STOCHASTIC LEADTIMES IN A ONE-WAREHOUSE, N-RETAILER INVENTORY SYSTEM WITH THE WAREHOUSE CARRYING STOCK

Adriano O. Solis[(a)], Charles P. Schmidt[(b)]

[(a)] Management Science Area, School of Administrative Studies, York University, Toronto, Ontario M3J 1P3, Canada
[(b)] Department of Information Systems, Statistics and Management Science, The University of Alabama, Tuscaloosa, Alabama 35487-0226 U.S.A.

[(a)] asolis@yorku.ca, [(b)] cschmidt@cba.ua.edu

ABSTRACT
This study involves the introduction of stochastic leadtimes between the warehouse and retail sites, in place of the original deterministic leadtimes in an earlier reported multi-echelon inventory model. We had previously investigated analytically the effects of stochastic leadtimes on required retailer base stock levels when the warehouse does *not* carry stock (e.g., serves as a cross-dock point). The current paper examines the effects of stochastic leadtimes in the traditional situation where the warehouse *does* carry stock. The model becomes mathematically intractable, and simulation studies become appropriate. A heuristic involving simulation experiments is devised for selecting a base stock policy, taking off from solutions to the deterministic leadtime model. Resulting average system inventory is found equal or very close to the optimal level. Our study suggests that the optimal solution to the deterministic leadtime case provides an appropriate starting point in searching for a solution in the stochastic leadtime case.

Keywords: inventory management, multi-echelon inventory system, stochastic leadtimes, simulation

1. INTRODUCTION

Graves (1996) developed a multi-echelon inventory model assuming what he refers to as *virtual allocation*. Each of the sites—a central warehouse (CW), transshipment sites, and retail sites—follows a base stock (or "order-up-to") policy. In the two-echelon case involving only a warehouse and N retailers, Graves specifies search procedures for the optimal combination of base stock levels at the CW and the retail sites minimizing the average system on-hand inventory under each of two service level criteria: a probability of no stockout criterion and a fill rate criterion.

The average system inventory under an optimal policy where the CW base stock level is positive (i.e., the CW carries stock) is expected to be lower than under a policy where the CW does not carry stock. This is in view of "statistical economies of scale" made possible by carrying, at least in part, a single-system inventory instead of simply N individual inventories—which Eppen and Schrage (1981) refer to as the *depot effect*. Based on a limited computational study, Graves makes the assessment that the depot effect in his model "seems fairly small except when there are many retailers." Solis (1998) expanded Graves' computational study, and found the depot effect to be felt more the larger the number of retail sites or the smaller the retail site demand rate.

Solis and Schmidt (2007) earlier introduced into Graves' model stochastic leadtimes between the CW and the retail sites in place of the original deterministic leadtimes, but dealing only with the case where the CW does *not* carry stock (e.g., serves as a cross-dock or transshipment point). Effects of stochastic leadtimes on required base stock levels at the retail sites were investigated analytically, taking two alternative treatments of stochastic leadtime distributions into consideration. Results of that earlier study suggest that it may be better to use the deterministic leadtime model with an accurately estimated mean leadtime than a stochastic model with a poorly estimated mean leadtime.

In the current study, the CW, in the traditional sense, *does* carry stock. The mathematics becomes intractable, and simulation studies are undertaken.

2. GRAVES' MODEL

The model involves an arborescent system with M inventory sites, i = 1, 2, ..., M. Each site j has a single internal supplier $i = \rho(j)$, with the exception of site 1, a central warehouse (CW) whose inventory is replenished by an external supplier. Customer demand occurs only at retail sites, which have no successor nodes. All other sites are storage and/or consolidation facilities, called transshipment sites. The unique path linking a retail site to the CW is the *supply chain* for the retail site.

The analysis involves a single item of inventory. The demand at each retail site j is an independent Poisson process with demand rate λ_j. $D_j(s,t)$ represents the demand over the time interval (s,t] for site j. The induced demand rate for any site other than the retail sites is the sum of demand rates at its immediate successors. For transshipment sites as well as the CW, $D_i(s,t)$ is the sum of the demands over the interval (s,t] at the sites supplied by site i (i.e., at the immediate successors of site i).

A multi-item distribution system is contemplated, with regularly scheduled shipments between sites. Each shipment is a consolidation of orders for various items in inventory, including the item under study. A schedule of preset times $p_j(m)$, m = 1, 2, ..., at which site j places its m^{th} replenishment order on its supplier, is followed. Positive leadtimes τ_j for shipments to each site j from its supplier are preset and known. The m^{th} shipment to site j is received at time $r_j(m) = p_j(m) + \tau_j > p_j(m)$. With fixed leadtimes, no order crossing will occur. Hence, $p_j(m) < p_j(m+1)$ and $r_j(m) < r_j(m+1)$. When inventory is in short supply, the supplier will ship less than the quantity ordered and make up for the shortfall on later shipments.

Each site j follows a base stock (or "order-up-to") policy. Initial inventory (at time 0) at site j is B_j, which is the base stock level for site j. Since customer demand is assumed to be fully backordered, site j places an order equal to $D_j[p_j(m-1), p_j(m)]$ at time $p_j(m)$.

$T_j(m)$ represents the *coverage* provided by the supplier to site j, on the occasion of the m^{th} order by site j. A quantity equal to $D_j[T_j(m-1), T_j(m)]$ is shipped by the supplier at time $p_j(m)$, which will later arrive at site j at time $r_j(m)$. If $T_j(m) < p_j(m)$, then the quantity $D_j[T_j(m), p_j(m)]$ remains on backorder. If $T_j(m) = p_j(m)$, then the supplier is able to fill the entire m^{th} order placed by site j. The external supplier is fully reliable and fills every order by the CW (site 1) exactly as scheduled, following a leadtime τ_1. In this case, $T_1(m) = p_1(m)$.

Graves' model differs from other existing models with its assumption of *virtual allocation*. Each site on the supply chain increases its next order quantity by one whenever a unit demand occurs at the retail site. At the same time, each site on the supply chain commits/reserves one unit of its inventory, if available, for shipment to the downstream site on the latter's next order occasion. Priority in the allocation of uncommitted stock is set according to the earliest demand occurrence. While implementing virtual allocation under current information technology is possible, it is not the common practice. It is assumed in the model because it proves to be tractable. It is found by Graves to be near-optimal in many cases.

