable retailer has been attributed to a number of logistical practices among which cross-docking is prominent (see Stalk et al. 1992).

Stockouts at each retailer are backlogged. Ordering and allocation costs are proportional with the order and allocation sizes. Allocation costs may be retailer specific. Inventory carrying and stockout costs depend on end-of-the-period inventory levels according to retailer-specific rates.

We pursue several objectives. *First,* we develop effective combined pricing, ordering, and allocation strategies. These are obtained from specific modifications of the optimal policy of an appropriately designed *approximate* model that is tractable (as opposed to the exact model). Its optimal solution provides analytically computable approximations for systemwide profits and other performance measures. Our *second* objective is to show that these approximate measures are very close to the *exact* ones under our proposed strategies (computed via simulations). The analytical and relatively efficient solution of the approximate model can thus be used effectively in various design studies to answer a variety of strategic questions, which is our *third* objective. The strategic questions of interest include: (i) the benefits of coordinated replenishments under dynamic pricing and how these benefits vary as the allocation decisions to the retailers can be postponed to a later point in the overall replenishment leadtime, or as the degree of heterogeneity among the retailers is increased and (ii) the impact on various performance measures of several system parameters, e.g., leadtimes, coefficients of variation of individual retailer's demand, price elasticities, etc.

While the design of our combined pricing and inventory strategies as well as the specification and solution of the approximate model build on techniques employed in simpler multiechelon inventory models, we face several *novel* challenges intrinsic to our environment with dynamically controllable prices. For example, in standard models with *exogenously* specified demand distributions, it is easy to handle leadtimes by using inventory *positions* (=inventory levels + orders in process) as opposed to inventory *levels* as the state variables. A standard accounting device is used, whereby each period is charged with the *expected* inventory costs one leadtime later. With dynamically controllable prices, these future costs depend on price decisions in *future* periods, and are therefore unspecified. As a second example, in standard models it has been shown that a close-to-optimal strategy can be designed using the very *ordering* strategy that is optimal in the approximate analytical model, without incorporating *any* adjustments. This strategy specifies orders on the basis of the *aggregate* systemwide inventory position only. Under dynamically controllable prices, increased imbalance in the inventory levels of the retailers can be expected, requiring adjustments of systemwide orders as a function of the prevailing imbalances.

The following is a brief review of the relevant literature; see Federgruen and Heching (2000) for details. The majority of the inventory literature assumes that the demand processes and pricing decisions are exogenously determined. The seminal paper here is Eppen and Schrage (1981). See Federgruen (1993) for a survey and Chen and Zheng (1994) and Aviv and Federgruen (1999) for more recent contributions. Most of the dynamic pricing literature addresses single-item, single-location systems. See Bitran and Mondschein (1993, 1997) and Gallego and van Ryzin (1994, 1997) and its references for settings where a *single* order covers the selling season. See Rothschild (1994), Grossman et al. (1977), McLennan (1984), Balvers and Cosimano (1990), and Braden and Oren (1994) for models in which no inventory can be carried from one period to the next. We refer to Federgruen and Heching (1999) and its references for a discussion of the more general case where inventories may be carried and replenishment orders placed—with zero leadtime—in some or all periods. Chen et al. (2001a,b) appear to be the first to consider integrated pricing and inventory control in a multilocation setting; however, they assume that demands occur at a *constant deterministic* rate that depends on the retail price charged.

The remainder of this paper is organized as follows. In §2 we introduce the basic model. In §3 we discuss some case studies, highlighting the benefits of combined pricing and inventory control and those of coordinated replenishments. In §4, we give a brief analysis of the deterministic version of the model. We use the results of this deterministic model to design the approximations and heuristics for the stochastic model. In §5 we develop a tractable model as an approximation to the original system. In §6 we derive structural properties of the optimal policy for this approximate model. Based on the approximate analytical model, in §7 we develop efficient heuristic strategies for the initial model. Finally, in §8 we describe a second numerical study with several objectives. First, it gauges the effectiveness of the proposed heuristics and the accuracy of the approximate model. Second, we investigate how different parameters impact various performance measures. Third, the study provides insights into several important strategic questions.

2. Model and Notation

We first specify the model and introduce the basic notation. The distribution system consists of a distribution center, via which J retailers, indexed $j = 1, \ldots, J$, are replenished. The planning horizon has $T < \infty$ periods, indexing each period as the number of periods remaining until the end of the horizon. Extensions to infinite horizon settings, with total discounted or long-run average profits as the objective, can be established along the lines of the single-location case in Federgruen and Heching (1999).

At the beginning of each period, several simultaneous decisions need to be made: (i) the size of a new replenishment order (if any) to be placed by the distribution center, (ii) the price to be charged in the current period, and (iii) the allocation to the retailers of any order arriving at the distribution center. (Recall, the distribution center carries no inventories.) The order leadtime consists of L periods, and l_j ($j = 1, \ldots, J$) is the allocation leadtime from the distribution center to retailer j. Demand at retailer j in period t is described by the following stochastic demand function:

$$d_{jt}(p_t) = [\gamma_{jt} + \delta_{jt}p_t]\epsilon_{jt} + [\eta_{jt} + \theta_{jt}p_t], \quad j = 1, \ldots, J, \quad t = 1, \ldots, T, \quad (1)$$

where p_t denotes the price charged in period t and the vector $\epsilon^t = \{\epsilon_{jt} : j = 1, \ldots, J\}$ has a general multivariate distribution with finite mean, allowing for arbitrary correlations between the retailers. For example, positive correlation may be explained by the weather, or general economic conditions, e.g., the consumer confidence index, increasing or decreasing demand at all retailers in that region; negative correlation may be due to road construction near one retailer, diverting traffic to a nearby store, or by store-specific advertising. $\{\epsilon^t : t = 1, \ldots, T\}$, and hence demands in consecutive periods, are independent. This specification encompasses two frequently considered special cases: (i) the additive model where $\gamma_{jt} = 1$ and $\delta_{jt} = 0$, i.e., $d_{jt}(p_t) = \eta_{jt} + \theta_{jt}p_t + \epsilon_{jt,}$ and (ii) the multiplicative model where $\eta_{jt} = \theta_{jt} = 0$, i.e., $d_{jt}(p_t) = [\gamma_{jt} + \delta_{jt}p_t]\epsilon_{jt}$. Extensions to more general demand functions can be made, as in the single-retailer case; see Federgruen and Heching (1999), §1. We assume that $\gamma_{jt}, \eta_{jt} \geq 0$ and $\delta_{jt}, \theta_{jt} \leq 0$, reflecting downward-sloping demand curves.

The cost structure is described by the following parameters:

c_{0t} = the ordering cost per unit ordered by the distribution center in period t ($t = 1, \ldots, T$);
c_{jt} = the shipment cost per unit allocated by the distribution center to retailer j in period t ($j = 1, \ldots, J$; $t = 1, \ldots, T$);
h^+_{jt} (h^-_{jt}) = the carrying (backlogging) cost per unit stored (backlogged) at retailer j at the end of period t ($j = 1, \ldots, J$; $t = 1, \ldots, T$).

The objective is to maximize expected infinite horizon profits, discounted with a factor $\alpha \leq 1$. Let b_t = the order capacity in period t. General capacity limits $\{b_t\}$ allow us to model limited ordering opportunities or backup agreements (i.e., $b_t > 0$ for only a few periods), with only a few, perhaps two, possible ordering periods and the supplier reserving a given quantity for the second or later ordering opportunity; see, e.g., Eppen and Iyer (1997) and Fisher and Raman (1996). The price, p_t, is selected from an interval $[p_{t,\min}, p_{t,\max}]$, $t = 1, \ldots, T$. In some cases only markdowns are permitted. Such price restrictions require minor extensions.

At the beginning of a period, the distribution center determines an order size. Simultaneously, any order placed L periods earlier arrives and is allocated among the J retailers. Immediately thereafter, each retailer j receives any shipment allocated to him l_j periods earlier. Demands in a period occur at the end of the period. Let:

x_{jt} = the inventory level at retailer j at the beginning of period t, before arrival of shipments from the distribution center, $j = 1, \ldots, J$; $t = 1, \ldots, T$;
z_{jt} = the allocation to retailer j assigned at the beginning of period t, $j = 1, \ldots, J$; $t = 1, \ldots, T$;
w_t = the order placed by the distribution center at the beginning of period t, $t = 1, \ldots, T$.

3. Benefits of Integrated Planning and Coordinated Ordering

In this section, we demonstrate the benefits of integrated price and inventory control and those associated with coordinated replenishments, with the help of a numerical study based on data from a national retailer of basic and fashion items. The retailer operates several divisions, each composed of a chain of stores. The divisions are located throughout the United States. Each chain uses a common price at any point of time. Procurements are coordinated via systemwide orders placed with outside vendors. Orders arrive at regional distribution centers from where they are allocated to the individual stores after, at most, 48 hours in the distribution center. We focus on a single item, a specific style of women's underwear. For this item, the sales data did not exhibit any significant seasonality pattern; fluctuations in sales are partially explained by price variations.

We have modeled the demand process as a multiplicative model, i.e., a special case of (1) with $\eta_{jt} = \theta_{jt} = 0$, $\gamma_{jt} = \gamma_j$, and $\delta_{jt} = \delta_j$ for all j, t, i.e.,

$$d_{jt}(p_t) = [\gamma_j + \delta_j p_t]\epsilon_{jt}. \tag{2}$$

Here, the vectors $\{\epsilon^t\}$ are independent and identically distributed as the random vector ϵ. The multiplicative model in (2) assumes that the coefficients of variation

Table 1 Demand Function Parameters for Initial Set of Five Retailers

	Location 1	Location 2	Location 3	Location 4	Location 5
η_j	212	100	72	81	31
θ_j	−36	−17	−12	−14	−5
c.v.	0.56	0.51	0.51	0.47	0.74

(c.v.s) of the demands are independent of the prices charged, an assumption satisfied by the sales data. Without any loss of generality, we normalize the ϵ variables so that $\mathsf{E}(\epsilon_j) = 1$ for all $j = 1, \ldots, J$. In our problem sets, we alternatively assume that all of the components of the vector of random variables ϵ have (i) normal $(1, \sigma)$ or (ii) gamma (λ, λ) distribution with density $f(x) = e^{-\lambda x}\lambda^{\lambda}x^{\lambda-1}/\Gamma(\lambda)$. The choice of the normal distribution is supported by the shape of the distributions of the standardized residuals in the regression equations used to estimate the parameters. However, the gamma distribution provides, in most cases, a close fit as well. (The family of gamma distributions can be used for arbitrary c.v. values. The normal distribution, on the other hand, is ill suited when the c.v. value is above 0.5 because the probability of the normal distribution adopting negative values is too large, in this case.) Finally, to simplify the specification of the problem instances, we assume that demands are independent across different stores. (The actual data show moderately positive correlations. In such cases, the value of coordinated ordering may be somewhat reduced because there will be fewer instances in which large demands at one retailer are "offset" by small demands at another.) The error term process $\{\epsilon^t\}$ is thus characterized by the parameters $\sigma_j = \text{stdev}(\epsilon_{jt}) = \text{c.v.}(d_{jt})$ for all j and t when the distribution is normal, and $\lambda_j = [\text{stdev}(\epsilon_j)]^{-2} = [\text{c.v.}(d_{jt})]^{-2}$ when it is gamma.

