TRANSPORTATION CONTRACTS AND SAFETY STOCKS FOR JUST-IN-TIME DELIVERIES

Candace Arai Yano
Department of Industrial & Operations Engineering
University of Michigan
Ann Arbor, MI

Yigal Gerchak
Department of Management Science
University of Waterloo
Waterloo, Ontario, Canada

March 1989
Technical Report 89-8
Revised November 1989

TRANSPORTATION CONTRACTS AND SAFETY STOCKS FOR JUST-IN-TIME DELIVERIES

ABSTRACT

Just-in-time manufacturing operations often establish transportation contracts to ensure the timely delivery of inbound freight. If the usage of parts is constant and known, it is easy to decide on the characteristics of the transportation contract taking into account costs and other practical considerations. If the usage of parts is not perfectly predictable, however, the transportation contract and safety stock levels must be established in view of the potential costs of emergency shipments (premium freight) and shortages. We address the problem of simultaneously determining characteristics of a transportation contract and safety stocks to support just-in-time operations at minimum cost at the interface between an assembly facility and a single high volume supplier from which it receives direct shipments. Heuristic procedures are developed and some general conclusions are drawn from numerical examples for a prototypical demand process.

TRANSPORTATION CONTRACTS AND SAFETY STOCKS FOR JUST-IN-TIME DELIVERIES

1. INTRODUCTION

Major assembly facilities that have established just-in-time arrangement with their suppliers are increasingly using scheduled shipments to ensure reliable deliveries. A survey by Lieb and Miller (1988) indicates that 78 percent of manufacturing firms have reduced the number of (transportation) carriers since the implementation of just-in-time, and 73 percent have negotiated specific contracts with carriers to support their just-in-time operations. These contracts have a number of benefits. The stability provided by these contracts decreases uncertainty in the transportation lead time. In addition, it permits suppliers to develop production schedules that are compatible with the delivery schedule, which in turn tends to reduce the amount of inventory at each supplier. Such contracts also contribute to reducing the overall cost of transportation because the transporter can schedule both labor and equipment in advance.

Our study was motivated by a set of applications at a U.S. automobile manufacturer. In the U.S. automobile industry, three types of just-in-time arrangements are common. For low-volume suppliers, either there are scheduled less-than-truckload shipments, or "milk runs" are arranged in which one truck stops at several suppliers to consolidate goods for delivery to a single assembly facility. For medium volume suppliers, especially those that produce components for several assembly facilities, there are usually full truckload (mixed-destination) shipments into break-bulk facilities. There, components from several suppliers are sorted, and components are consolidated by destination, allowing full truckload (mixed-origin) inbound shipments to the assembly facilities. For high-volume suppliers, or those which manufacture expensive components, direct shipments are common. In the former case, there is no need for consolidation. In the latter case, security considerations prohibit unneccesary handling, and the additional cost and effort of maintaining security tends to lead to larger shipments . It is these direct-shipment cases that we consider here.

A typical just-in-time transportation contract for a high-volume supplier specifies the timing (i.e., day, or day and hour) and size (e.g., number of trucks, number of rail cars) of each shipment. Because of the inventory reduction resulting from the just-in-time arrangement, suppliers are normally required to ship whatever has been ordered, if the components are available (i.e., short shipments are strongly discouraged). Occasionally, because demands for the various parts fluctuate, there may be more goods to ship than will fit into the contracted vehicles. In these situations, one or more "emergency" vehicles must be arranged at premium rates, assuming that the time between shipments is long enough to warrant emergency shipments.

Although just-in-time concepts are intended for systems with nearly constant demand, some demand uncertainty exists and must be incorporated appropriately into planning models. In the applications that motivated our investigation, most suppliers receive daily, and in some cases, real-time, reports on the quantities to be shipped over the subsequent few days or weeks. The quantities are based upon various characteristics (colors, options, etc.) of specific orders to be assembled. For example, the number of air conditioners to be delivered on a given day would typically be the actual number of orders requiring air conditioners on that day. Schedule changes at the assembly facility occur frequently, however, for a variety of reasons. In most cases, the changes are simply alterations in the assembly sequence because of rework. In other cases, shortages of a component or a machine failure may necessitate a modified schedule which excludes orders requiring that component or operation. Thus, although each order is eventually assembled, the actual timing often deviates from the planned timing. Since the planned assembly sequence is changing in real-time, the planned shipment quantities are also. As a consequence of this, the actual shipment quantities not only vary from shipment to shipment, but are, in reality, uncertain until very near the time of shipment.