A random variable requiring attention in this study is $A_j(t)$, which denotes the *available inventory at site j at time t*—the on-hand inventory at site j at time t that has not yet been committed for shipment to another site. $A_j(t) < 0$ indicates outstanding orders (or backorders) at site j at time t. The beginning inventory at site j at time t = 0 is the base stock level B_j. Hence, we have $A_j(0) = B_j$. This leads Graves to the following important result:

<u>Theorem</u> *If* $r_j(m) \le t < r_j(m+1)$, *then*

$$A_j(t) = B_j - D_j[T_j(m),t]. \qquad (1)$$

Suppose that site i is the internal supplier to site j. At time $p_j(m)$, site j places its m^{th} order with site i for a quantity equal to $D_j[p_j(m-1), p_j(m)]$. Since $T_j(m)$ represents the coverage provided by the supplier on the occasion of the m^{th} order by site j, then site i ships to site j at time $p_j(m)$ a quantity equal to $D_j[T_j(m-1), T_j(m)]$, which will arrive at site j at time $r_j(m)$. If site i has sufficient stock and is able to fill the entire m^{th} order placed by site j, then $T_j(m) = p_j(m)$. Otherwise, $T_j(m) < p_j(m)$. In the latter case, $T_j(m)$ equals the time at which site i would run out of available inventory to allocate to site j (i.e., the time when a demand occurrence at some site k, which may or may not be site j, supplied by site i reduces available inventory at site i to zero).

Consider the *relevant* shipment to the supplier i, for which $r_i(n) \le p_j(m) < r_i(n+1)$. At time $p_j(m)$, site i has received its n^{th} shipment, but not yet its $(n+1)^{th}$ shipment. Define $S_i(n)$ to be the *depletion* or *runout time* for this relevant (n^{th}) shipment to site i. That is, based upon its receipt of the n^{th} shipment, site i is able to cover (or replenish) the demand processes of its successor sites up through the runout time $S_i(n)$. If $S_i(n)$ occurs after $p_j(m)$, then $T_j(m) = p_j(m)$. However, if $S_i(n)$ occurs before $p_j(m)$, then $T_j(m) = S_i(n)$. Thus,

$$T_j(m) = \min \{p_j(m), S_i(n)\}. \qquad (2)$$

The base stock B_i takes the role of a buffer for demand at site i after the coverage time $T_i(n)$ provided by the supplier to site i. This buffer B_i is depleted at the runout time $S_i(n)$. Graves refers to the difference $S_i(n) - T_i(n)$ as the *buffer time* provided by B_i, and establishes that

$$S_i(n) - T_i(n) \sim \text{gamma}(\lambda_i,B_i). \qquad (3)$$

Graves then focuses on a two-echelon system consisting of sites 1 (the CW) and j (the retail sites), reducing the supply chain of interest to just two sites, 1 and j. A *single-cycle* ordering policy is in place: each retail site j orders a fixed number of times for every order placed by the CW. Hence, if θ_1 denotes the length of the CW order cycle and θ_j that of the retail site, then the ratio θ_1/θ_j is a positive integer. The ordering policy is assumed to be *nested*: every time the CW receives a shipment, all retail sites place an order.

Consider an arbitrary (n^{th}) CW order cycle. Graves simplifies the analysis by setting time zero equal to $p_1(n)$, the time at which the n^{th} CW order is placed with the external supplier. Assuming the CW receives delivery of this order at the end of leadtime τ_1, the retail site orders at time τ_1 (which signals the start of the full CW order cycle). Graves draws attention to the last retail site order—i.e., the $(\theta_1/\theta_j)^{th}$ order—within the CW order cycle. The last order within this CW order cycle is placed by retail site j at time $p_j = \tau_1 + \theta_1 - \theta_j$, and is received by the retail site at time $p_j + \tau_j$, where τ_j is the leadtime between the CW and retail site j. The resulting available inventory will be used to cover demand until the next order (placed at time $p_j + \theta_j$) arrives at the retail site (at time $t_r = p_j + \theta_j + \tau_j$, which is equal to $\tau_1 + \theta_1 + \tau_j$). The instant of time t_r^- *before* this replenishment at $t_r = \tau_1 + \theta_1 + \tau_j$ proves crucial to the analysis. [In this case, the antecedent of the previous Theorem, i.e., $r_j(m) \le t_r^- < r_j(m+1)$, holds.]

Treating this $(\theta_1/\theta_j)^{th}$ order as the m^{th} order for the retail site j within the n^{th} CW order cycle, the indices m and n are henceforth dropped for notational convenience. Rewriting (2), the coverage provided by the m^{th} shipment to retail site j is given by $T_j = \min\{p_j, S_1\}$, where p_j = the time at which the (m^{th}) order was placed by retail site j, and S_1 is the runout time for the relevant (n^{th}) shipment to the CW. Having set time zero equal to p_1 for convenience, and assuming the external supplier to the CW to be reliable (so that $T_1 = p_1$), we have $S_1 - T_1 = S_1$. Hence, in this case, the runout time S_1 is equal to the buffer time $S_1 - T_1$ provided by the base stock B_1 at the CW. It follows from (3) that $S_1 \sim \text{gamma}(\lambda_1, B_1)$.

Graves establishes the first two moments of T_j, and uses these to specify the mean and variance of the random variable $D_j[T_j,t]$, where t is some specified point in time. Of particular interest is the distribution of $D_j[T_j,t_r]$, where $t_r = \tau_1 + \theta_1 + \tau_j$ is the critical point in time discussed above. Graves refers to $D_j[T_j,t_r]$ as *uncovered demand* (up to time t_r): demand at retail site j not covered by the $(\theta_1/\theta_j)^{th}$ shipment from the CW.

Graves reports to have computationally found the negative binomial distribution, having the same first two moments as $D_j[T_j,t]$, to provide a fairly accurate approximation to the distribution of $D_j[T_j,t]$, while presenting little evidence in support of his assertion. He cites two earlier multi-echelon inventory studies (Graves 1985; Lee and Moinzadeh 1987) in which a negative binomial approximation had also been found effective. Solis, Schmidt, and Conerly (2007) provided a mathematical analysis of the effectiveness of the approximation in the current model, where the maximum absolute deviation between cumulative probabilities of the approximate and exact distributions reach about 0.01 in one of 64 test cases, but much less in most others.

Based upon Graves' specification that the negative binomial distribution used to approximate the distribution of $D_j[T_j,t]$ has the same first two moments as the latter, the discrete density function (d.d.f.) of this approximate distribution may be characterized by $f(x;r,p) = {}_{r+x-1}C_x \, p^r (1-p)^x I_{\{0,1,2,\ldots\}}(x)$, with parameters

$$r = (t - E[T_j])^2 / \mathrm{Var}[T_j] \quad (4)$$

and

$$p = (t - E[T_j]) \,/\, \{(t - E[T_j]) + \lambda_j \mathrm{Var}[T_j]\}. \quad (5)$$

When $\mathrm{Var}[T_j] \to 0$, however, the distribution of $D_j[T_j,t]$ approaches a Poisson distribution with parameter $\lambda_j(t - E[T_j])$.

2.1. Probability of No Stockout as Service Criterion

Each time the retail site places an order with the CW, there is a risk of stockout. This risk exists especially towards the end of each retail order cycle, just before the next shipment to the retail site is received.