The parameters η, θ, and σ were determined via a generalized least squares regression analysis (to account for heteroscedasticity). Table 1 lists these for an initial set of five locations. The weekly holding cost rates h_j^+ amount to 2% of the procurement cost rates c_j. (This cost rate includes the cost of capital, maintenance, insurance, and handling costs.) To assure high service levels, we have set the backlogging cost rates equal to at least double the variable procurement costs. Because the stores are located in the same general area, they experience identical per-unit cost rates, i.e., $h^+ = \$0.06$, $h^- = \$6.00$, and $c = \$3.03$.

We first show that integrated price and inventory control may result in significant profit increases when compared to what appears to be the best one can do in a traditional, *sequential* planning approach where prices are set *first* followed by a systemwide replenishment policy to respond to these prices *next.* To this end, we compare the expected profits under the best of our proposed (heuristic) strategies with $\pi(p)$, those under the best-known replenishment strategy responding to a prespecified vector of prices $p = (p_1, \ldots, p_T)$. Several strawmen could be used for the latter, e.g., the prices that maximize gross profits, exclusive of inventory related costs. A better strawman is the price vector chosen to optimize the deterministic version of our model, which is discussed in the next section and incorporates all inventory-related costs, thus representing a limited form of integration of pricing and inventory control. It seems that the optimal fixed vector $\bar{p}$ that maximizes $\pi(p)$ can only be found by an intractable complete enumeration of all $|P|^T$ vectors, with $|P|^T$ as the number of possible price levels because no structural properties of the function $\pi(\cdot)$ are known to be exploitable. ($|P| = 50$ as the retail price varies between \$3.00 and \$5.50 in increments of 5 cents.)

Under a fixed-price vector $\bar{p}$, the problem reduces to the two-echelon inventory problem in Eppen and Schrage (1981) and Federgruen and Zipkin (1984a,c) whose heuristics incur expected costs extremely close to a lower bound of the optimal costs. The simulated average profit under the resulting, sequentially determined strategy is referred to as SEQ. In Figure 1, we compare the SEQ values with the profits under the integrated strategy (denoted by INT) fitting gamma distributions to the parameters in Table 1, as described above. (The numbers within parentheses in Figure 1, as well as those in Figures 2 and 3, refer to the percentage improvement for INT as compared with SEQ.) Note that INT outperforms SEQ by 1.20%, a considerable improvement of gross profits for most retailers. If the c.v.'s of the stores' demands were 25 or 50% larger, the number goes up to 2.2 and 3.52%,

Figure 1 Comparison of INT and SEQ Profits

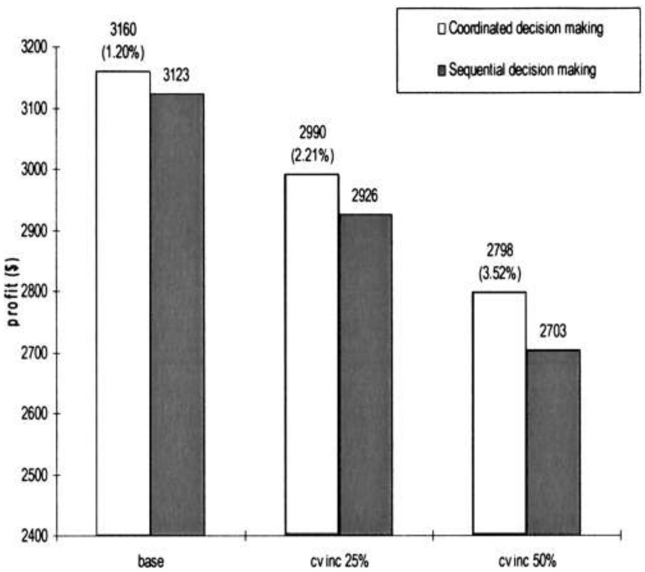

respectively.[4] The profit improvements in Figure 1 arise because of the following: The sequential strategy adopts a constant price of $4.45 during the entire season, in this instance with stationary data. Our recommended integrated strategy implements a target price of $4.55 except for the last three weeks with target prices of $4.60, $4.80, and $5.10.[5] However, under our policy the target price is used only if the systemwide inventory position is sufficiently low. If, due to low sales figures, the inventory level is above a certain threshold value, the retail price is reduced as a function of the prevailing inventory. For example, in our simulations (with 3,000 replicas) the average price is $4.77 and $4.78 in the last two weeks.

Under random demands, a higher target price is desirable. This reduces both the mean and the standard deviation of weekly demands and the expected inventory costs. This cost reduction more than offsets the loss in expected revenues. The larger the c.v.'s, the larger the weight of the inventory-related costs in the profit expression and the higher the target retail prices, p_t^*, are set: If the c.v. values are increased by 25% (50%), p_t^* is set at $4.55 ($4.60) in the first 17 weeks, followed by an increase to $4.60 ($4.65), $4.70 ($4.85), $4.95 ($5.15), and $5.20 ($5.40) in the last four weeks. Random demands call for safety stocks. Toward the end of the sales season, the retailers face the very significant risk of these safety stocks remaining unsold, with complete loss of their purchase value. The integrated policy therefore calls for a reduction of the replenishment quantities (order-up-to levels) in the last few weeks of the season and a parallel reduction of the mean and standard deviation of the demand via a price increase. In addition, the integrated policy often deviates from the target price p_t^* when relatively high inventory levels call for markdowns. For example, if the c.v. value is 50% above the base case, the average implemented prices in the last two weeks are $4.84 and $4.71, while the target prices, p_t^*, are $5.15 and $5.40, respectively. In other words, it is essential for prices to react to observed inventory levels as opposed to being fixed a priori. Moreover, the benefits of integrated planning decrease with the length of the season. Figure 2 displays for our base instance the benefits for sales seasons of $T = 5$, 10, 15, and 21 weeks. Under $T = 10$ (5) the integrated approach increases profits by 2.1% (4.6%)!

The *integrated* policy with contingent pricing, while of a fairly simple structure, adjusts the price as a function of the prevailing aggregate inventory position in the system. More specifically, it adopts, for each period t, a pair of target values (y_t^*, p_t^*) such that the list price p_t^* is implemented when it is feasible to set the aggregate inventory position at y_t^*; on the other hand, when the aggregate inventory position is above y_t^*, a lower price may be prescribed to stimulate demand and decrease inventories. (In the presence of capacity limits, a further modification in the structure is needed; see §6 for details.) A final question of interest is how much would be lost if a simpler

[4] Even when the c.v. values are increased by 50%, they range from 0.6 to 1.11. Such coefficients of variation are by no means exceptionally large. Some of the locations in our sample exhibited c.v. values for single store-demand of up to 1.73. Agrawal and Smith (1996) report on sales data for a collection of 41 SKUs of basic men's slacks sold in 24 different stores. Dividing the year into ''peak'' and ''off-peak'' weeks gives rise to a sample of 1,968 demand distributions that the authors partition into eight categories. The average c.v. values in the eight categories range from 0.64 to 3.

[5] The integrated strategy adopts, for each period t, a pair of target values (y_t^*, p_t^*) such that the target price p_t^* is implemented when it is feasible to set the aggregate inventory position at y_t^*.

Figure 2 **Comparison of INT and SEQ Profits for Varying Lengths of Sales Season**

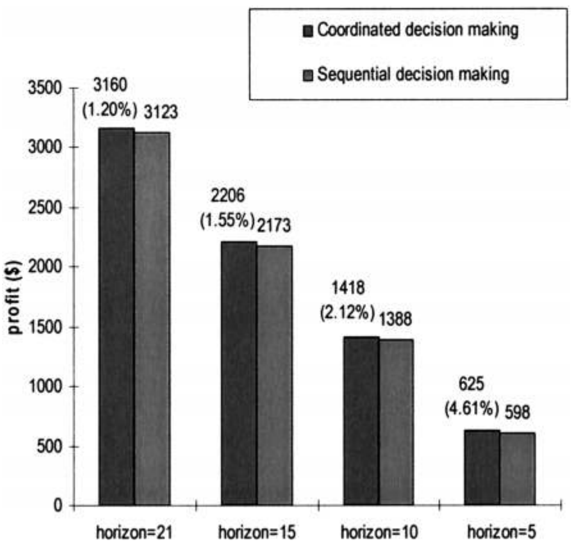

Figure 3 **Comparison of INT and Simplified Sequential Profits**

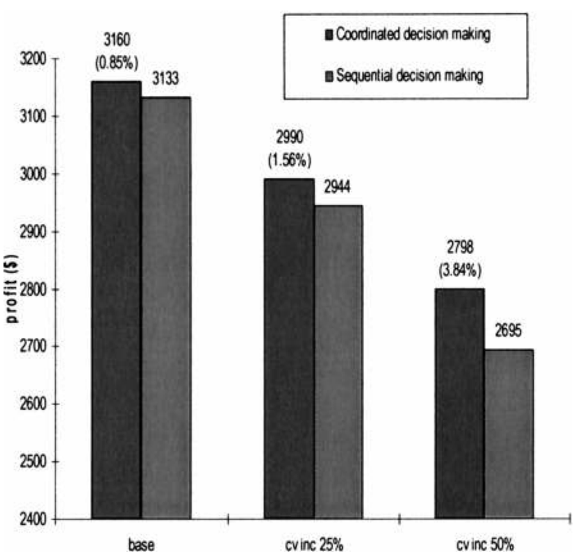

policy with the *fixed* vector of list prices $\{p_1^*, \ldots, p_T^*\}$ were adopted, complemented with the best known systemwide replenishment policy to respond to *this* price vector. For the three instances displayed in Figure 3, the profit losses resulting from the adoption of *this sequential* policy with *uncontingent* pricing are 0.85, 1.56, and 3.84%. Note, in addition, that this version of a sequential policy requires the solution of our *integrated* model for the determination of the price vector p^*. Finally, as an additional interesting benchmark, in 1997 our national retailer employed an *integrated* though more erratic policy, starting with a price of \$4.87 for the first eight weeks, followed by a promotion for five weeks with markdowns in excess of 50%, several weeks with a price of \$4.99, and an end-of-season clearance sale with major markdowns. (Price variations, while perhaps driven by considerations that are distinct from those treated here, were not based on a given generalizable policy structure. Therefore, it is impossible to evaluate this approach in a systematic manner.)

We complete this section with a comparison of the performance of (I) our coordinated systems with common but contingent pricing, (II) uncoordinated replenishments with common and uncontingent (scheduled) pricing, and (III) uncoordinated replenishments with location-specific and contingent pricing. The optimal profit values under (I) and (III) dominate the optimal value under (II) because each relaxes the strategy space of (II). The comparison between the optimal profit values under (I) and (III) is less clear-cut: (III) enjoys the advantage of being able to differentiate prices by location. On the other hand, the coordinated system selects lower inventories for any given service level (or parameters $\{h_j^+, h_j^-\}$) because the risks associated with demand uncertainty can be pooled, at least during the first common leadtime stage. A second advantage of coordinated replenishments comes from the limited order capacity in each period. If retailers manage their inventories independently, a number of retailers may request orders in the same period, with the total of the orders exceeding the capacity. This can result in a delay in receipt of orders for at least some of the retailers or possibly a (partial) cancellation of an order. We show that independent retailers are sometimes more profitable, in particular when the first leadtime L is small and the retailers rather heterogeneous. The coordinated system (I) is more likely to perform better when the differences between the retailers' attributes are moderate or the leadtime L or the variability of demands is larger.