These circumstances create a need for some type of flexibility in the shipping arrangements. Since the timing of contracted shipments must be specified in advance, it is difficult to provide this flexibility through modifications in the timing of shipments. Instead, it can be provided by (i) allowing some slack space in each truck trailer, on the average, to accommodate fluctuations and/or (ii) arranging for additional shipments on an "as needed" basis. In addition to flexibility in the shipping arrangements, it may be desirable to hold some safety stock at the assembly facility to accommodate "normal" schedule fluctuations, especially if the supplier is instructed to ship a quantity equal to the *planned* requirements. (These shipping instructions, or minor variations thereof, are quite common in the automobile industry.) If the assembly sequence were to remain as initially planned, there would be no need for safety stock if the suppliers are completely reliable. Because of the wide variety of options and colors, however, there still would be fluctuations in the component usage rates from day to day, and the exact usage rates would not be known far enough in advance to use as input in determining an annual (or longer) transportation contract. Thus, from the perspective of the tranportation planner, the daily requirements appear as observations from a random demand process, and it is practical to model them as such.

2. PROBLEM DESCRIPTION AND RELATED LITERATURE

In this paper, we develop an optimization-based model and heuristic solution procedure for the problem of establishing a transportation contract between an assembly facility and one of its high-volume suppliers. The procedure simultaneously determines the characteristics of the transportation contract (time between shipments and number of vehicles per shipment) and the amount of safety stock at the assembly facility. To keep the model simple, we use a single aggregate component, but the extension to multiple components is straightforward.

The objective in the optimization model is the minimize the total expected cost per unit time of the following:

(i) cost of the contracted shipments;

(ii) cost of emergency shipments;

(iii) cycle stock holding costs;

(iv) safety stock holding costs; and

(v) shortage costs.

We assume that transportation is by truck, which is the most common mode of just-in-time deliveries. Since the cost of truck transportation is very insensitive to the utilization of the truck trailer, we charge a fixed cost per truck movement between the supplier and the assembly facility. This assumption is consistent with existing common carrier pricing policies in the context of a long-term contract. The fixed cost may vary by supplier, of course. Similarly, emergency shipments are charged per truck movement, irrespective of the utilization of the trailer. It is assumed that the contracted shipments are less expensive per shipment than the emergency shipments, which is true in practice.

The assembly facility is assumed to consist of a paced or unpaced assembly line, or another type of assembly system in which the shortage of *any* part might force the facility to stop production, or at a minimum, create a costly disruption. Thus, emergency shipments are always arranged whenever the goods overflow the contracted vehicles. Shortages occur only if the safety stock provision is not large enough to cover the fluctuation. Shortage costs are charged on a "per unit short" basis, since the severity of the consequences is correlated with the number of units short.

The usage of the aggregate component by the assembly facility follows a stationary stochastic process, which can be interpreted to mean that the forecast remains constant, but there are forecast errors, and the forecast errors have the same distribution over time. This is a reasonable representation for the scenarios that motivated this research, because the sequencing procedure for the assembly facility attempts to sequence orders so that important components related to customer-selected options are used at a reasonably constant rate. Moreover, actual data obtained from one automobile company indicates that actual order quantities, even at the component level, appear to follow a stationary stochastic process, except during plant shutdowns and model-year transition periods.

We assume that the assembly facility uses an (S,T) procurement policy for the aggregate component. That is, it orders up to S units every T time units, where T is the time interval between deliveries. Although an assembly facility may not use this particular policy, the actual order quantities tend to exhibit characteristics similar to those of orders that would be generated by such a policy. The value of S is determined through a tradeoff between expected shortage costs and expected inventory holding costs.

The ordering information is transmitted electronically and therefore there is no delay associated with the ordering process itself. Also, we assume that the transportation time between the supplier and the assembly facility is very short relative to the time between deliveries. Thus, there is always at most one order outstanding, and the impact of the delivery leadtime upon the optimal amount of safety stock in the system is considered to be negligible. (The extension to positive lead times is straighforward but unnecessarily complicates the exposition.)