Graves notes that, for a nested, single-cycle ordering policy, the probability of the retail site stocking out is greatest for the last order—that is, the $(\theta_1/\theta_j)^{th}$ order—within the CW order cycle. Thus, to set the base stock levels to achieve a given probability α of the retail site *not* stocking out within the CW order cycle, it suffices to consider the probability of stockout at the $(\theta_1/\theta_j)^{th}$ order occasion. Based on the earlier discussion,

$$\Pr\{ A_j(t_r^-) \ge 0 \} \ge \alpha, \quad (6)$$

where $t_r = \tau_1 + \theta_1 + \tau_j$, will need to be assured.

This leads to a computational procedure that searches over possible settings of the base stock level B_1 at the CW. For each B_1, the minimum base stock level B_j at the retail sites that would yield (6) is to be determined. Following (1), the requirement (6) translates into

$$\Pr\{ D_j[T_j,t_r] \le B_j \} \ge \alpha. \quad (7)$$

The negative binomial distribution, negbin(r,p), with parameters r and p as specified by (4) and (5), is used as approximation to the distribution of $D_j[T_j,t_r]$. The probability $\Pr\{D_j[T_j,t_r] \le B_j\}$ is approximated by $\sum_{x=0}^{B_j} {}_{r+x-1}C_x \, p^r (1-p)^x$, starting with $B_j = 1$, and incrementing B_j by 1 until (7) is satisfied. [Graves did not specify at what level of $B_1 > 0$ the entire search could be allowed to terminate. Solis (1997) established such a stopping condition.]

The base stock level B_1 which yields the lowest average system inventory is finally selected. (In the case of ties, the smallest value of B_1 is preferred—there being no difference assumed in Graves' model between holding costs at the CW and at the retail sites.)

Graves provides the following approximation to expected system on-hand inventory:

$$\text{average inventory} = B_1 + \sum_{1}^{N} B_j - 0.5\lambda_1\theta_1 - \lambda_1\tau_1. \quad (8)$$

He clarifies that (8) should actually be corrected for counting retail backorders at negative inventory—by adding back the time-weighted backorders at the retail sites. However, he points out that, for reasonable service levels, the expected backorder component is very small and insensitive to the inventory policy. Accordingly, this expected backorder component is ignored in using (8) to calculate average system on-hand inventory.

2.2. Fill Rate as Service Criterion

In the case of a fill rate criterion, Graves notes that:

- For a single-cycle ordering policy, the number of backorders over the CW order cycle is equal to the sum of the backorders at the end of each of the θ_1/θ_j retail order cycles, just before receipt of the next shipment to the retail site. (Graves once again points out that this is not completely accurate as it double counts any backorders that may persist for more than one retail cycle.)
- For a nested, single-cycle ordering policy, the expected backorders will be greatest for the last retail order—that is, the $(\theta_1/\theta_j)^{th}$ order—within the CW order cycle.
- Expected backorders over a CW order cycle may be approximated by expected backorders pertaining to the $(\theta_1/\theta_j)^{th}$ retail order within the CW order cycle, since for "realistic" fill rates (> 0.95) effectively all of the backorders occur at this last retail order within the CW order cycle.

We recall that the last order within the CW order cycle is placed by the retail site at time $p_j = \tau_1 + \theta_1 - \theta_j$. This last retail order within the CW order cycle (which arrives at the retail site at time $p_j + \tau_j = \tau_1 + \theta_1 - \theta_j + \tau_j$) may not last until the next order arrives at the retail site (at time $t_r = \tau_1 + \theta_1 + \tau_j$). Backorders would occur if the base stock B_j at the retail site is inadequate.

As already noted, for "realistic" fill rates, we approximate the expected backorders over the CW order cycle by evaluating expected backorders at $t_r = \tau_1 + \theta_1 + \tau_j$.

$E[\{A_j(t_r)\}^-]$, where the symbol y^- stands for $\max\{0,-y\}$, represents expected backorders at time $t_r = \tau_1 + \theta_1 + \tau_j$ (just before the next order arrives). A computational procedure similar to that for the probability of no stockout service criterion arises. For each B_1, we search for the minimum base stock level B_j at the retail sites that would yield

$$E[\{A_j(t_r)\}^-] = \left\{ \sum_{x=0}^{B_j} [(B_j - x)\, f(x)] \right\} - \left\{ B_j - \lambda_j\, (t_r - E[T_j]) \right\} \leq (1 - \beta)\, \lambda_j\theta_1, \qquad (9)$$

where $\lambda_j\theta_1$ represents mean demand at the retail site over the CW order cycle. The CW base stock level B_1 which yields the lowest average system inventory is chosen.

2.3. Graves' Computational Study

In his computational study, Graves used test scenarios all based on a single system demand rate $\lambda_1 = 36$. Identical retail sites are assumed, with the number N of retail sites being 2, 3, 6, or 18. Hence, the retail site demand rates λ_j are 18, 12, 6, or 2, respectively. The length of the retail site order cycle is fixed at $\theta_j = 1$. Four different parameter combinations $< \theta_1, \tau_1, \tau_j >$, involving the length of the CW order cycle (θ_1), the leadtime (τ_1) from the external supplier to the CW, and the leadtime (τ_j) from the CW to the retail site, are tested. This resulted in *16 test scenarios*, summarized in Table 1.

Table 1: Summary of Graves' Test Scenarios

Scenario	Length of CW Order Cycle θ_1	External Supplier Leadtime τ_1	CW to Retailer Leadtime τ_j	No. of Retailer Sites N	Retail Site Demand Rate λ_j
1	2	1	1	18	2
2				6	6
3				3	12
4				2	18
5	2	1	5	18	2
6				6	6
7				3	12
8				2	18
9	5	4	1	18	2
10				6	6
11				3	12
12				2	18
13	5	4	5	18	2
14				6	6
15				3	12
16				2	18

For the probability of no stockout service criterion, four different levels of α were used: 0.80, 0.90, 0.95, and 0.975. Similarly, four different fill rate levels β were tested: 0.95, 0.98, 0.99, and 0.999. Thus, for each service criterion, a total of 64 test cases were utilized by Graves.

3. STOCHASTIC LEADTIMES

Graves (1996) suggested a number of possible extensions of his model, including stochastic leadtimes τ_j from the CW to the retail sites. The current study investigates effects of stochastic leadtimes τ_j from the CW to the retail sites when the CW carries stock ($B_1 > 0$).