Table 2 Demand Function Parameters for Base Set of Retailers

	Location 1	Location 2	Location 3	Location 4	Location 5
η_j	183	67	245	66	81
θ_j	−30	−9	−46	−11	−14
c.v.	0.06	0.35	0.19	0.33	0.38

To show that coordination benefits quickly dominate, we use a set of instances for the following five locations with *low* c.v. values, if anything, biasing the case *against* ''coordination'';

Because the c.v's are below 0.5, we fit normal distributions to the data. We consider at one extreme the base scenario for this set of retailers and at the other a transformed scenario with *fully identical* retailers, each with a demand function $d(p_t) = [\eta + \theta p_t]\epsilon_t$ where $\eta = (1/J)\ \Sigma_{j=1}^{J}\ \eta_j$, $\theta = (1/J)\ \Sigma_{j=1}^{J}\ \theta_j$, and with $\sigma^2(\epsilon_t) = (1/J)\ \Sigma_{j=1}^{J}\ \sigma_j^2$. We also consider two intermediate cases. In the first, the retailers share the deterministic part of their demand functions but retain their individual σ_j-values. In the second case, all retailers have the same c.v., $\sigma = [(1/J)\ \Sigma_{j=1}^{J}\ \sigma_j^2]^{1/2}$, and the demand functions' deterministic parts are $\bar{d}_j(p) = w_j\bar{d}(p)$ with $\{w_j\}$ appropriately selected weights and $\bar{d}(p)$ the deterministic part of the aggregate demand. (In the latter case, the retailers differ from each other by a scale factor only.) In Table 3, we evaluate the four cases for seven total leadtime values $\tau = 0, \ldots, 6$, setting $L = \lfloor \tau/2 \rfloor$, by simulation of the corresponding policies. π(I) represents the profit value under our (best) integrated pricing and inventory strategy, ϕ^*, π(III) represents the sum of the retailers' *optimal* profit values when acting independently, and π(II) when the retailers use the common price vector p^* of ϕ^* and each replenishes his inventory under the resulting demand process, *independently.* An ''I'' (''NT'') in the column labeled ''c.v.'' or ''μ'' denotes that the c.v.'s, or the deterministic parts of the demand functions, are identical (nonidentical).

For the first seven instances, with fully nonidentical retailers, π(III) dominates π(I) for all leadtime combinations. For this set of instances, the retailers exhibit a large degree of heterogeneity, so that the benefits of location-specific pricing dominate those that result from coordinated replenishments. On the other hand, the relative gaps between π(I) and π(III) increase with L because the benefits of coordinated replenishments increase as L (the *common* first-order leadtime) increases, for a given allocation leadtime l. The larger L is, the longer the period of time during which detailed allocations to the individual retailers can be postponed and the demand risks faced by each retailer can be pooled. For the last set of seven instances with fully identical retailers, the coordinated system (I) begins to outperform the system with independent retailers (III), as soon as L is positive. Again, the relative gaps between π(I) and π(III) increase in L, for a given allocation leadtime l. For the two intermediate cases, π(I) is at least as large as π(III) for *all* leadtime combinations. Figure 4 exhibits the average values of π(I), π(II), and π(III) for the seven leadtime combinations.

Table 3 Simulated Total Profit for Cases (I), (II), and (III)

L	l	c.v.	μ	π (I)	π (II)	π (III)
0	0	NI	NI	4348.71 ± 4.09	4336.50	4525.81
0	1	NI	NI	4259.83 ± 4.49	4256.86	4445.25
1	1	NI	NI	4227.03 ± 5.01	4193.07	4381.24
1	2	NI	NI	4128.88 ± 5.56	4142.86	4331.22
2	2	NI	NI	4132.76 ± 5.68	4100.37	4292.19
2	3	NI	NI	4069.53 ± 6.14	4060.67	4254.78
3	3	NI	NI	4051.46 ± 6.33	4019.86	4223.12
0	0	I	NI	4290.04 ± 6.98	4109.58	4117.35
0	1	I	NI	4185.19 ± 7.49	3965.95	3973.63
1	1	I	NI	4144.42 ± 8.15	3846.58	3860.91
1	2	I	NI	4043.9 ± 9.55	3752.49	3774.19
2	2	I	NI	4018.90 ± 10.08	3665.97	3700.84
2	3	I	NI	3903.37 ± 11.35	3598.12	3642.71
3	3	I	NI	3846.3 ± 12.40	3538.00	3604.73
0	0	NI	I	4309.99 ± 5.11	4280.59	4284.91
0	1	NI	I	4210.15 ± 5.68	4170.07	4176.47
1	1	NI	I	4165.85 ± 6.28	4091.02	4097.83
1	2	NI	I	4085.25 ± 6.89	4025.74	4033.61
2	2	NI	I	4073.52 ± 6.97	3971.04	3988.05
2	3	NI	I	4004.42 ± 7.54	3918.90	3949.90
3	3	NI	I	3902.95 ± 8.72	3876.93	3912.40
0	0	I	I	4319.61 ± 3.92	4358.50	4359.95
0	1	I	I	4270.60 ± 4.31	4280.55	4280.80
1	1	I	I	4251.84 ± 4.65	4220.30	4220.45
1	2	I	I	4191.84 ± 5.28	4170.90	4170.90
2	2	I	I	4183.15 ± 5.46	4132.85	4132.85
2	3	I	I	4131.93 ± 5.89	4096.35	4096.50
3	3	I	I	4126.00 ± 6.31	4058.30	4069.20

4. The Deterministic Model

We now discuss the deterministic version of our model. Considerably simpler than the stochastic model, it reduces to a simple mathematical program with linear constraints and a separable concave objective to be maximized. While interesting in its own right, we develop this model primarily because its optimal price path is used to design an effective approximate model for the *stochastic* model. Thus, let $\epsilon_{jt} = 0$ for all j, t, almost surely so that $d_{jt}(p_t) = \eta_{jt} + \theta_{jt}p_t$. Let x_{jst} = the part of the sales of retailer j in period t that are procured as part of an order placed by the distribution center in period s where $t = 1, \ldots, T$ and $s = L + l_j + 1, \ldots, t + L + l_j$. One easily determines the total (discounted) ordering, shipment, carrying, and backlogging cost, γ_{jst}, for any unit sold at retailer j in period t and procured as part of a systemwide order in period s. If $s = (>, <)\ t + L + l_j$, the unit is delivered just-in-time (carried in inventory, backlogged). The deterministic problem can thus be formulated as follows:

$$\text{(DET)}\max \sum_{t=1}^{T} \left\{ \left[\left(\sum_j \eta_{jt} \right) + \left(\sum_j \theta_{jt} \right) p_t \right] p_t - \sum_{j,s,t} \gamma_{jst} x_{jst} \right\}$$

$$\text{s.t.} \sum_s x_{jst} = d_{jt} = \eta_{jt} + \theta_{jt} p_t; \quad \text{all } j,\ t$$

$$\sum_{j,t} x_{jst} \le b_s; \quad \text{all } s\ p_t,\ x_{jst} \ge 0.$$

In the uncapacitated case (all $b_s = \infty$), the problem decomposes into T separate problems, one for each period. Moreover the tth period problem reduces to the maximization of a simple quadratic function in p_t: Let $\gamma^*_{jt} = \min_s \gamma_{jst}$ ($j = 1, \ldots, J$; $t = 1, \ldots, T$). The optimal price for period t, p^d_t is thus the unique maximizer of the quadratic function $\Sigma_j \{\eta_{jt} + \theta_{jt}p_t\}\{p_t - \gamma^*_{jt}\} = (\Sigma_j \theta_{jt})p_t^2 - \Sigma_j \gamma^*_{jt}\eta_{jt} + (\Sigma_j \eta_{jt} - \Sigma_j \theta_{jt}\gamma^*_{jt})p_t$; i.e., $p^d_t = [\Sigma_j (-\eta_{jt} + \theta_{jt}\gamma^*_{jt})]/2\, \Sigma_j \theta_{jt}$. In the capacitated case, any standard convex programming method can be invoked to solve the problem (DET). In the capacitated case with a single retailer, the problem reduces to the flexible pricing model of Chan et al. that can be solved by a simple greedy procedure. The numerical study in this paper shows that the benefits of varying the price over time increase as capacity becomes more constrained or the variability of capacity, or that of the demand functions, increase.

Figure 4 **Simulated Total Profit for Cases (I), (II), and (III)**

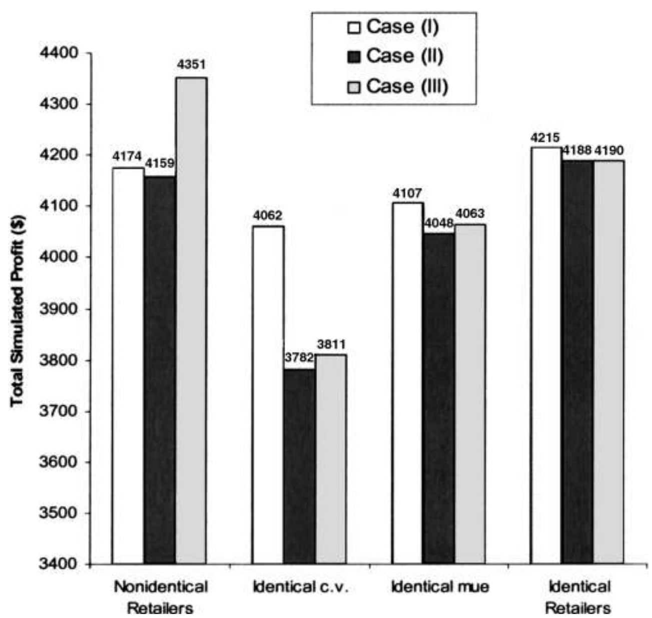

5. An Approximate Model

An *exact* dynamic program for the *original* model has a state space of dimension $J + L + \Sigma_{j=1}^{J} l_j$. We therefore develop an approximate model with an optimal policy of relatively simple structure. In models with exogenous prices, and hence demand distributions, it is possible to reduce the state space dimension to $(J + L)$ by a standard device, whereby period t is charged with the *expected* inventory costs one allocation leadtime later. This reduction is possible in standard models with backlogging because the distribution of the inventory level at retailer j at the end of period $t - l_j$ is determined once $\hat{x}_{jt}$ = retailer j's inventory position at the beginning of period t, but after that period's allocations = $\tilde{x}_{jt} + z_{jt}$ is known, where $\tilde{x}_{jt} = x_{jt} + \Sigma_{s=t+l}^{t-1} z_{js}$ denotes retailer j's inventory position at the beginning of period t *before* that period's allocation decisions, i.e., its physical inventory plus outstanding shipments. The expected costs charged to period t are thus given by

$$G_{jt}(\hat{x}_{jt}) = \alpha^{l_j} h^{+}_{j,t-l_j} \mathsf{E}[\hat{x}_{jt} - d_{jt} - d_{j,t-1} - \cdots - d_{j,t-l_j}]^{+} + \alpha^{l_j} h^{-}_{j,t-l_j} \mathsf{E}[d_{jt} + d_{j,t-1} + \cdots + d_{j,t-l_j} - \hat{x}_{jt}]^{+}. \quad (3)$$

In our setting, this accounting device fails because the leadtime demand depends on the prices selected in future periods $t-1, t-2, \ldots, t-l_j$, and not just on the price selected in period t itself. This difficulty precludes exact evaluation, even in a single-location setting. A simple approximation would consist of replacing the function G_{jt} by

$$\tilde{G}_{jt}(\hat{x}_{jt}, p_t) = \alpha^{l_j} h^{+}_{j,t-l_j} \mathsf{E}[\hat{x}_{jt} - d_{jt}(p_t) - d_{j,t-1}(p_t) - \cdots - d_{j,t-l_j}(p_t)]^{+} + \alpha^{l_j} h^{-}_{j,t-l_j} \mathsf{E}[d_{jt}(p_t) + d_{j,t-1}(p_t) + \cdots + d_{j,t-l_j}(p_t) - \hat{x}_{jt}]^{+},$$

i.e., the expected inventory costs one allocation leadtime later, if the currently selected price is maintained over the course of the leadtime.