To our knowledge, this particular problem has not been investigated before in a system with uncertain shipping requirements. However, the issues are well-understood. Jackson (1983) describes some of the implications of just-in-time implementation for logistics managers. He mentions the problem of achieving economies of scale in transportation (i.e., full truckload shipments) while simultaneously keeping inventory levels low, as well as the importance of close cooperation among shippers, customers, and suppliers. Herron (1979) discusses situations in which planned expediting (analogous to our emergency shipments) is preferable to incurring stockouts. He also describes the tradeoff among transportation costs, inventory costs, and stockout costs in deciding the best shipping frequency.

One paper that addresses a problem related to ours, but in a deterministic setting, is by Hill and Vollman (1986). Their paper deals with scheduling and routing trucks for vendor pickups in a just-in-time setting. They suggest that benefits from establishing coordinated schedules include the reduction of uncertainty and the ability to reduce safety buffers (i.e., inventory).

Another somewhat related set of problems are stochastic vehicle routing problems. Related work includes Cook and Russell (1978), Golden and Stewart (1978), Golden and Yee (1979), and Tillman (1969). While there are several variations of these problems, the general focus is to determine vehicle routes for a set of locations with random demands. The goal is to minimize expected travel time or distance, and the probabilistic aspects are incorporated by including chance constraints on completing the routes or considering the time or distance to complete the route with another vehicle. The stochastic vehicle routing problem most similar to ours is the garbage collection problem, of which a fairly complicated version is addressed by Cook and Russell. Here, all of the "demand" must be satisfied and extra vehicles must be dispatched if necessary. However, all of these models differ from ours in that they do not include inventory costs, nor is there a significant premium charged for unplanned routes (except for the extra time or distance). In addition, our problem does not require routing decisions.

A problem somewhat more closely related to ours is the inventory-routing problem considered by Federgruen and Zipkin (1984). In this single-period problem, there are multiple demand locations with random demands, and both the vehicle routes and delivery quantities must be chosen. The objective is to minimize transportation, inventory, and shortage costs. Here, "short-shipments" are allowed; that is, a portion of a location's request (or desired delivery quantity) may be backordered, if the goods will not fit onto the vehicles. Thus, the relative importance of shortages is considerably less in this problem than it is in ours. Although we have no routing decisions, our problem is complicated by the fact that we must determine the frequency of visits, which in turn, affects the demand distributions.

3. SUMMARY OF RESULTS

We view the contributions of this work as threefold: first, the introduction of a simple model which might provide the foundations for more complicated and more realistic models; and second, a procedure which might be used to provide a first-cut solution, which can then be refined with a more detailed analysis; and third, results from a computational study which can provide some insight into how various factors influence the characteristics of good solutions and their costs. The main aspects of the model have been described above. The reader is referred to sections 5 and 6 for additional details and a description of the heuristic procedure. In the remainder of this section, we will focus on the insights gained from the model and the computational study.

Recall that the decisions included the time between shipments and the number of vehicles per shipment. Most existing just-in-time operations use one truck per shipment under the premise that smaller shipments reduce inventory, which is certainly true. The primary reason for considering policies in which each contracted shipment consists of more than one truck is the effect of risk pooling (over time) on the number of emergency truck shipments. For example, suppose the current policy is to have one contracted truck per day. Although both cycle and safety stock will increase with a policy of "two trucks every other day," the number of exposures to possible emergency shipments is decreased by half and the resulting savings may more than offset the inventory cost increase.

One of the primary results from our model is that for any given number of trucks per shipment, n, the time between shipments, T, is smaller than the value from an appropriate economic order quantity (EOQ) computation. Both the premium freight and safety stock considerations encourage a smaller time between shipments, and hence, more capacity slack, on the average, in the contracted shipments.

In our computational study, we used a Brownian motion demand process, which is the continuous-time equivalent of independent and identically distributed Normal demands in a discrete time framework. We investigated a variety of problems with parameters that are reflective of the applications that motivated the study. The most surprising finding is that having more than one truck per shipment is desirable under the following conditions:

(i) high demand variability (coefficient of variation greater than 10%),

(ii) high emergency shipment costs (two or more times the contracted price),

(iii) low inventory holding costs in comparison to normal transportation costs, and

(iv) relatively low shortage costs in comparison to the inventory holding costs.

In these instances, we found that the savings from using the solution from the heuristic procedure rather than the solution with one truck per shipment were as much as 40% of the total cost. The savings resulted from risk pooling over time, which allowed both an increase in the average utilization of the contracted vehicles (and hence fewer total truck movements) *and* fewer emergency shipments. The reduction in the number of emergency shipments is a consequence of having much less frequent, but only slightly larger, overflows.