The situation where τ_j is no longer preset and known is a departure from the original deterministic leadtime τ_j in Graves' model, as in many others in the literature on multi-echelon inventory systems—and, particularly more so, from the zero leadtime assumed in some other models (McGavin, Schwarz, and Ward 1993; Nahmias and Smith 1994). A zero leadtime may be plausible where regularly scheduled deliveries from the CW are made overnight while the retail sites are closed. Graves (1996) likewise assumes regularly scheduled shipments. He offers the motivation that, in a multi-item distribution system where each item occupies only a portion of a truckload, a fixed replenishment schedule allows consolidation of item shipments and, accordingly, transportation economies. This scenario may appear to allow a fixed, common positive leadtime for shipments between the CW and the retail site—where one truck services one retail site within the latter's order cycle.

While the order occasion $p_j(m)$ remains fixed, the time $r_j(m)$ at which the order is received by site j may vary. Hence, the leadtime $\tau_j = r_j - p_j$ may be treated as a random variable that varies according to a number of possible factors, for example, truck, road, and weather conditions, or even loading times for the consolidated, multi-item shipments. The realized time of receipt r_j and, as a result, the leadtime τ_j do not depend on the demand process or the quantity ordered of the specific item under consideration. For Graves' original results to continue to apply in the stochastic leadtime case, the condition that orders do not cross—that is, $r_j(m-1) \le r_j(m)$ for any m—will need to hold.

The treatment here of τ_j is based on $\theta_j = 1$ (i.e., a retail site order cycle length equal to one time unit), as used in all 16 test scenarios in Graves' computational study, but may be extended to any possible value of θ_j. We wish to ensure that the distribution of leadtimes τ_j satisfies the requirement that there is "no order crossing"—specifically, that $r_j(m) < r_j(m+1)$ for any m. This requires the range of leadtimes (maximum leadtime less minimum leadtime) to be at most θ_j. In our treatment, therefore, we require this range of leadtimes to be at most 1. Such a requirement is less restrictive (with respect to the variability of leadtimes) when the mean leadtime is small relative to θ_j, and becomes more restrictive otherwise.

In his computational study, Graves used the fixed leadtimes $\tau_j = 1$ and $\tau_j = 5$. Corresponding to Graves' deterministic leadtime $\tau_j = 1$, we may, for instance, consider leadtimes that vary within the interval between 0.5 and 1.5, or between 0.25 and 1.25, or between 0.9 and 1.9, among other such possibilities for which the range is 1 (equal to θ_j) and $\tau_j = 1$ is included between the minimum and maximum leadtimes. [The range may as well be 0.5 or 0.9, or some other value $< \theta_j$. A range of 1 is, in the case where $\theta_j = 1$, the largest that would be allowed under the no order crossing requirement.]

The lower and upper limits of this interval of leadtimes would depend upon what may actually provide a good approximation to the practice/experience with respect to, among other factors, the consolidation of shipments, the dispatching of delivery vehicles from the CW, and the return of the vehicles to the CW. Since the shipments are made from the CW to the retail sites, an assumption that such shipments are made under a fair degree of control would appear to be a reasonable one to make, in view of the operational requirements of period-to-period consolidated, multi-item shipments—even as the actual leadtimes are subject to some amount of variation due to factors such as loading/ unloading times, weather and road conditions, and unexpected vehicular problems. Hence, we may reasonably expect that leadtime intervals will not be dispersed too widely and, as a result, that orders will not cross.

In our analysis, we shall use a leadtime interval of (0.5,1.5) arbitrarily, in place of the deterministic leadtime $\tau_j = 1$. The analysis will essentially proceed in the same manner regardless of the final interval that actually applies. Similarly, the interval (4.5,5.5) is used in place of the deterministic leadtime $\tau_j = 5$. Likewise, regardless of the original deterministic leadtime (be it $\tau_j = 1$ or $\tau_j = 5$), the analysis will proceed in the same manner. Within this interval of (0.5,1.5) or (4.5,5.5), we shall look into various stochastic leadtime distributions and their effects on required base stock levels.

Corresponding to Graves' fixed leadtime $\tau_j = 1$, we consider stochastic leadtimes such that $\tau_j - 0.5 \sim$ beta(a,b), where the p.d.f. of τ_j is specified by

$$g(\tau_j) = [1/B(a,b)]\, (\tau_j - 0.5)^{a-1}\, (1.5 - \tau_j)^{b-1}\, I_{(0.5,1.5)}(\tau_j), \quad (10)$$

with the so-called beta function $B(a,b) = \Gamma(a)\times\Gamma(b)/\Gamma(a+b)$. A random variable with a beta(a,b) distribution varies between 0 and 1, and has a mean of a/(a+b). This distribution is skewed to the right or left depending upon whether a/(a+b) < 0.5 or a/(a+b) > 0.5. The mean of τ_j, when $\tau_j - 0.5 \sim$ beta(a,b), is equal to a/(a+b) + 0.5 and can be anywhere on the interval (0.5,1.5). In this case, the distribution of leadtimes τ_j would be skewed to the right or left when the mean of τ_j is less than 1 or greater than 1, respectively. [When a = b = 1, the distribution of τ_j reduces to a uniform distribution over the interval (0.5,1.5).]

3.1. The Case Where the CW Actually Carries Stock

When the CW is allowed to carry some of the system inventory, the average system inventory associated with the optimal combination $< B_1, B_j >$ is smaller than the required average system inventory when $B_1 = 0$. Eppen and Schrage (1981) refer to this benefit of the CW carrying stock as the *depot effect*.

Solis and Schmidt (2007) earlier investigated analytically the effects of stochastic leadtime distributions as specified above for the case where the CW does *not* carry stock ($B_1 = 0$). When the CW *does* carry some of the system inventory ($B_1 > 0$), the model does not readily lend itself to the same kind of mathematical analysis.

Accordingly, we conducted simulation studies in order to evaluate the effects of stochastic leadtimes τ_j on the optimal combination $< B_1, B_j >$ that minimizes average system inventory. Baseline simulation models were developed in SIMAN V using the Arena modeling system (see Pegden, Shannon, and Sadowski 1995; Kelton, Sadowski, and Sturrock 2007), for both deterministic and stochastic leadtimes. Preliminary runs indicated that service levels under given $< B_1, B_j >$ combinations may be significantly affected when deterministic leadtimes between the CW and the retail sites (as in Graves' model) are replaced with leadtimes following uniform and various beta distributions.

The simulation experiments were envisioned to study how optimal levels of B_1 and B_j, determined for the fixed leadtime case, would need to be adjusted to compensate for stochastic leadtimes—such that the specified service level is met under either the probability of no stockout criterion or the fill rate criterion.

For each of the test scenarios, experimental and model frames were developed for both the deterministic leadtime and stochastic leadtime cases. Most of the experiments conducted consisted of 100 replications, with each replication involving 100 CW order cycles. Hence, a typical experiment covered 10,000 CW order cycles.