This approximation may be adequate if price changes are infrequent, as in stationary instances with long planning horizons. When price changes are more frequent, the approximation is too crude. Instead, we exploit the fact established in simpler single-location models, that the optimal price path tends to be closely approximated by the optimal price path under deterministic demands; see, e.g., the numerical results in Gallego and van Ryzin (1997) and Federgruen and Heching (1999). For their setting, Gallego and van Ryzin show that the optimal deterministic price path is asymptotically optimal as the initial stock level and the horizon length tend to infinity.

We thus develop an alternative approximate cost function, $\hat{G}_{jt}(\hat{x}_{jt}, p_t)$, assuming that prices over the leadtime change in the same proportion as the optimal deterministic prices $\{p_t^d\}$:

$$\hat{G}_{jt}(\hat{x}_{jt}, p_t) = \alpha^{l_j} h^{+}_{j,t-l_j} \mathsf{E}\left[\hat{x}_{jt} - d_{jt}(p_t) - d_{j,t-1}\left(p_t \cdot \frac{p^d_{t-1}}{p^d_t}\right) - \cdots - d_{j,t-l_j}\left(p_t \cdot \frac{p^d_{t-l_j}}{p^d_t}\right)\right]^{+} + \alpha^{l_j} h^{-}_{j,t-l_j} \mathsf{E}\left[d_{jt}(p_t) + d_{j,t-1}\left(p_t \cdot \frac{p^d_{t-1}}{p^d_t}\right) + \cdots + d_{j,t-l_j}\left(p_t \cdot \frac{p^d_{t-l_j}}{p^d_t}\right) - \hat{x}_{jt}\right]^{+}.$$

$\tilde{G}_{jt} \equiv \hat{G}_{jt}$ iff the deterministic price path $\{p_t^d\}$ is constant over the course of the leadtime.

To facilitate the derivations, we first develop the approximate model for the case where $L = 0$. Under the approximate cost functions $\hat{G}_{jt}(\cdot, \cdot)$, the model has the vector $\tilde{x}^t = (\tilde{x}_{1t}, \tilde{x}_{2t}, \ldots, \tilde{x}_{Jt})$ as its state at the beginning of period t, and the set of feasible actions in state $\tilde{x}$ is given by: $A_t(\tilde{x}) = \{(\hat{x}, p): \tilde{x} \le \hat{x}; \sum_{j=1}^J \hat{x}_j \le \sum_{j=1}^J \tilde{x}_j + b_t; p_{t,\min} \le p \le p_{t,\max}\}$.

$v_t(\tilde{x})$, the expected maximum total profit in periods $1, \ldots, t$ when starting period t in state $\tilde{x}$, satisfies the following recursion: $v_0(\cdot) \equiv 0$ and

$$v_t(\tilde{x}) = \sum_{j=1}^{J} c_{jt}\tilde{x}_j + \max_{(\hat{x},p)\in A_t(\tilde{x})} \left\{ \alpha p \sum_{j=1}^{J} \mathsf{E} d_{jt}(p) - c_{0t}\left(\sum_{j=1}^{J} \hat{x}_j - \sum_{j=1}^{J} \tilde{x}_j\right) - \sum_{j=1}^{J} c_{jt}\hat{x}_j - \sum_{j=1}^{J} \hat{G}_{jt}(\hat{x}_j, p) + \alpha \mathsf{E} v_{t-1}[\hat{x} - d^t(p)] \right\}, \quad (4)$$

where $d^t(p) = (d_{1t}(p), d_{2t}(p), \ldots, d_{Jt}(p))$ is the vector of demands in period t. The first term in the expression within curled brackets denotes period t's expected revenues, assumed to be received at the end of the period. The terms $\sum_{j=1}^J c_{jt}\tilde{x}_j - \sum_{j=1}^J c_{jt}\hat{x}_j - c_{0t}(\sum_{j=1}^J \hat{x}_j - \sum_{j=1}^J \tilde{x}_j)$ account for all variable order and shipment costs in period t. The recursion can be simplified by the following transformation: Let $v_t^*(\tilde{x}) = v_t(\tilde{x}) - \sum_{j=1}^J c_{jt}\tilde{x}_{jt}$. Substituting $v_t^*(\cdot)$ for $v_t(\cdot)$ we obtain:

$$v_t^*(\tilde{x}) = \max_{(\hat{x},p)\in A_t(\tilde{x})} \left\{ \alpha p \sum_{j=1}^{J} \mathsf{E} d_{jt}(p) - c_{0t}\left(\sum_{j=1}^{J} \hat{x}_j - \sum_{j=1}^{J} \tilde{x}_j\right) - \sum_{j=1}^{J} c_{jt}\hat{x}_j - \sum_{j=1}^{J} \hat{G}_t(\hat{x}_j, p) \right.$$

$$+ \alpha \sum_{j=1}^{J} c_{j,t-1} \mathsf{E}[\hat{x}_j - d_{jt}(p)]$$

$$+ \alpha \mathsf{E} v_{t-1}^{*}[\hat{x} - d^{t}(p)] \Big\}$$

$$= \max_{(\hat{x},p) \in A_t(\tilde{x})} \Big\{ \alpha \sum_{j=1}^{J} (p - c_{j,t-1}) \mathsf{E} d_{jt}(p)$$

$$+ \sum_{j=1}^{J} (\alpha c_{j,t-1} - c_{jt}) \hat{x}_j$$

$$- c_{0t} \left(\sum_{j=1}^{J} \hat{x}_j - \sum_{j=1}^{J} \tilde{x}_j \right) - \sum_{j=1}^{J} \hat{G}_{jt}(\hat{x}_j, p)$$

$$+ \alpha \mathsf{E} v_{t-1}^{*}[\hat{x} - d^{t}(p)] \Big\}. \tag{5}$$

This dynamic program's state space is J-dimensional, and therefore continues to be intractable for all but the smallest values of J. As our second and final approximation step, we relax the feasible action sets $\{A_t(\tilde{x})\}$ to sets $\{\hat{A}_t(\tilde{x})\}$ by aggregating the constraints $\hat{x} \geq \tilde{x}$ into the single constraint $\Sigma_{j=1}^{J} \hat{x}_j \geq \Sigma_{j=1}^{J} \tilde{x}_j$. An induction argument verifies that after this aggregation step, a one-dimensional dynamic program arises, with $\tilde{X} = \Sigma_{j=1}^{J} \tilde{x}_j$, the aggregate inventory position, as the single-state component. Its value-functions are defined recursively via $V_0^{*} \equiv 0$ and:

$$V_t^{*}(\tilde{X}) = \max_{\tilde{X} \leq \hat{X} \leq \tilde{X} + b_t;\, p_{t,\min} \leq p \leq p_{t,\max}}$$

$$\Big\{ \alpha \sum_{j=1}^{J} (p - c_{j,t-1})) \mathsf{E} d_{jt}(p) - c_{0t}(\hat{X} - \tilde{X})$$

$$- R_t(\hat{X}, p) + \alpha \mathsf{E} V_{t-1}^{*}\left(\hat{X} - \sum_{j=1}^{J} d_{jt}(p) \right) \Big\}, \tag{6}$$

where

$$R_t(\tilde{X}, p) = \min_{\hat{x}} \sum_{j=1}^{J} [\hat{G}_{jt}(\hat{x}_j, p) + (c_{jt} - \alpha c_{j,t-1}) \hat{x}_j] \tag{7}$$

$$\text{s.t.} \sum_{j=1}^{J} \hat{x}_j = \hat{X}. \tag{8}$$

In the next section we will establish that the dynamic program (6) can, in fact, be interpreted as the combined inventory control and pricing problem of a single-location periodic review system, and that the optimal strategy has a number of important structural properties.

We now turn to the general case where $L \geq 0$. It is easily verified that after the above two approximation steps, the following dynamic program arises:

$$V_t(\tilde{X}, w_{t+L}, \ldots, w_{t+1})$$

$$= \max_{\{w,\hat{x},p\}} \Big\{ \alpha \mathsf{E} \sum_{j=1}^{J} (p - c_{j,t-1}) d_{jt}(p) - c_{0t} w$$

$$- R_t(\tilde{X} + w_{t-L}, p)$$

$$+ \alpha \mathsf{E} V_{t-1}[\hat{x} - d_t(p), w_{t+L-1}, \ldots, w_{t+1}, w] \Big\}. \tag{9}$$

This dynamic program has a state space of dimension $(L + 1)$, independent of J but still impractical. The above accounting device can again be invoked, charging to period t, not $R(\tilde{X} + w_{t+L}, p)$, but rather its expected value one order leadtime (L periods) later, i.e.,

$$\alpha^{L} \mathsf{E}\{R_{t-L}[\tilde{X}_t + w_{t+L} + w_{t+L-1} + \cdots + w_t - D_t(p) - D_{t-1}(p_{t-1}) - \cdots - D_{t-L}(p_{t-L})]\}, \tag{10}$$

where $D_t(p) = \Sigma_{j=1}^{J} d_{jt}(p)$. As with $G_{jt}(\cdot)$, (10) depends on all prices charged over the order leadtime, which are unknown at the start of period t. Similar to the approximation $\hat{G}_{jt}(\cdot, \cdot)$ for $G_{jt}(\cdot)$, the first of our two approximation steps, we replace, if $L > 0$, (10) by:

$$\alpha^{L} \mathsf{E}\left\{ R_{t-L}\left[X_t + w_{t+L} + w_{t+L-1} + \cdots + w_t - D_t(p) - D_{t-1}\left(p \cdot \frac{p_{t-1}^{d}}{p_t^{d}} \right) - \cdots - D_{t-L}\left(p \cdot \frac{p_{t-L}^{d}}{p_t^{d}} \right) \right] \right\}.$$

To simplify the notation let:

$\tilde{Y}_t = \tilde{X}_t + \Sigma_{s=t+1}^{t+L} w_s$, the aggregate inventory position at the start of period t *before* a replenishment order is placed, i.e., all units in stock at the retailers, on order from the supplier or in transit between the distribution center and the retailers, minus the retailers' backlog.