The numerical results also suggest that the larger the variability of demand, the smaller the average utilization of the contracted vehicles should be (i.e., there should be more planned slack). This result is not surprising. What is surprising is that we obtained some solutions with average utilization levels of less than 60% for sets of parameters which are reflective of some actual costs and demand patterns. Our model provides a quick-and-dirty method to assess the savings from dampening the variations. This dampening could be achieved, for example, through a reduction of the product line, product redesign to increase commonality, or marketing strategies to smooth the usage rates of components related to options.

We also found that the solutions were affected little, if at all, by the magnitude of the shortage cost, as long as the shortage costs are high enough to require emergency shipments rather than making short-shipments. Thus, in practical settings, it would be possible to determine the characteristics of the transportation contract ignoring safety stock considerations, then set the safety stock levels given the transportation contract.

4. DISCUSSION

We have developed an optimization-based model and a heuristic solution procedure to simultaneously determine the characteristics of a transportation contract (frequency of shipment and number of vehicles per shipment) and the amount of safety stock to support a just-in-time manufacturing system with random demand for parts. Numerical results from the model suggest that, because of the possible need for emergency shipments at premium rates, full-truckload and "one vehicle per shipment" policies are not always optimal. By providing an appropriate amount of slack in the contracted shipments, the best tradeoff between the cost of the contracted shipments and the cost of emergency shipments can be achieved. When demand variability is sufficiently high, it is optimal to ship more than one vehicle at a time. By doing so, the system is able to take advantage of risk pooling over time to reduce the number of emergency shipments.

Further research is needed to generalize this model to incorporate more realistic costs, especially for emergency shipments. It may be necessary to use alternate means of transportation, such as air freight, to ensure timely delivery. The charges for air freight normally depend upon the weight and volume of the shipment. Thus, the "fixed charge per emergency truck" penalty would need to be modified accordingly. Research is also needed to generalize our model to explicitly consider multiple items.

The remainder of this paper is designed for the technical reader, and is organized as follows. A formulation is given in Section 5. A few properties of the optimal solution and a heuristic solution procedure are presented in Section 6. In Section 7, we provide a few examples and discuss results of a computational study from which we draw some insights. Section 8 concludes the paper with analogies between our results and those of simple inventory models, and additional directions for future research.

5. MODEL FORMULATION

The following formulation is based upon the assumptions stated earlier in the paper.

Notation:

j = item index ($j = 1,...J$),

T = time between shipments (in years),

n = number of contracted vehicles per shipment,

S_j = order-up-to point for item j,

A = transportation cost per vehicle for regular shipment,

C = transportation cost per vehicle for emergency shipment,

h_j = annual inventory holding cost for one unit of item j,

π_j = shortage cost for item j (per unit),

v_j = volume per unit of item j (measured as fraction of a truckload),

$X_j(T)$ = demand for item j in an interval of duration T (random variable), and

$f_j(X_j;T)$ = density of demand for item j in an interval of duration T.

We assume that $E[X_j(T)] = \mu_j T$ (i.e, the demand process is stationary).

The problem is to choose values of n, T, and S_j, j = 1,...J so as to

$$\min \ \{An/T + \Sigma_j \{h_j [\mu_j T/2 + \int_{-\infty}^{S_j} (S_j - x_j) f_j (x_j;T) dx_j] + (\pi_j/T) \int_{S_j}^{\infty} (x_j - S_j) f_j (x_j;T) dx_j\}$$

$$+ (C/T) \sum_{k=0}^{\infty} P(\Sigma_j v_j X_j(T) > n + k)\} \qquad (1)$$

subject to $n \geq 0$, integer

$T \geq 0$

$S_j \geq 0$ for all j

All terms in the objective function represent annual costs. The first term is the cost of the contracted truck movements. The second term is the cycle stock holding cost, while the third term is the expected cost of holding safety stock (or equivalently, the holding cost associated with expected end-of-cycle inventory). The fourth term is the expected cost of shortages. Note that the integral expression in the fourth term is the expected number of units short in a cycle of duration T, and thus it must be divided by T to obtain the expected number of shortages per unit time. Finally, the fifth term is the expected cost of the emergency vehicles, and it is obtained by observing that

$$E\{(\lceil \Sigma_j v_j X_j(T) - n \rceil)^+\} = \sum_{k=0}^{\infty} P(\Sigma_j v_j X_j(T) > n + k)\}$$

where $\lceil y \rceil$ denotes the ceiling of y, i.e., the quantity is rounded to the next larger integer, and $(x)^+ = \max\{0,x\}$. The expression takes this form because we must charge for the entire emergency vehicle, even if only part of it is used, and the number of emergency vehicles must be non-negative.