The Seeds element in SIMAN (Pegden, Shannon and Sadowski 1995) was used to specify separate random number streams for the various sources of randomness—the independent Poisson demand processes at the retail sites, and the stochastic leadtimes τ_j. Thus, experiments with the same test case become comparable, with random demand occurrences and random leadtimes repeating themselves with every run of the experiment, as values of the parameters B_1 and B_j are changed.

Confidence intervals for system service levels (average probability of no stockout and system fill rate)—as reported by Arena's statistical analysis module, which rounds reported figures off to the nearest 0.1—vary in width, and are generally tighter with higher service levels and more retail sites.

3.2. Heuristic for Finding a Policy $< B_1, B_j >$ that Minimizes Average System Inventory

We have devised a heuristic for finding a policy $< B_1, B_j >$ which yields an average system on-hand inventory that appears to closely approximate the smallest possible level of inventory such that the specified service level is satisfied.

From (8), an approximation to expected system inventory when there are N identical retail sites is given by average system inventory = $B_1 + NB_j - 0.5\lambda_1\theta_1 - \lambda_1\tau_1$, which varies only with B_1 and B_j. Hence, to find the minimum average system inventory, we may simply search for the minimum level of echelon base stock = $B_1 + NB_j$.

Let us first introduce additional notation: $e\{< B_1, B_j >\}$ = echelon base stock at $< B_1, B_j >$ = $B_1 + NB_j$; $s\{< B_1, B_j >\}$ = mean service level observed from a simulation experiment involving policy $< B_1, B_j >$; and $< B_1^d, B_j^d >$ = optimal policy for the deterministic leadtime case, found using Graves' search procedure.

Let $< B_1^h, B_j^h >$ represent the current policy as the heuristic searches for an approximation to the optimal policy $< B_1^s, B_j^s >$ for the given stochastic leadtime case—i.e., a policy that minimizes echelon base stock (and, hence, average system inventory). We further let $\delta = B_j^s(0) - B_j^d(0)$—where $B_j^d(0)$ and $B_j^s(0)$ denote the computed optimal retailer base stock levels, if the CW does not carry stock, for the deterministic and stochastic leadtime cases, respectively. The values of $B_j^s(0)$ were earlier determined by Solis and Schmidt (2007) for the same test cases we consider in the current study.

We specify our heuristic as follows:

<u>*Heuristic*</u>

Step 1. Initialize: $< B_1^h, B_j^h > \leftarrow < B_1^d, B_j^d + \delta >$. If $s\{< B_1^h, B_j^h >\} > \alpha$ or β, then go to step 3.

Step 2. Adjust, via simulation experiments, B_j^h upward until $< B_1^h, B_j^h >$ first satisfies $s\{< B_1^h, B_j^h >\} > \alpha$ or β.

Step 3. Adjust B_1^h downward, via simulation experiments (using, if necessary, some bisection method), until such point that any further reduction in B_1^h will result in $s\{< B_1^h, B_j^h >\} < \alpha$ or β. Select the policy $< B_1^h, B_j^h >$ corresponding to the smallest B_1^h considered such that $s\{< B_1^h, B_j^h >\} > \alpha$ or β.

We express the condition that the specified service level be satisfied in terms of $s\{< B_1^h, B_j^h >\} > \alpha$ or β—instead of using the more familiar ≥ inequality—in light of the simulation experiment. This is in recognition of $\Pr[\, s\{< B_1^h, B_j^h >\} = \alpha \text{ (or } \beta) \,] = 0$.

A motivation for step 1 (the initialization of $< B_1^h, B_j^h >$) is our intuition—which we are unable to establish mathematically—with respect to δ. On the one hand, if shifting from deterministic to stochastic leadtimes τ_j leads to an increase ($\delta > 0$) in the required level of B_j for the case where the CW does not carry any stock ($B_1 = 0$), our expectation is that the required increase in B_j (for given $B_1 > 0$)—when the CW does carry some of the system inventory—must be *at least* equal to δ for the given service level to be satisfied. This would seem to make sense because, when $B_1 > 0$, the retail site appears to be at greater risk of experiencing stockouts/backorders. On the other hand, if the stochastic leadtimes result in a decrease ($\delta < 0$) in the required B_j for the case where $B_1 = 0$, we would expect the magnitude of the corresponding decrease in B_j (for given $B_1 > 0$) to be *at most* equal to $|\delta|$.

We provide here a simple case study to illustrate application of the heuristic.

<u>*Case Study*</u>

Consider one warehouse supplying three retail stores selling, among others, a certain large screen TV model. Demand at each retail store is independent of demand at every other retail store, averaging 12 units per week and observed to essentially follow a Poisson distribution. The warehouse places replenishment orders with the electronics manufacturer every two weeks, with a guaranteed delivery leadtime of one week. Each retail site places a replenishment order once a week. The average delivery leadtime from the warehouse to each retail site averages one week from the time the order was placed, but has been found to vary between 0.5 and 1.5 weeks, more or less according to a beta(6,2) distribution. The desired probability of stockout is not more than 5%.

The above illustrative case leads to the following computational example:

Example 1

Given N = 3, $\lambda_j = 12$, $\theta_1 = 2$, $\theta_j = 1$, $\tau_1 = 1$, and $E[\tau_j] = 1$, we have scenario 3 (refer to Table 1). Graves (1996) reported $< B_1^d, B_j^d >$ = <56, 39> as the optimal base stock policy, with echelon base stock = 56 + (3 × 39) = 173. However, if the warehouse does not carry stock, $B_j^d(0) = 60$ and echelon base stock = 3 × 60 = 180. With $\alpha = 0.95$ and $\tau_j - 0.5 \sim$ beta(6,2), Solis and Schmidt (2007) have reported $B_j^s(0) = 63$, leading to an echelon base stock of 3 × 63 = 189. Thus, $\delta = B_j^s(0) - B_j^d(0) = 63 - 60 = 3$.

Step 1. Initialize: $< B_1^h, B_j^h > \leftarrow < B_1^d, B_j^d + \delta >$ = <56, 42>, with echelon base stock = 173 + 3δ = 182. s{<56, 42>} = 92.8% < α.

Step 2. $B_j^h \leftarrow B_j^h + 1 = 43$. s{<56, 43>} = 94.8% < α. $B_j^h \leftarrow B_j^h + 1 = 44$. s{<56, 44>} = 96.2% > α. Echelon base stock at <56, 44> is 182 + 6 = 188.

Step 3. In adjusting $B_1^h = 56$ downward, we find that s{<54,44 >} = 95.2% > α, while s{<53, 44>} = 94.7% < α. $< B_1^h, B_j^h > \leftarrow$ <54, 44>, with echelon base stock = 188 - 2 = 186.