$\hat{Y}_t = \tilde{Y}_t + wt = \tilde{X}_t + \Sigma_{s=t}^{t+L} w_s$, the systemwide inventory position at the beginning of period t *after* an order is placed by the depot.

We thus replace (10) by:

$$\hat{R}_t(\hat{Y}, p) = \alpha^L \mathsf{E} R_{t-L}\left[\hat{Y} - D_t(p) - D_{t-1}\left(p \cdot \frac{p_{t-1}^d}{p_t^d}\right) - \cdots - D_{t-L}\left(p \cdot \frac{p_{t-L}^d}{p_t^d}\right), p\right], \tag{11}$$

which can be obtained in closed form as the one-step, expected-cost function of a single-location model with a specific leadtime demand distribution. This results in a one-dimensional dynamic program, similar to (6), with $\tilde{Y}_t$ as the single-state component:

$$V_t(\tilde{Y}) = c_{0t}\tilde{Y} + \max_{\{\tilde{Y} \le \hat{Y} \le \tilde{Y} + b_t;\, p_{t,\min} \le p \le p_{t,\max}\}} \left\{\alpha \sum_{j=1}^{J} (p - c_{j,t-1}) \mathsf{E} d_{jt}(p) - c_{0t}\hat{Y} - \hat{R}_t(\hat{Y}, p) + \alpha \mathsf{E} V_{t-1}\left[\hat{Y} - \sum_{j=1}^{J} d_{jt}(p)\right]\right\}. \tag{12}$$

It is once again convenient to eliminate the first term to the right of (12). This can be achieved via the transformation $V_t^*(\tilde{Y}) = V_t(\tilde{Y}) - c_{0t}\tilde{Y}$ similar to that transforming (4) into (5):

$$V_t^*(\tilde{Y}) = \max_{\{\tilde{Y} \le \hat{Y} \le \tilde{Y} + b_t; p_{t,\min} \le p \le p_{t,\max}\}} H_t(\hat{Y}, p),$$

where

$$H_t(\hat{Y}, p) = \alpha \sum_{j=1}^{J} (p - c_{j,t-1} - c_{0,t-1}) \mathsf{E} d_{jt}(p) - (\alpha c_{0,t-1} - c_{0,t})\hat{Y} - \hat{R}_t(\hat{Y}, p) + \alpha \mathsf{E} V_{t-1}^*\left[\hat{Y} - \sum_{j=1}^{J} d_{jt}(p)\right]. \tag{13}$$

6. Structural Properties

The structure of the approximate dynamic program (13) is similar to that describing the combined pricing and inventory control problem for a single-location problem; see Federgruen and Heching (1999). We now show that an optimizing strategy exists that is a base-stock/list price policy, a term coined by Porteus (1990). In each period t, there exists a pair of target values (y_t^*, p_t^*) such that it is optimal to increase the aggregate systemwide inventory position $\tilde{Y}$ to a level as close as possible to y_t^*, and to charge a price p_t^* if $\hat{Y} = y_t^*$, a price $p_t \ge p_t^*$ if $\hat{Y} < y_t^*$, and a price $p_t \le p_t^*$ otherwise. More specifically: (i) If $\tilde{Y} < y_t^* - b_t$, then $\hat{Y} = \tilde{Y} + b_t$ and $p_t \ge p_t^*$; (ii) if $y^* - b_t \le \tilde{Y} \le y^*$, then $\hat{Y} = y_t^*$ and $p_t = p_t^*$; (iii) if $\tilde{Y} > y_t^*$, then $\hat{Y} = \tilde{Y}$ and $p_t \le p_t^*$. We need the following lemma.[6]

Lemma 1. *Fix $t = 1, \ldots, T$. The function $R_t(\hat{X}, p)$ is jointly convex and supermodular.*

Proof. (a) The functions $\hat{G}_{jt}(\hat{x}_{jt}, p_t)$ are jointly convex; see, e.g., Lemma 1 in Federgruen and Heching (1999). This implies that for all $t = 1, \ldots, T$ the function $R_t(\cdot, \cdot)$ is *jointly convex* as well: For any $(\hat{X}^{(1)}, p_1)$ and $(\hat{X}^{(2)}, p_2) \in \Re^2$, let $\hat{x}^{(1)}$ and $\hat{x}^{(2)}$ achieve the minimum in (6) for $\hat{X} = \hat{X}^{(1)}$ and $\hat{X} = \hat{X}^{(2)}$, respectively. Observe that $(\hat{x}^{(1)} + \hat{x}^{(2)})/2$ satisfies constraint (8) with right-hand side $\hat{X} = (\hat{X}_1 + \hat{X}_2)/2$ and is therefore a feasible solution to the mathematical program {(7), (8)}. Thus,

$$\begin{aligned}
R_t\left(\frac{\hat{X}^{(1)} + \hat{X}^{(2)}}{2}, \frac{p_1 + p_2}{2}\right) &\le \sum_{j=1}^{J} \hat{G}_{jt}\left(\frac{\hat{x}_j^{(1)} + \hat{x}_j^{(2)}}{2}, \frac{p_1 + p_2}{2}\right) + \sum_{j=1}^{J} (c_{jt} - \alpha c_{j,t-1})\left(\frac{\hat{x}_j^{(1)} + \hat{x}_j^{(2)}}{2}\right) \\
&\le \frac{1}{2}\left[\sum_{j=1}^{J} \hat{G}_{jt}(\hat{x}_j^{(1)}, p_1) + \sum_{j=1}^{J} \hat{G}_{jt}(\hat{x}_j^{(2)}, p_2)\right] + \frac{1}{2}\sum_{j=1}^{J} (c_{jt} - \alpha c_{j,t-1})\hat{x}_j^{(1)} + \frac{1}{2}\sum_{j=1}^{J} (c_{jt} - \alpha c_{j,t-1})\hat{x}_j^{(2)} \\
&= \frac{1}{2}[R_t(\hat{X}_1, p_1) + R_t(\hat{X}_2, p_2)],
\end{aligned}$$

where the second inequality follows from the joint convexity of $\hat{G}_t(\cdot, \cdot)$. To show supermodularity, suppose the random variables $\{\epsilon_{jt}\}$ are continuous with unbounded support. (Similar proofs exist for other

[6]A bivariate function $F(x, y)$ is supermodular (submodular) if, for all $x^{(2)} > x^{(1)}$, the difference function $F(x^{(2)}, y) - F(x^{(1)}, y)$ is nondecreasing (nonincreasing) in y.

distributional assumptions.) In this case, it is easily verified and well known that the function $\hat{G}_{jt}(\hat{x}_j, p)$ is differentiable with respect to $\hat{x}_j$ $(j = 1, \ldots, J)$, and that $\hat{R}_t(\hat{X}, p)$ is differentiable with respect to $\hat{X}$. By Lemma 1 in Federgruen and Heching (1999), $\hat{G}_{jt}(\hat{x}_j, p)$ is strictly convex in $\hat{x}_j$ $(j = 1, \ldots, J)$. The Kuhn-Tucker conditions for the optimum $\hat{x}(p)$ in the mathematical program {(7), (8)} are thus necessary and sufficient. Let $\lambda(\hat{X}, p)$ denote the dual price of (8). Then,

$$\frac{\partial R_t(\hat{X}, p)}{\partial \hat{X}} = \lambda(\hat{X}, p) = \frac{\partial \hat{G}_{jt}[\hat{x}_j^*(p), p]}{\partial \hat{x}_j} + (c_{jt} - \alpha c_{j,t-1}), \quad \text{for all } j = 1, \ldots, J. \tag{14}$$

It suffices to show that $\partial R_t / \partial \tilde{x}$ is nondecreasing in p. Assume to the contrary that for some $p^{(2)} > p^{(1)}$, $\lambda(\hat{X}, p^{(2)}) < \lambda(\hat{X}, p^{(1)})$. It follows from a straightforward adaptation of Theorem 2 in Federgruen and Heching (1999) that the functions $\hat{G}_{jt}(\hat{x}_j, p)$ are supermodular as well as strictly convex. Thus, for all $x_j \geq \hat{x}_j^{(1)} \geq \hat{x}_j(p^{(1)})$,

$$\frac{\partial \hat{G}_{jt}(x_j, p^{(2)})}{\partial x_j} \geq \frac{\partial G_{jt}(\hat{x}_j^{(1)}, p^{(2)})}{\partial x_j} \geq \frac{\partial \hat{G}_{jt}(\hat{x}_j^{(1)}, p^{(1)})}{\partial x_j} = 0,$$

so that $\hat{x}_j(p^{(2)}) < \hat{x}_j(p^{(1)})$ for all j, and hence $\Sigma_{j=1}^{J} \hat{x}_j(p^{(2)}) < \Sigma_{j=1}^{J} p^{(1)}) = \hat{X}$, contradicting the feasibility of $\hat{x}(p^{(2)})$ when $p = p^{(2)}$ in the mathematical program {(7), (8)}.

THEOREM 1. *Fix $t = 1, \ldots, T$.*

(a) The function $H_t(\hat{Y}, p)$ is jointly concave in $\hat{Y}$ and p and has a finite maximizer (y_t^*, p_t^*). The function $V_t^*(\tilde{Y})$ is concave in $\tilde{Y}$.

(b) The optimal price $p_t(\tilde{Y})$, to be charged in period t, is nonincreasing in the period's starting systemwide inventory position $\tilde{Y}_t$.

(c) A base-stock/list-price policy, with base-stock/list-price combinations $\{(y_t^*, p_t^*): t = 1, \ldots, T\}$ is optimal in the dynamic program (13).

PROOF.

(a) In view of Lemma 1 and the proof of Lemma 1 in Federgruen and Heching (1999) it is easily verified from (7) that the functions $\hat{R}_t(\hat{Y}, p)$ are jointly convex. Similarly, one can show that the first two terms to the right of (13) are concave in p and $\hat{Y}$, respectively. The joint concavity of the function $\hat{H}_t(\cdot, \cdot)$ can now be proven inductively; it has a finite unconstrained maximizer because $\lim_{|\hat{Y}| \to \infty} H_t(\hat{Y}, p) = \infty$ and $H_t(\hat{Y}, p) = O(|\hat{Y}|)$, which can be verified by induction as well.