6. PROPERTIES OF THE OPTIMAL SOLUTION AND A SOLUTION PROCEDURE

We first note that the third and fourth terms of the objective function represents an imbedded newsboy-like problem in which n does not appear. Moreover, the problem is separable by item. Thus, for each item j, the imbedded subproblem is to find S_j so as to

$$\min \; h_j \int_{-\infty}^{S_j} (S_j - x_j) f_j (x_j;T)\, dx_j + (\pi_j/T) \int_{S_j}^{\infty} (x_j - S_j) f_j (x_j;T)\, dx_j \;\; = G_j(T). \qquad (2)$$

Using standard approaches, we find the solution as

$$S_j^*(T) = F_{Xj(T)}^{-1}[\pi_j/(\pi_j + h_jT)], \qquad (3)$$

which can be substituted into the original objective function, leaving a function of T.

We next analyze the properties of G(T), the optimal value of the newsboy portion of the objective function, dropping the subscript for ease of exposition.

The expression in (2) can be rewritten as

$$h\{ F_{X(T)}^{-1}[\pi/(\pi+hT)] - \mu T \} + (h+ \pi/T) \int_{F_{X(T)}^{-1}[\pi/(\pi+hT)]}^{\infty} P[X(T) > x]\, dx.$$

The derivative with respect to T can be shown to equal

$$-h\mu - (\pi/T^2) \int_{F_{X(T)}^{-1}[\pi/(\pi+hT)]}^{\infty} P[X(T) > x]\, dx$$

$$+ (h + \pi/T) \int_{F_{X(T)}^{-1}[\pi/(\pi+hT)]}^{\infty} \partial P[X(T) > x]/\partial T\, dx.$$

The first two terms are clearly negative. The sign of the third term depends upon the distribution of $X(T)$. If we assume that $hT \leq \pi$ (i.e, $\pi/(\pi + hT) \geq 0.5$), then in most instances, the second integral is positive. Thus, we cannot make any definitive statements about the sign of the first derivative. Empirical evidence (see Figure 1) suggests that $G(T)$ is a concave increasing function of T. Conceptually, the concavity is the result of risk pooling over time as T increases.

FIGURE 1

We now turn to an examination of the last term in the objective function. Once again, it is quite difficult to derive properties analytically. An empirical investigation reveals that for a given n, this function has the form illustrated in Figure 2. For "small" values of T, the sum in (1) is zero or nearly zero, so the entire expression is very small. The function then begins to increase, and is asymptotically (but actually soon becomes) concave because of the effects of risk pooling over time.

FIGURE 2

Although the objective function is not convex, each portion of the objective function appears to be relatively well-behaved. The reasons for this are as follows. For small values of T, the derivative of the function is dominated by the first term, which is convex decreasing in T. Thus, we might expect the function to be decreasing as T increases from zero. As T becomes large, the derivative of the sum of the last four terms, for which the sum is "almost" concave increasing, dominates. Consequently, the function should eventually increase.

It is not clear, however, how the embedded newsboy problems influence the shape of the objective function, especially when the cost and demand characteristics of the items vary. Thus, we limit the experimental study in the next section to the case of one item or one family of items for which the safety-stock related costs can be modeled in an aggregate manner.

At this point, it would be desirable to find an expression for T^* as a function of n. This would permit us to search over n (which is constrained to be an integer) to find a solution quickly. This is difficult to do, however, so we concentrate on obtaining other properties that can be used to accelerate a two-dimensional search over n and T.

Since the optimal expected cost of the embedded newsboy problem and the expected cost of emergency vehicles are increasing with T, T satisfies

$$T \leq \{2An/h\mu\}^{0.5} \ (= T_{max}),$$

where the right hand side is the solution to the problem if we were to ignore the newsboy-related and emergency vehicle costs. This constraint limits the range of T over which we must search for each value of n. In addition, economic factors generally dictate that the vehicles should not be more than 100% utilized on the average, since otherwise emergency shipment costs will be exorbitant. Consequently, in most situations it is also safe to assume that

$$T \leq n/\mu$$

and one might choose to limit the search to $T \leq \min\,[n/\mu, \{2An/h\mu\}^{0.5}]$.