The average observed probability of no stockout at the selected policy <54, 44>—based on our typical simulation experiment consisting of 100 replications—does not lead to a statistically significant conclusion that the true mean probability is, indeed, no less than 95%. [We created a separate file (using the Files element in SIMAN) in a worksheet file structure, and used Excel to compute statistics (and report them beyond Arena's one digit after the decimal point).] At the policy <54, 44>, we computed a mean of 95.183(%), with a standard deviation of 1.791. The t statistic for a right hand-tailed hypothesis test, given a sample size of 100, is only 1.024 (which is less than the critical value of 1.6604 at a 5% level of significance). When we expanded our experiment to 400 replications, we observed a mean of 95.186 and a standard deviation of 1.645. The computed t statistic of 2.259 exceeds the critical value.

Table 2 presents figures relevant to our illustration, which will allow us to evaluate how well the heuristic performs, in this case, relative to the "true" minimum level of echelon base stock. Since echelon base stock at the selected policy <54, 44> is 186 units, we now focus our attention on the next lower level of echelon base stock (185 units). Table 2 enumerates all possible (integer valued) policies $< B_1, B_j >$ such that $e\{< B_1, B_j >\} = B_1 + 3B_j$ = 185 or 186, and such that $B_j \geq B_j^{min} = 36$. B_j^{min} is a lower bound on B_j, determined by way of a minor heuristic developed by Solis (1997). Below B_j^{min}, the specified service level α will definitely *not* be satisfied. The average observed probability of no stockout $s\{< B_1, B_j >\}$, based on our typical experiment (consisting of 100 replications involving 100 CW order cycles per replication), is shown for each of the listed policies. Moreover, for some policies of interest, 95% confidence intervals for the mean probability of no stockout are shown below the values of $s\{< B_1, B_j >\}$. [Arena outputs only two-sided confidence intervals. Figures with one digit after the decimal point are as reported using Arena. Those showing two decimal places, in the case of certain policies of interest, are as computed using Excel.]

Table 2: Average Observed Probabilities of No Stockout (%) – Example 1

Echelon Base Stock = 185			Echelon Base Stock = 186		
B_1	B_j	Average Observed Prob of No Stockout (%)	B_1	B_j	Average Observed Prob of No Stockout (%)
			0	62	92.9
2	61	92.4	3	61	93
5	60	92.4	6	60	93.1
8	59	92.5	9	59	93.3
11	58	92.7	12	58	93.4
14	57	92.9	15	57	93.5
17	56	93	18	56	93.7
20	55	93.2	21	55	93.9
23	54	93.3	24	54	94
26	53	93.5	27	53	94.1
29	52	93.6	30	52	94.3
32	51	93.8	33	51	94.3
35	50	93.8	36	50	94.5
38	49	94	39	49	94.69
41	48	94.1	42	48	94.69
		(93.8, 94.5)			
44	47	94.2	45	47	94.82
		(93.9, 94.6)			
47	46	94.4	48	46	94.99
		(94.0, 94.8)			
50	45	94.48	51	45	95.05
		(94.10, 94.86)			
53	44	94.66	54	44	95.18
		(94.28, 95.03)			(94.83, 95.53)
56	43	94.76	57	43	95.31
		(94.40, 95.12)			
59	42	94.77	60	42	95.29
		(94.42, 95.13)			
62	41	94.79	63	41	95.3
		(94.41, 95.17)			
65	40	94.57	66	40	95.04
		(94.16, 94.98)			
68	39	94.2	69	39	94.62
		(93.8, 94.6)			
71	38	93.5	72	38	93.9
		(93.1, 93.9)			
74	37	92.3	75	37	92.5
77	36	90.4	78	36	90.6

In Example 1, we note that the upper limit of the 95% confidence interval exceeds the specified 95(%) service level at a number of policies $< B_1, B_j >$ for which $e\{< B_1, B_j >\}$ = 185—specifically, at the policies <53, 44>, <56, 43>, <59, 42>, and <62, 41>. Using Excel, we compute for <53, 44> a 94.657 average probability of no stockout, with a standard deviation of 1.911. The resulting t statistic for a left hand-tailed test of hypothesis on the mean is -1.797, which is less than the critical value at a 5% level of significance. Tests are inconclusive for the other three policies. By again expanding our experiment with each of these three policies, we find, at a 5% level of significance, that mean service levels are below α = 0.95. [For instance, for the policy <62, 41>, an expanded experiment of 200 replications showed a mean of 94.763 and a standard deviation of 1.876, yielding a t statistic of -1.784 (which is below the critical value for a left hand-tailed test of hypothesis at a 5% level of significance).] Hence, we are able to infer, at a 5% level of significance, that no policy $< B_1, B_j >$ involving 185 units of echelon base stock satisfies α = 0.95.

Having found, at a 5% level of significance, that the selected policy <54, 44>—with 186 units of echelon base stock—satisfies α = 0.95, and that no policy involving 185 units of echelon base stock is able to satisfy the given service level, we now make important theoretical observations that apply to the general situation.

Lemma 1 *For any given* B_j, *the probability of no stockout associated with the policy* $< B_1, B_j >$ *increases monotonically with* B_1.

Proof: In section 2, we had noted that $S_1 \sim \text{gamma}(\lambda_1, B_1)$. It follows that $E[S_1] = B_1/\lambda_1$, and $E[S_1]$ increases with B_1. Since $T_j = \min\{p_j, S_1\}$, then $E[T_j]$ also increases with B_1. Accordingly, the expected width of the time interval $(T_j, t_r]$, given τ_j, decreases with B_1. Therefore, for any given level of retailer base stock B_j, $\Pr\{ (D_j[T_j,t_r] \mid \tau_j) \le B_j \}$ increases with B_1, as does the probability of no stockout, $\Pr\{ D_j[T_j,t_r] \le B_j \}$. *q.e.d.*

Lemma 2 *Take any level* E *of echelon base stock. Consider the sets* P_E *and* P_{E-1} *of all (integer-valued) policies* $< B_1, B_j >$ *at the echelon base stock levels* E *and* E-1*, respectively. Then the maximum probability of no stockout possible for all policies in* P_{E-1} *will be less than the maximum probability of no stockout possible for all policies in* P_E.

Proof: Let s_{E-1} be the maximum probability of no stockout for all policies in P_{E-1} and s_E be the maximum probability of no stockout for all policies in P_E. Assume that $s_{E-1} \ge s_E$. Let the probability s_{E-1} be at some specific policy $< B_1, B_j >$ in P_{E-1}. Then the policy $< B_1+1, B_j >$ is one of the policies in the set P_E, since $e\{< B_1, B_j >\} = B_1 + NB_j = E-1$ and $e\{< B_1+1, B_j >\} = (B_1+1) + NB_j = E$. If s^+ denotes the probability of no stockout at $< B_1+1, B_j >$, then $s^+ \le s_E$—and, therefore, $s^+ \le s_{E-1}$ according to the assumption we made. This runs contrary to Lemma 1, which tells us that s^+ (the probability of no stockout at $< B_1+1, B_j >$) must be greater than s_{E-1} (the probability of no stockout at $< B_1, B_j >$). We conclude that $s_{E-1} < s_E$. *q.e.d.*

Returning to Example 1, it follows from Lemma 2 that the maximum service level possible for any echelon base stock level less than 185 will be less than that for E = 185, at which no "candidate" policy satisfies α = 0.95. Hence, we conclude that the minimum level of echelon base stock that satisfies $\alpha = 0.95$ is 186 units.