(b) By Topkis (1978), it suffices to show that H_t is submodular. This can be shown inductively because all one-step expected profit terms; i.e., the first three terms to the right in (13), are submodular. While obvious for the first two terms, to verify the supermodularity of $\hat{R}_t(\hat{Y}, p)$, fix $\hat{Y}^{(2)} > \hat{Y}^{(1)}$ and $p^{(2)} > p^{(1)}$, let $C_t(\epsilon) = \Sigma_{s=t}^{t-L} [\gamma_{jt}\epsilon_{jt} + \eta_{jt}]$ and $\hat{\xi}_s = p_s^d / p_t^d$, and substitute (1) into (11) to get:

$$\begin{aligned}
&\hat{R}_t(\hat{Y}^{(2)}, p^{(2)}) - \hat{R}(\hat{Y}^{(1)}, p^{(2)}) \\
&= \mathsf{E}_{\{\epsilon_{js}\}}\Bigg[R_t(\hat{Y}^{(2)} + C_t(\epsilon) - \sum_{s=t}^{t-L} (\theta_{js} + \delta_{js}\epsilon_{js})p^{(2)}\hat{\xi}_s, p^{(2)}\hat{\xi}_{t-L}) \\
&\qquad - R_t(\hat{Y}^{(1)}) + C_t(\epsilon) \\
&\qquad - \sum_{s=t}^{t-L} (\theta_{js} + \delta_{js}\epsilon_{js})p^{(2)}\hat{\xi}_s, p^{(2)}\hat{\xi}_{t-L}\Bigg] \\
&\geq \mathsf{E}_{\{\epsilon_{js}\}}\Bigg\{R_t\Bigg[\hat{Y}^{(2)} + C_t(\epsilon) \\
&\qquad - \sum_{s=t}^{t-L} (\theta_{js} + \delta_{js}\epsilon_{js})p^{(1)}\hat{\xi}_s, p^{(2)}\hat{\xi}_{t-L}\Bigg] \\
&\qquad - R_t\Bigg[\hat{Y}^{(1)} + C_t(\epsilon) \\
&\qquad - \sum_{s=t}^{t-L} (\theta_{js} + \delta_{js}\epsilon_{js})p^{(1)}\hat{\xi}_s, p^{(2)}\hat{\xi}_{t-L}\Bigg]\Bigg\} \\
&\geq \mathsf{E}_{\{\epsilon_{js}\}}\Bigg\{R_t\Bigg[\hat{Y}^{(2)} + C_t(\epsilon) - \sum_{s=t}^{t-L} (\theta_{js} + \delta_{js}\epsilon_{js})p^{(1)}\hat{\xi}_s, p^{(1)}\hat{\xi}_{t-L}\Bigg] \\
&\qquad - R_t\Bigg[\hat{Y}^{(1)} + C_t(\epsilon) \\
&\qquad - \sum_{s=t}^{t-L} (\theta_{js} + \delta_{js}\epsilon_{js})p^{(1)}\hat{\xi}_s, p^{(1)}\hat{\xi}_{t-L}\Bigg]\Bigg\}.
\end{aligned}$$

In view of the supermodularity and convexity of R_t in its first argument (see Lemma 1) both inequalities hold for any realization of the random variables $\{\epsilon_{js}\}$ and, hence, for their expectations.

(c) In view of part (a), the pair (y_t^*, p_t^*) is optimal when feasible. This verifies (ii). In parts (i) and (iii), we first show the identities for $\hat{Y}_t$: (1) Every pair $(\hat{Y}_t, p_t)$ with $\hat{Y}_t \neq \tilde{Y}_t + b_t$, hence $\hat{Y}_t < \hat{Y}_t + b_t$, is not optimal because by the joint concavity of $H_t(\cdot, \cdot)$ the point $(\hat{Y}_t + b_t, p')$ on the line connecting $(\hat{Y}_t, p_t)$ with (y_t^*, p_t^*) is superior to $(\hat{Y}_t, p_t)$. Similarly, for (iii), when $\tilde{Y}_t > y_t^*$, every pair $(\hat{Y}_t, p_t)$ with $\hat{Y}_t \neq \tilde{Y}_t$, hence $\hat{Y}_t > \tilde{Y}_t$, is not optimal because the point $(\tilde{Y}_t, p')$ on the line connecting (y_t^*, p_t^*) with $(\hat{Y}_t, p_t)$ is superior to $(\hat{Y}_t, p_t)$. The remainder of parts (i) and (iii) follow from part (b).

Given the optimality of a base-stock/list-price policy in each period, the approximate dynamic program (12) is thus fairly easily solved; see §6 in Federgruen and Heching (1999). We conclude this section with a few structural properties of the optimizing policy, in particular the impact of shifts in the demand functions $d_{jt}(\cdot)$ and the variability of demands, on the optimal prices and inventory levels. Consider, for example, the impact of an increase in η_{jt} in the demand function (1), for some period t and retailer j, i.e., a parallel outward shift of the demand curve, perhaps due to increased market penetration and brand recognition. The stocking level for this period at this retailer can be expected to increase, for any starting inventory, provided it continues to charge the same price or, more generally, as long as it does not increase the price too excessively. The same can be conjectured when the slope of a demand function decreases in absolute value (i.e., some θ_{jt} or δ_{jt} increases), reflecting less price-sensitive demand, perhaps due to reduced availability of alternative brands or variants. We expect a similar response to an increase in the magnitude of the noise terms in (1), i.e., when some ϵ_{jt} is replaced by $\bar{\epsilon}_{jt} = k\epsilon_{jt}$ for some $k > 1$. The following theorem, proven in Federgruen and Heching (2000), verifies these conjectures: Let $\Gamma_{jt} \doteq (\gamma_{jt}, \delta_{jt}, \eta_{jt}, \theta_{jt})$; $j = 1, \ldots, J$; $t = 1, \ldots, T$.

Theorem 2.

(a) Fix $t = 1, \ldots, T$. Assume for some retailer $j^* = 1, \ldots, J$, the string Γ_{j^*t} is replaced with $\bar{\Gamma}_{j^*t} \geq \Gamma_{j^*t}$. Let $p_t(\tilde{Y}_t | \bar{\Gamma}_{j^*t})$ denote the optimal price to be charged when the parameter string in retailer j^*'s demand function is $\bar{\Gamma}_{j^*t}$ and $p_t(\tilde{Y}_t | \Gamma_{j^*t})$ if it is Γ_{j^*t}. Define $\hat{Y}_t(\tilde{Y}_t | \bar{\Gamma}_{j^*t})$ and $\hat{Y}(\tilde{Y}_t | \Gamma_{j^*t})$ in a similar way. If $p_t(\tilde{Y}_t | \Gamma_{j^*t}) \leq p_t(\tilde{Y}_t | \Gamma_{j^*t})$, then $\hat{Y}_t(\tilde{Y}_t | \bar{\Gamma}_{j^*t}) \geq \hat{Y}_t(\tilde{Y}_t | \Gamma_{j^*t})$.

(b) Fix a retailer j and period t. Assume that

(i) the demand curve for this retailer and period shifts up, i.e., η_{jt} or γ_{jt} increases,

(ii) the demand curve for this retailer and period flattens, i.e., δ_{jt} and θ_{jt} increases toward 0, or

(iii) the noise term in the demand curve for this retailer and period is almost surely enlarged by a constant factor $k > 1$ (ϵ_{jt} is replaced by $\bar{\epsilon}_{jt} = k\epsilon_{jt}$, with $k > 1$). The optimal response to any one of these changes in the approximate model (12) is to increase the systemwide order size in period t, regardless of this period's starting inventory position, but provided that the same or a lower price is charged. Conversely, the optimal response is to increase this period's price, provided that the same or lower orders are placed in this period.

The impact of increased demand variability on the optimal base-stock and list-price levels has been discussed extensively in the past, albeit for single-period models only: Mills (1959) treats the additive demand model, Karlin and Carr (1962) the multiplicative case, and Young (1978) the general structure in (1). See also Petruzzi and Dada (1990). Even for the single-period models, the results are confined to comparisons with the extreme case of zero variability, while Theorem 2 provides comparisons between *arbitrary* pairs of variability levels in general multiperiod models.

7. Heuristic Strategies

Any strategy to govern our system consists of three parts: (a) a pricing rule for each period, (b) a systemwide ordering policy that prescribes the aggregate order quantity in each period, and (c) a mechanism to allocate incoming orders at the distribution center to the individual retailers. The approximate model suggests the following heuristic strategy:

Heuristic (H^A). In period t, select a pair $(\hat{Y}, p)$ that achieves the maximum in (12) for the prevailing aggregate inventory position $\tilde{Y}$. Select p as the price for the upcoming period, and place an order of size $W = (\hat{Y} - \tilde{Y})$ with the outside supplier. This leaves us with the third-strategy component, i.e., the allocation mechanism. In the approximate model, it is optimal to allocate any incoming order so as to minimize expected costs in the very first period in which these allocations have an impact, i.e., one shipment leadtime hence, assuming that the price selected in the current period is adjusted in the same proportions as $\{p_t^d\}$, the optimal prices of the deterministic model, throughout the leadtime. See (4) for a specification of the approximate cost functions $\hat{G}_{jt}(\cdot, \cdot)$. It is therefore reasonable to employ the same allocation mechanism, i.e., to minimize (7), with p replaced by the price selected for this period. To ensure that the allocations are feasible, it is necessary and sufficient that the vector $\hat{x}$ satisfies (8) as well as the constraints

$$\hat{x} \geq \tilde{x}_j, \qquad j = 1, \ldots, J. \tag{15}$$

The resulting set of values $\{(\hat{x}_j - \tilde{x}_{jt}) : j = 1, \ldots, J\}$ represents the shipment quantities in period t to the individual retailers, under this allocation mechanism. This allocation mechanism is easy to implement and similar to mechanisms suggested for two-echelon systems with predetermined prices; see, e.g., Federgruen and Zipkin (1984a,c). The mathematical program, minimizing (7) subject to the constraints (8) and (15), is easily solvable. Note that the objective is separable and convex; see the proof of Lemma 1(a). Several highly efficient algorithms can be employed to solve this problem, see, for example, Zipkin (1980).

While simple, this combined pricing, ordering, and allocation heuristic suffers from the limitation that both the pricing and ordering decisions are based entirely on aggregate inventory information, disregarding potential inventory imbalances between retailers. For example, if the aggregate inventory position is reasonably large, while some retailers face low inventories, it may be desirable to maintain a higher price or even to increase the price to avoid stockouts at the exposed retailers. Heuristic H^A, on the other hand, may implement a lower price because the *aggregate* inventory level is relatively high. Alternatively, the challenge posed by inventory imbalances may be addressed by adjusting the order quantity via a procedure that explicitly accounts for the needs of the individual retailers, on the basis of their individual inventory positions.