We now derive some properties of the optimal value of n for a given T. Observe that $G_j(T)$ is not a function of n. Thus, for fixed T, we have the following optimization problem to find n:

$$\min\,\{An/T + (C/T)\sum_{k=0}^{\infty} P[\Sigma_j\, v_j X_j(T) > n + k]\} \qquad (4)$$

subject to $\quad n \geq 1$, integer.

Clearly the first term is linearly increasing with n, while the second term is decreasing with n. The first difference of (4) with respect to n is

$$A/T - (C/T)\, P[\Sigma_j v_j X_j(T) > n].$$

Since for fixed T, $P[\Sigma_j v_j X_j(T) > n]$ is non-increasing with n, the optimal solution is to choose

$$n^*(T) = \sup \{n: P[\Sigma_j v_j X_j(T) > n] > A/C\} + 1$$

$$\text{or } \inf \{n: P[\Sigma_j v_j X_j(T) \le n] \ge (C-A)/C\},$$

or any value between these two. (The solution may not be unique if $\Sigma_j v_j X_j(T)$ is not stochastically strictly increasing.) The expressions for $n^*(T)$ indicate that $n^*(T) \ge n^*(T+\Delta T)$ if $\Sigma_j v_j X_j(T)$ is stochastically increasing. This fact can be used in a search procedure, when applicable. In particular, if $\Sigma_j v_j X_j(T)$ is stochastically increasing, then $n^*(T) \ge n^*(T_{max})$ for all $T < T_{max}$.

Since the objective function is not convex, we cannot use standard non-linear programming techniques to solve the problem. Moreover, we have already seen that even the first order conditions are quite complicated. For these reasons, we use a search procedure in which we search over n, and find a near-optimal value of T for each n. Since there is an upper bound (T_{max}) on the maximum value of T, this can be done quite easily even if the cost function is not unimodal in T for a fixed n. Then, if there is reason to believe that $z(n, T^*(n))$ is unimodal in n, we can start with n = 1 and increment it until $z(n, T^*(n))$ begins to increase. On the other hand, if there is evidence that $z(n, T^*(n))$ is not unimodal in n, then we can simply compute $z(n, T^*(n))$ for all practical values of n. In most instances, storage capacity constraints will limit the maximum value of n.

7. EXPERIMENTAL INVESTIGATION

Our experimental investigation was designed to test the reasonableness of our assumptions about characteristics of the objective function and to understand how the parameters affect the optimal solution. Throughout our investigation, we assume that the demand process is represented by Brownian motion. Although Brownian motion may not be an accurate representation of all demand processes, it has three desirable properties: (i) two controllable moments, (ii) tabulated loss functions (because of its relation to the Normal distribution), and (iii) unimodality of the density of demand during any time period. Moreover, the assumption of Normal demands during time periods of known duration is quite common in the inventory literature, and there is empirical support for this approximation for high volume products.

In the first set of experiments, our goal was to ascertain whether the assumption of unimodality of the objective function *for a given n* is realistic. For the problem parameters given in Table 1, we computed the objective function values for values of T between $0.05T_{max}$ up to T_{max} in increments of $0.05T_{max}$ and checked the costs for unimodality. In all cases, the function was unimodal on this grid. Although this does not ensure unimodality in general, this provides strong support for the assumption of unimodality for the case of Brownian motion.

The problem parameters were selected to be consistent with the range of actual costs and demand characteristics in just-in-time operations in which trucks are the primary mode of transportation. The two values of A represent approximate costs of short- and medium-distance shipments. The ratios of C to A reflect penalty costs from using "on demand" shipments rather than contracted shipments on the low end, up to air freight charges for a package representing a small portion of a truckload at the high end. The values of h represent annual inventory holding costs for a truckload of commodities ranging from inexpensive, bulky raw materials to medium-priced electronic components shipped in protective containers. The parameter π was chosen so that the resulting service levels (probability of avoiding a stockout in a delivery cycle) would be 80% or larger for the anticipated values of T. The values of μ (number of truckloads per year) reflect a range of suppliers with deliveries as infrequent as once per month or as frequent as twice a week. Finally, the values of σ were chosen consistently with the notion that just-in-time systems should have relatively stable demand.