Interpretation of Heuristic Solution for the Case Study

To achieve a 95% probability of not stocking out, the optimal base stock policy with deterministic leadtime $\tau_j = 1$ is $< B_1^d, B_j^d >$ = <56, 39>, with echelon base stock = 173 units. Using this deterministic leadtime optimal policy as the starting point, the search heuristic under stochastic leadtimes yielded the policy $< B_1^h, B_j^h >$ = <54, 44>, with echelon base stock of 54 + (3 × 44) = 186 units. This base stock policy satisfies the specified 95% probability of not stocking out and yields the optimal echelon base stock in the given case when the warehouse is made to carry stock. It compares favorably with the optimal retailer base stock level of 63 units, leading to an echelon base stock of 3 × 63 = 189 units, if the warehouse does not carry stock (Solis and Schmidt 2007).

The next example involves the fill rate criterion.

Example 2

Consider scenario 7 (N = 3), with τ_j - 4.5 ~ beta(2,6) and β = 0.99. We know that $< B_1^d, B_j^d >$ = <59, 91>, with echelon base stock = 332; $B_j^d(0) = 112$; $B_j^s(0) = 110$; and $\delta = B_j^s(0) - B_j^d(0) = 110 - 112 = -2$.

Step 1. Initialize: $< B_1^h, B_j^h > \leftarrow < B_1^d, B_j^d + \delta >$ = <59, 89>, with echelon base stock = 332 + 3δ = 326. s{<59, 89>} = 99.05% > β. In this case, we proceed to step 3.

Step 3. In trying to adjust $B_1^h = 59$ downward, we find that s{<58, 89>} = 98.98% < β. We keep the current policy $< B_1^h, B_j^h >$ = <59, 89>, with echelon base stock = 326.

In step 1, at the policy <59, 89>, our typical experiment yielded an average observed fill rate of 99.050(%), with a standard deviation of 0.523. The computed t statistic was only 0.95, which falls below the critical value for the right hand-tailed test of hypothesis at a 5% level of significance. However, in expanding the experiment to 400 replications, we computed a mean fill rate of 99.049 and a standard deviation of 0.535. The resulting t value was 1.818, and we infer that the mean fill rate satisfies $\beta = 0.99$ at a 5% level of significance.

A heuristic lower bound on B_j, as developed by Solis (1997), yielded a value of $B_j^{min,h} = 83$. Table 3 shows the average observed fill rates at policies of interest.

At the echelon base stock level of 325 (one unit below the level at the selected policy <59, 89>), we consider, for instance, the policy <46, 93>. Our typical experiment yielded an average observed fill rate of 99.035(%), with a standard deviation of 0.488. The resulting t statistic of 0.71 is below the critical value at a 5% level of significance for a right hand-tailed test of hypothesis about the mean fill rate. However, when we expanded the experiment to 500 replications, we found a mean of 99.044 and a standard deviation of 0.507, yielding a t value of 1.953, which exceeds the critical value at a 5% level of significance. We have thus found one policy—among a number of policies, in fact—satisfying $\beta = 0.99$, and at which the echelon base stock level of 325 is one unit lower than the level of 326 at the policy <59, 89> selected by the heuristic.

At the next lower echelon base stock level of 324, there are eight policies in all (<33, 97>, <36, 96>, <39, 95>, <42, 94>, <45, 93>, <48, 92>, <51, 91>, and <57, 89>) for which left hand-tailed tests of hypothesis concerning the mean fill rate are not conclusive based on our typical experiments. However, by increasing the number of replications, we are able to infer at a 5% level of significance that all policies with echelon base stock = 324 result in mean fill rates that are lower than β = 0.99. [For example, at the policy <51, 91>, our typical experiment yielded a mean of 98.954(%), with a standard deviation of 0.529, resulting in a t statistic of only -0.862. With 400 replications, the observed mean was 98.947, with a standard deviation of 0.551. The computed t value in the latter case was -1.918, which is to the left of the critical value for the left hand-tailed test of hypothesis at a 5% level of significance.]

Table 3: Average Observed Fill Rates (%) – Example 2

Echelon Base Stock = 324			Echelon Base Stock = 325			Echelon Base Stock = 326		
B_1	B_j	Average Observed System Fill Rate (%)	B_1	B_j	Average Observed System Fill Rate (%)	B_1	B_j	Average Observed System Fill Rate (%)
0	108	98.76	1	108	98.86	2	108	
3	107	98.79	4	107	98.87	5	107	
6	106	98.8	7	106	98.89	8	106	
9	105	98.81	10	105	98.9	11	105	
12	104	98.82	13	104	98.905	14	104	
					(98.80, 99.01)			
15	103	98.84	16	103	98.928	17	103	
					(98.83, 99.03)			
18	102	98.87	19	102	98.947	20	102	
					(98.85, 99.04)			
21	101	98.88	22	101	98.959	23	101	
					(98.86, 99.05)			
24	100	98.89	25	100	98.975	26	100	
					(98.88, 99.07)			
27	99	98.9	28	99	98.98	29	99	
					(98.89, 99.07)			
30	98	98.908	31	98	98.991	32	98	
		(98.81, 99.01)			(98.90, 99.09)			
33	97	98.919	34	97	99.004	35	97	
		(98.82, 99.02)			(98.91, 99.10)			
36	96	98.928	37	96	99.013	38	96	
		(98.83, 99.03)			(98.92, 99.11)			
39	95	98.942	40	95	99.024	41	95	
		(98.84, 99.04)			(98.93, 99.12)			
42	94	98.953	43	94	99.031	44	94	
		(98.85, 99.05)			(98.94, 99.13)			
45	93	98.954	46	93	99.035	47	93	
		(98.85, 99.05)			(98.94, 99.13)			
48	92	98.949	49	92	99.026	50	92	
		(98.85, 99.05)			(98.93, 99.12)			
51	91	98.954	52	91	99.026	53	91	
		(98.85, 99.06)			(98.93, 99.13)			
54	90	98.933	55	90	99.01	56	90	
		(98.83, 99.04)			(98.91, 99.11)			
57	89	98.913	58	89	98.985	59	89	99.05
		(98.80, 99.02)			(98.88, 99.09)			(98.95, 99.15)
60	88	98.87	61	88	98.94	62	88	
					(98.83, 99.05)			
63	87	98.81	64	87	98.88	65	87	
66	86	98.71	67	86	98.77	68	86	
69	85	98.57	70	85	98.63	71	85	
72	84	98.38	73	84	98.43	74	84	
75	83	98.1	76	83	98.14	77	83	

It is difficult to mathematically make theoretical observations for the fill rate criterion as we had done for the probability of no stockout criterion in Lemmas 1 and 2. This difficulty arises because expected backorders do not work out as "neatly" as the probability of not stocking out. Nevertheless, "empirical evidence" from our simulation experiments suggests that, for given B_j, the (observed) fill rate associated with policy < B_1, B_j > is likewise monotone increasing as B_1. We state this observation as well as a statement analogous to Lemma 2 as *conjectures*.