Heuristic (H^D). Heuristic (H^D) maintains the pricing rule and allocation mechanisms of the basic heuristic (H^A) but specifies each period's order by the distribution center on the basis of *disaggregate* inventory information. We describe the rule first when $L = 0$. Assume that in period t the approximate model prescribes a price p_t as well as an order that increases the systemwide inventory position to a level $\hat{Y}_t$. We first disaggregate this quantity into a vector $s_t^* = (s_{1t}^*, \ldots, s_{Jt}^*)$, with s_t^* as the optimal disaggregation of $\hat{Y}_t$ in the (relaxed) allocation problem (7)–(8). More specifically,

$$\begin{aligned} &R_t(\hat{Y}_t, p_t) \\ &= \sum_{j=1}^{J} [\hat{G}_{jt}(s_{jt}^*, p_t) + (c_{jt} - \alpha c_{j,t-1}) s_{jt}^*] \\ &= \min_{\hat{x}} \left\{ \sum_{j=1}^{J} [\hat{G}_{jt}(\hat{x}_j, p_t) + (c_{jt} - \alpha c_{j,t-1})\hat{x}_j] : \sum_{j=1}^{J} \hat{x}_j = \hat{Y}_t \right\}. \end{aligned}$$

We now compare s_{jt}^*, the "ideal" order-up-to level for retailer j, with the current inventory position $\tilde{x}_{jt}$ and add any positive gap $[s_{jt}^* - \tilde{x}_{jt}]^+$ to the systemwide order quantity, until reaching the capacity level b_t. In other words, we set the order size $w_t^D = \min\{b_t, \Sigma_j [s_{jt}^* - \tilde{x}_{jt}]^+\}$. Clearly,

$$\begin{aligned} w_t^D &= \min\left\{b_t, \sum_j [s_{jt}^* - \tilde{x}_{jt}]^+\right\} \geq \min\left\{b_t, \sum_j (s_{jt}^* - \tilde{x}_{jt})\right\} \\ &= \min\{b_t, \hat{Y}_t - \hat{X}_t\} = \hat{Y}_t - \hat{X}_t = w_t^A, \end{aligned}$$

where $w_t^D > w_t^A$ in the presence of inventory imbalances, i.e., if for one or more retailers $\tilde{x}_{jt} > s_{jt}^*$.

In case $L > 0$, s_t^* is determined so as to allocate $\hat{Y}_t$ with the goal of minimizing total expected inventory holding and backlogging costs a complete order leadtime later. In other words, $\hat{Y}_t$ is allocated so as to minimize total expected inventory costs in the period in which any new order becomes available for allocation to the retailers. s_t^* achieves the minimum in the problem

$$\min\left\{\sum_{j=1}^{J} \hat{G}_{j,t-L}\left(y_j, p_t \cdot \frac{p_{t-L}^d}{p_t^d}\right) : \sum_{j=1}^{J} y_j = \hat{Y}_t\right\} \quad \text{and}$$

$$w_t^D = \min\left\{b_t, \sum_{j=1}^{J} [s_{jt}^* - \hat{x}_{jt}]^+ - [w_{t+1}^D + \cdots + w_{t+L}^D]\right\}$$

with $w^D_{t+1}, \ldots, w^D_{t+L}$ the order quantities determined in the earlier periods $t + L, \ldots, t + 1$. See Federgruen and Heching (2000) for a discussion of an additional heuristic, (H^P).

7.1. Two-Stage Heuristics

Both heuristics (H^A) and (H^D) use the prices prescribed by the optimal base-stock/list-price strategy for the approximate model, without any modifications. In the case of (H^A) the order quantity is identical to that prescribed by the optimal strategy in the approximate model, and for (H^D) it is sometimes adjusted upwards to account for inventory imbalances among the retailers. Recall, the approximate analytical model uses approximate expected cost functions $\hat{G}_{jt}(\cdot, \cdot)$, which are accurate except when leadtimes are long or frequent and sizeable price changes occur. To develop effective heuristics for the latter type of problem instances, we consider the following two-stage process.

Step 1. Solve the approximate combined inventory control/pricing model so as to obtain the optimal list-price sequence $\{p^*_t : t = T, T - 1, \ldots, 1\}$.

Step 2. Fix $p_t = p^*_t$ for all $t = 1, \ldots, T$ so that retailer j's demand distribution in period t is now given by the random variable $d^*_{jt} = d_{jt}(p^*_t)$, $j = 1, \ldots, J$ and $t = 1, \ldots, T$. The problem now reduces to the "one-warehouse, multi-retailer problem" with predetermined demand distributions addressed by Aviv and Federgruen (1999), among others. The approximate analytical model suggested by these authors to solve this problem is a special case of the approximate model in §5, which arises when all price variables p_t are set at the level $p_t = p^*_t$ throughout. Let V^{**} denote its optimal value. Determine the (modified) base-stock policy that optimizes this model, and let $\hat{\hat{Y}}_t$ denote the systemwide inventory position prescribed by this policy in a given state in period t.

We now specify two "price-first" variants of heuristics (H^A) and (H^D) that we will refer to as $(H^{PF,A})$ and $(H^{PF,D})$. Both set $p_t = p^*_t$ for every period t throughout, modifying all cost functions $\hat{G}_{jt}(\cdot, \cdot)$, $R_t(\cdot, \cdot)$, and $\hat{R}_t(\cdot, \cdot)$, accordingly. Both use myopic allocations in every period t. $(H^{PF,A})$ $[(H^{PF,D})]$ specifies the order quantity in the same way (H^A) $[(H^D)]$ does except that in period t the order-up-to level of the combined base-stock/list-price strategy is replaced by $\hat{\hat{Y}}_t$.

Table 4 Demand Function Parameters for Additional Set of Retailers

	Location 1	Location 2	Location 3	Location 4	Location 5
η_j	80	67	47	70	33
θ_j	−13	−12	−7	−12	−5
c.v.	0.26	0.29	0.19	0.22	0.33

8. Numerical Study

Our numerical study has several goals. *First,* it gauges the effectiveness of the four proposed heuristics using the approximate model as the benchmark, as well as the accuracy of the profit and cost estimates obtained by the approximate model using the simulated performance measures of the heuristic strategies as its benchmark. (As mentioned, except for the case of zero leadtimes, the approximate model is not guaranteed to result in a strict upper bound for the optimal profit value.)

Second, the study investigates how different system parameters impact various performance measures. These parameters include the desired service level, the variability of weekly demands, the impact of non-stationarities in the demand patterns arising, for example, because of seasonalities, the magnitude of the leadtimes, and the number of retailer locations.

Third, we use the study to provide insights into a number of important strategic questions:

(a) What are the benefits of reduced leadtimes in the system?

(b) What are the benefits of postponed geographic differentiation, i.e., of extending L while maintaining a constant total procurement leadtime $\tau = L + l$?

Questions (a) and (b) have been studied for traditional inventory systems with exogenously specified demand distributions, but not for settings with simultaneous inventory and price control.

Most of our study uses the national retailer data for the five locations described in Table 2. Where we investigate the impact of the *number* of retailers, we augment this quintuplet with the five locations described in Table 4, whose average c.v. value is similar to that of the first quintuplet of locations. Throughout this study, we assume that the error terms have a normal distribution.

For the item analyzed in this study, the total procurement leadtime, $\tau = L + l$, is four weeks. The company followed the practice of allocating systemwide orders *immediately* to the individual stores, i.e., $L = 0$ and $l = 4$. As discussed above, the advantages of coordinated replenishments can be increased to the extent the point of geographic differentiation can be postponed, i.e., when L can be increased and $l = 4 - L$ decreased. We therefore systematically consider the five possible integer partitions of the total procurement leadtime τ: (I) $L = 0$, $l = 4$, (II) $L = 1$, $l = 3$, (III) $L = 2$, $l = 2$, (IV) $L = 3$, $l = 1$, and (V) $L = 4$, $l = 0$.

We have evaluated a set of 109 instances, which include the 28 instances in Table 4. Of the remaining 81, 45 instances arise by considering all possible points of differentiation (I)–(V) in combination with five values for h^- (specifically, the base value $h^- = 6$ and $h^- = 8, 10, 12, 14$) and in combination with five vectors of c.v. values (specifically, the base c.v. values, c.v.^{0}, as well as 0.5c.v.^{0}, 0.75c.v.^{0}, 1.25c.v.^{0}, and 1.5c.v.^{0}). For the five possible points of differentiation, we evaluate the system with all 10 locations and the system with 5 fully identical retailers with equalized stochastic demand functions, as in §3—but 50% larger coefficients of variation—thus giving rise to another 10 instances. Fourteen instances investigate seven possible values for the total leadtime $\tau = 0, \ldots, 6$ for the systems with 5 and 10 retailers, always assuming $L = \lfloor \tau/2 \rfloor$. The remaining 12 instances investigate the impact of nonstationarities, assuming the demand distributions follow a periodic pattern with a cycle of four periods and seasonality factors of 1.5, 2.0, 1.25, and 1.0 for the first, second, third, and fourth week of each cycle, respectively. The seasonality factors apply to the intercepts $\{\eta_j\}$ only, i.e., we continue to assume that the slopes of the demand functions and hence the marginal price sensitivities remain stationary. Considering all 5 points of differentiation for $\tau = 4$ and 7 possible values for τ, with $L = \lfloor \tau/2 \rfloor$, we obtain a total of 12 instances.

For each instance, we have generated 3,000 replicas to obtain sufficiently narrow confidence intervals. We have observed that (H^D) and ($H^{PF,D}$) systematically outperform (H^A) and ($H^{PF,A}$), respectively. We therefore confine ourselves, henceforth, to the former pair of heuristics.

The average gap of the best heuristic vis-a-vis V^* is 1.56% and the median gap is 1.14%. Similarly, the average gap vis-a-vis V^{**} is 1.39% and the median gap is 1.15%. Heuristic ($H^{PF,D}$) usually outperforms or performs as well as (H^D), but in some cases (H^D) dominates. The latter instances tend to have small-order leadtimes L. Here, (H^D) has the upper hand because uncertainties about future price changes during the order leadtime have little or no impact on the proper specification of the order quantities, while the specification of the order quantities w^D, via the disaggregation procedure described in §7, is sufficiently powerful to adequately address any prevailing inventory imbalances among the retailers. Clearly, (H^D) has the additional advantage of allowing for fully dynamic price setting as opposed to the price-setting procedure under ($H^{PF,D}$). On the other hand, there are significantly more instances where ($H^{PF,D}$) dominates (H^D). We therefore recommend that in each instance the best of the two heuristics (H^D) and ($H^{PF,D}$) be implemented, and we conclude that V^* can be used as an accurate profit approximation. The proposed combined heuristic comes close to achieving the V^* or V^{**} value. Our experience indicates that the performance of the heuristics improves with the length of the planning horizon. In Figure 5 we show the average gap vis-a-vis V^*, for the best performing heuristic among all instances with a given possible integer partition of the total procurement leadtime $L + l = 4$. Note, the average gaps decrease as L increases except when L increases from zero to one. This decrease in the gap is not surprising: The approximate model relaxes the original system to a single-location model, and as L increases and l decreases, the original system becomes increasingly similar to a single-location system.

We now turn to the managerial questions listed at the beginning of this section. In carrying out various comparisons, we base these on the *simulated* expected performance measures associated with our proposed strategies, as opposed to the analytical approxima-

Figure 5 **Average Gap Between Approximate Model and Best Performing Heuristic**

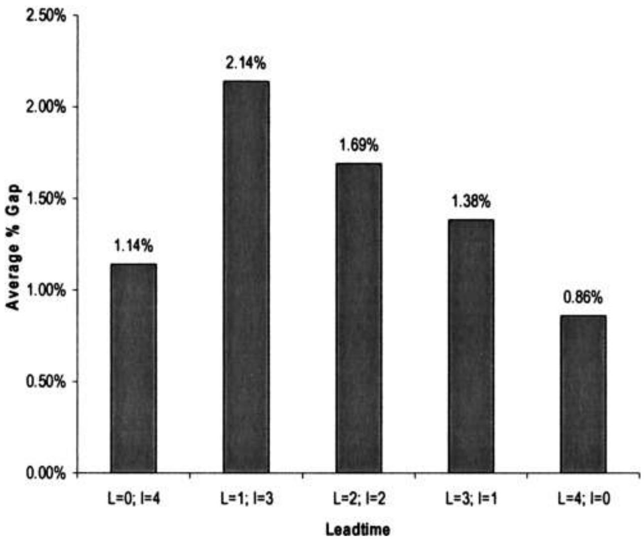

Figure 6 **Average Simulated Profit for Varying Values of Coefficient of Variation**

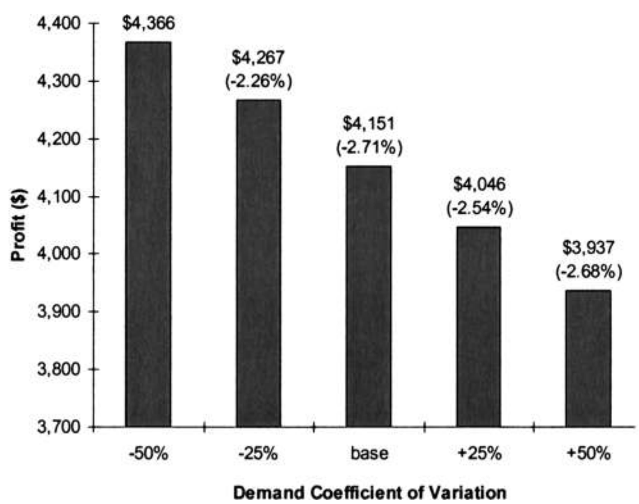

tions V^* or V^{**}, their demonstrated accuracy not withstanding.

Figure 6 exhibits how simulated profits decrease as the demand c.v. increases. Each of the bars represents the average profit values across all instances with the corresponding c.v. vector. The numbers within parentheses in Figure 6, as well as those in Figures 7 and

Figure 7 **Average Simulated Profit for Varying Degrees of Geographic Differentiation**

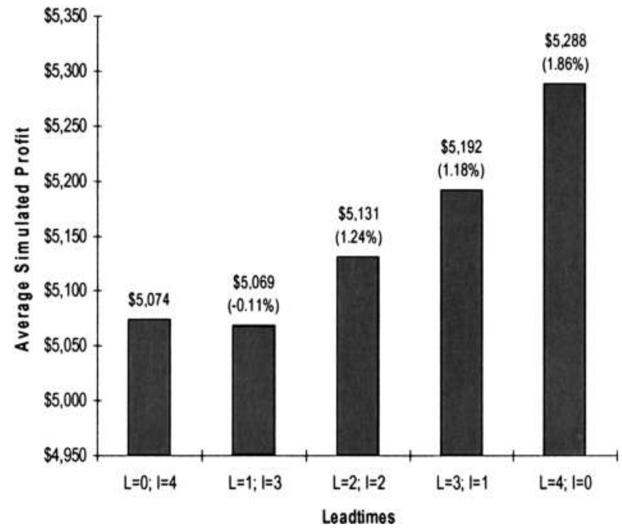

8, denote the percentage improvement of the average profit value when compared with the value associated with the bar to its left. The average percent decrease in profits is 2.55% when going from the lowest to the highest c.v. levels. Most of the decrease in profits results from increasing inventory-related costs due to increased safety stocks and increasing incidences of stockouts due to greater demand uncertainty. The average inventory cost is $186, $199, $364, $453, and $536 for the five c.v. levels, respectively, with an average increase from one c.v. value to the next of 33.22%.

Figure 7 shows, for systems with $\tau = L + l = 4$, how the systemwide profits increase as the point of geographic differentiation is extended. The figure displays the average across all five location instances with a given leadtime combination of both total profits and their percent improvements as L is increased by one unit at a time. The average percent improvement for a one-unit increase in L is 1.04%, again largely due to reduced inventory holding and backlogging costs.

Figure 8 exhibits the impact of *total* procurement leadtimes on total systemwide profits. (Again, each bar represents the average value across all relevant instances.) As one would expect, total systemwide

Figure 8 Average Simulated Profit for Varying Values of T

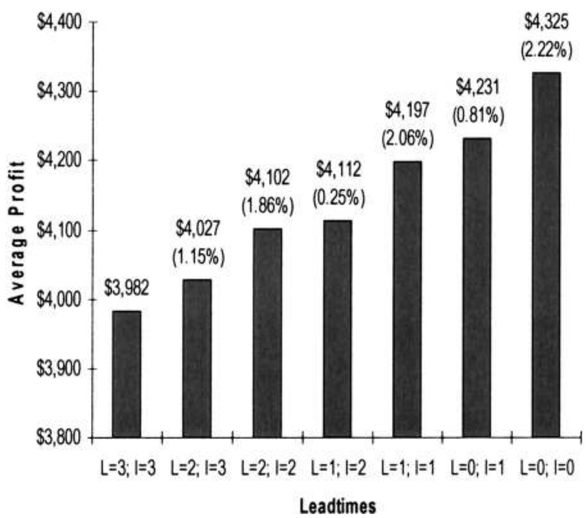

profits increase as the total procurement leadtime decreases. The average relative increase in total profits due to a one-week decrease in total procurement leadtime is 1.39%.

References

Agrawal, N., S. Smith. 1996. Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Res. Logist.* **43** 839–861.

Aviv, Y., A. Federgruen. 1999. Capacitated multi-item inventory systems with random and seasonally fluctuating demands: Implications for postponement strategies. *Management Sci.* **47** 512–531.

Axsater, S. 1993. Continuous review policies for multi-level inventory systems with stochastic demand. *Logistics of Production and Inventory,* ch. 4. S. Graves, A. Rinnooy Kan, and P. Zipkin, eds. *Handbooks in Operations Research and Management Science,* vol. 4, Elsevier, Amsterdam, The Netherlands.

Balvers, R.J., T.F. Cosimano. 1990. Actively learning about demand and the dynamics of price adjustment. *The Econom. J.* **100** 882–898.

Bitran, G.R., S.V. Mondschein. 1993. Pricing perishable products: An application to the retail industry. MIT Working Paper #3592-93, Cambridge, MA.

——, ——. 1997. Periodic pricing of seasonal products in retailing. *Management Sci.* **43** 64–79.

Braden, D.J., S.S. Oren. 1994. Nonlinear pricing to produce information. *Marketing Sci.* **13** 310–326.

Chan, L.M.A., D. Simchi-Levi, J. Swann. 2000. Flexible pricing strategies to improve supply chain performance. Working paper, Northwestern University, Chicago, IL.

Chen, F., Y.S. Zheng. 1994. Lower bounds for multi-echelon stochastic inventory systems. *Management Sci.* **40** 1426–1443.

——, A. Federgruen, Y.S. Zheng. 2001a. Coordination mechanisms for a distribution system with one supplier and multiple retailers. *Management Sci.* **47** 693–708.

——, ——, ——. 2001b. Near optimal pricing and replenishment strategies for a retail/distribution system. *Oper Res.* **49** 839–853.

Eppen, G.D., A.V. Iyer. 1997. Backup agreements in fashion buying—The value of upstream flexibility. *Management Sci.* **43** 1469–1484.

——, L. Schrage. 1981. Centralized ordering policies in a multiwarehouse system with leadtimes and random demand. L. Schwarz, ed. *Multi-Level Production/Inventory Control Systems: Theory and Practice.* North-Holland, Amsterdam, The Netherlands, 51–69.

Federgruen, A. 1993. Centralized planning models for multi-echelon inventory systems under uncertainty. S.C. Graves et al., eds. *Logistics of Production and Inventory, Handbooks in Operations Research and Management Science,* vol. 4. North Holland, Amsterdam, The Netherlands.

——, A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Oper. Res.* **47** 454–475.

——, ——. 2000. Multi-location combined pricing and inventory control. Working paper, Unabridged version, Columbia University, New York.

——, P. Zipkin. 1984a. Approximation of dynamic, multi-location production and inventory problems. *Management Sci.* **30** 69–84.

——, ——. 1984b. Computational issues in an infinite-horizon multiechelon inventory model. *Oper Res.* **32**(4) 818–836.

——, ——. 1984c. Allocation policies and cost approximation for multilocation inventory systems. *Naval Res. Logist. Quart.* **31** 97–131.

Fisher, M., A. Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Oper. Res.* **44** 87–99.

Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40** 999–1020.

——, ——. 1997. A multi-product dynamic pricing model with applications to network yield management. *Oper. Res.* **45** 24–41.

——, P.H. Zipkin. 1999. Stock positioning and performance estimation in serial production-transportation systems. *Manufacturing and Service Oper. Management* **1** 77–88.

Grossman, S., R.E. Kihlstrom, L.J. Mirman. 1977. A Bayesian approach to the production of information and learning by doing. *Rev. Econom. Stud.* **44** 533–547.

Kane, A. 2000. Private communication with authors.

Karlin, S., C.R. Carr. 1962. Prices and optimal inventory policy. K.J. Arrow, S. Karlin, and H. Scarf, eds. *Studies in Applied Probability and Management Science.* Stanford University Press, Stanford, CA, 159–172.

McLennan, A. 1984. Price dispersion and incomplete learning in the long run. *J. Econom. Dynamics Control* **7** 331–347.

Mills, E.S. 1959. Uncertainty and price theory. *Quart. J. of Econom.* **73** 116–130.

Newman, Kevin. 2000. Amazon accused of ''sneakiness''—Giant internet retailer admits to charging different customers different prices. ABCnews.com (September 27).

On the Web, price tags blur. 2000. *Washington Post* (September 27) A1.

Petruzzi, N.C., M. Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Oper. Res.* **47** 183–194.

Porteus, E.L. 1990. Stochastic inventory theory. D.P. Heyman and M.J. Sobel, eds. *Handbooks in OR and MS,* vol 2. Elsevier Science Publishers B.V., Amsterdam, The Netherlands, 605–652.

Rothschild, M. 1974. A two-armed bandit theory of market pricing. *J. Econom. Theory* **9** 185–202.

Stalk, G., P. Evans, L. Shulman. 1992. Competing on capabilities: The new rule of corporate strategy. *Harvard Bus. Rev.* 57–69.

Svoronos, A., P. Zipkin. 1988. Estimating the performance of multi-level inventory systems. *Oper. Res.* **36** 57–72.

Thisse, J., X. Vives. 1988. On the strategic choice of spacial price policy. *The Amer. Econom. Rev.* **78** 122–137.

Topkis. 1978. Minimizing a submodular function on a lattice. *Oper. Res.* **26** 305–321.

Wolverton, T. 2000. Pricing glitch wreaks havoc in Amazon DVD store. CNET News (September 12).

Young, L. 1978. Price, inventory, and the structure of uncertain demand. *New Zealand Oper. Res.* **6** 157–177.

Zipkin, P. 1980. Simple ranking methods for allocation of one resource. *Management Sci.* **26** 34–43.

The consulting Senior Editor for this manuscript was Evan Porteus. This manuscript was received on March 31, 2000, and was with the authors 843 days for 6 revisions. The average review cycle time was 57 days.