<u>*Conjecture 1*</u> *For any given* B_j, *the fill rate associated with the policy* < B_1, B_j > *increases monotonically with* B_1.

<u>*Conjecture 2*</u> *Take any level* E *of echelon base stock. Consider the sets* P_E *and* P_{E-1} *of all (integer-valued) policies* < B_1, B_j > *at the echelon base stock levels* E *and* E-1, *respectively. The maximum fill rate possible for all policies in* P_{E-1} *will be less than the maximum fill rate possible for all policies in* P_E.

Returning to Example 2, since no policy having echelon base stock = 324 (or less) satisfies β = 0.99, we conclude by virtue of Conjecture 2 that no policy with echelon base stock less than 325 will satisfy the specified fill rate. Hence, the optimal echelon base stock level in this case is 325 units, and the policy <59, 89> selected by the heuristic yields an inventory level that is one unit more than optimal.

The above two examples we have provided illustrate how well our heuristic appears to perform with respect to selecting a policy < B_1, B_j > which yields an echelon base stock level (and, thus, an average system on-hand inventory level) that is close to the smallest possible level at which the specified α or β is satisfied. In these, as well as other test cases we have thus far investigated—albeit limited, in view of the computing effort involved, to only a portion of Graves' 64 test cases under either service criterion—we have found the selected policy to result in inventory levels that are generally either equal to or one unit more than optimal. Considering average inventory levels of roughly between 100 and 400 in our limited test cases, our heuristic seems to perform fairly well.

The choice of the optimal policy < B_1^d, B_j^d > in the deterministic leadtime case as the starting point for the heuristic appears to be an appropriate one. We are unable to provide a guarantee, however, that the heuristic would err by no more than one unit away from the optimal level of average system inventory. We have found that the average observed service levels (based on comparable simulation experiments) at policies having the same echelon base stock are not always unimodal. One, two or more relative maxima may exist within a given echelon base stock level E. We had earlier noted (by way of Lemma 1 or Conjecture 1) that, for given B_j, observed service levels increase as B_1 (and, hence, as the echelon base stock level E = B_1 + NB_j) increases; however, we were unable to observe any uniformity or pattern in the behavior of such increases across varying B_1 values (for a given B_j) or across different B_j values.

Lemma 2 or Conjecture 2 states that the maximum service level (probability of no stockout or fill rate, as the case may be) possible for echelon base stock level E will be larger than the maximum for E-1. We have observed, however, that the maxima at different echelon base stock levels may be located at various levels of B_j. The service level observed for some selected policy < B_1, B_j >, having echelon base stock level E, may satisfy the specified service level α or β fairly closely (in particular, in view of the steps in our heuristic). Yet, the maximum service level at the next lower echelon base stock level E-1 (which could well be at some policy < B_1', B_j' > ≠ < B_1-1, B_j >)—or that at the even lower echelon base stock level E-2, for that matter—may still be able to satisfy α or β.

4. CONCLUSION

In this study, we extended Graves' one-warehouse, N-retailer model by introducing, as Graves suggested, stochastic leadtimes τ_j between the CW and the retail sites in place of the original deterministic leadtimes. We investigated the effects of stochastic leadtimes—on the optimal base stock policy < B_1, B_j > for the deterministic leadtime case, in which the CW actually carries stock—by way of simulation studies. We have devised a heuristic for selecting, in the stochastic leadtime case, a policy < B_1, B_j > that seeks to minimize average system inventory. The heuristic appears, based on test cases we have investigated, to select a policy at which the average system inventory is equal or very close to the optimal average inventory level.

REFERENCES

Eppen, G. and Schrage, L., 1981. Centralized ordering policies in a multi-warehouse system with lead times and random demand, In: L.B. Schwarz, ed. *Multi-level production/inventory control systems: theory and practice*, TIMS Studies in the Management Sciences, 16, Amsterdam: North Holland, 51-67.

Graves, S.C., 1985. A multi-echelon inventory model for a repairable item with one-for-one replenishment, *Management Science*, 31 (10), 1247-1256.

Graves, S.C., 1996. A multiechelon inventory model with fixed replenishment intervals. *Management Science*, 42 (1), 1-18.

Kelton, W.D., Sadowski, R.P., and Sturrock, D.T., 2007. *Simulation with Arena*, 4th ed., New York, NY: McGraw-Hill.

Lee, H.L. and Moinzadeh, K., 1987. Two-parameter approximations for multi-echelon repairable inventory models with batch ordering policy. *IIE Transactions*, 19, 140-149.

McGavin, E.J., Schwarz, L.B., and Ward, J.E., 1993. Two-interval inventory-allocation policies in a one-warehouse, N-identical-retailer distribution system. *Management Science*, 39 (9), 1092-1107.

Nahmias, S. and Smith, S.A., 1994. Optimizing inventory levels in a two-echelon retailer system with partial lost sales. *Management Science*, 40 (5), 582-596.

Pegden, C.D., Shannon, R.E., and Sadowski, R.P., 1995. *Introduction to simulation using SIMAN*, 2nd ed., New York, NY: McGraw-Hill.

Solis, A.O., 1997. Evaluation of the negative binomial approximation and stochastic leadtimes in a multi-echelon inventory model. Thesis (PhD). The University of Alabama.

Solis, A.O., 1998. On the depot effect in a multi-echelon inventory model with fixed replenishment intervals. *Proceedings of the Twenty-Seventh Annual Meeting of the Western Decision Sciences Institute*, pp. 537-540, April 7-11, Reno (Nevada, USA).

Solis, A.O., and Schmidt, C.P., 2007. Stochastic leadtimes in a one-warehouse, N-retailer inventory system with the warehouse not carrying stock. *European Journal of Operational Research*, 181 (2), 1004-1013.

Solis, A.O., Schmidt, C.P., and Conerly, M.D., 2007. On the effectiveness of the negative binomial approximation in a multi-echelon inventory model: a mathematical analysis. *Proceedings of the International Conference of Numerical Analysis and Applied Mathematics*, American Institute of Physics Conference Proceedings 936, pp. 531-534, September 16-20, Corfu (Greece).