TABLE 1

Next, we selected 32 sets of parameters for which we might expect $n^* > 1$ in order to test for unimodality of $z(n, T^*(n))$. The selected problems have the following characteristics: (1) high emergency shipment costs; (2) high coefficient of variation; (3) low inventory holding costs; and (4) low shortage costs. These problem characteristics tend to make it desirable to make multiple-truck shipments so as to benefit from the effect of pooling variances over time, and from reduction of the frequency of exposures to potential emergency shipments, without incurring substantial inventory holding and shortage costs. We used the same grid search on T as described above.

Results from the second set of problems suggest that $z(n,T^*(n))$ is very well behaved for the case of Brownian motion. Consequently, a search of the type described above is likely to provide very good solutions. Of interest is the fact that in *all 32 problems*, the "optimal" value of n is greater than one, and for one problem $n^* = 10$. (Since we perform a grid search on T, $T^*(n)$ is approximate and consequently $z(n,T^*(n))$ is also.) Thus, instances of $n^* > 1$ may not be very unusual. This also suggests that reducing the coefficient of variation of the demand process and finding inexpensive means of providing for emergency shipments (when necessary) may be critical in implementing just-in-time economically. (Although inventory holding and shortage costs affect the optimal solution, these parameters are generally less controllable than the other two.) We also note that when $n^* >> 1$, $z(n,T^*(n))$ is very insensitive to n near the optimal solution, and in many of these cases, the solution with $n = n^*$ represents a 40% reduction in costs over the solution with $n = 1$.

To provide the reader a sense of how n^* changes with the parameters, we present the data and the optimal solutions for these 32 problems in Table 2. We also present the resulting average vehicle utilization, which provides information on the amount of transportation slack in the system, and the service level (probability of avoiding a stockout). Even within this small set of problems, some patterns are clear. First, as transportation costs become more expensive relative to inventory and shortage costs, the number of vehicles per shipment increases. This permits the system to take advantage of risk pooling over time which has two effects: to increase the average vehicle utilization for contracted shipments, and to reduce the expected number of emergency vehicles. Second, the shortage cost has little effect on n^* and T^*, and principally affects the amount of safety stock in the system. This may suggest that the characteristics of the transportation contract can be established first, and then the safety stocks set accordingly. More research is needed to test such a heuristic.

TABLE 2

Third, the larger the standard deviation of the demand process, the lower the optimal average utilization tends to be. This result is consistent with the intuition that more slack is required in a system with higher variability. Finally, as expected, an increased C/A ratio tends to lead to a slightly larger number of vehicles per shipment, principally to allow for more slack in the vehicles, which in turn reduces the frequency of emergency shipments.

8. CONCLUSIONS

There is some similarity between these results and those obtained in more traditional inventory models. In continuous review, order quantity-reorder point models, for problems with high demand variability, it is not unusual to obtain optimal solutions for the order quantity that are much larger than the EOQ. By increasing the order quantity, the frequency of exposures to potential stockout situations is decreased, thereby allowing a reduction in the amount of safety stock. This tradeoff was discussed by Magee (1956). In our problem, increasing the time between shipments (and correspondingly, the number of vehicles per shipment) does not appreciably affect the safety stock, but it has a substantial effect on the number of emergency shipments.

There is also a slight similarity to risk pooling results obtained from the (R,T) policy described by Johnson and Montgomery (1974). In this single item inventory model, demand follows some type of stationary stochastic process. There is a fixed charge per order, costs for cycle stock and safety stock, and per unit costs for shortages. The problem is to find the time between orders (T), and the order-up-to point (R). In this problem, increasing T allows for risk pooling over time, which leads to a reduction in the total safety stock and shortage costs per unit time, as well as the obvious reduction of total setup costs.

We have assumed here that the assembly facility uses an (S,T) policy, and practical considerations may preclude the use of more complex policies. However, in instances where short-shipments are allowed, the optimal policy is not of the order-up-to form, since the cost of emergency shipments may make short-shipments optimal. Research is needed to determine characteristics of the optimal inventory policy in these situations, and its impact on the nature of the transportation contract.

Acknowledgement

This research was supported in part by a U.S. automobile manufacturer through a contract to the University of Michigan.

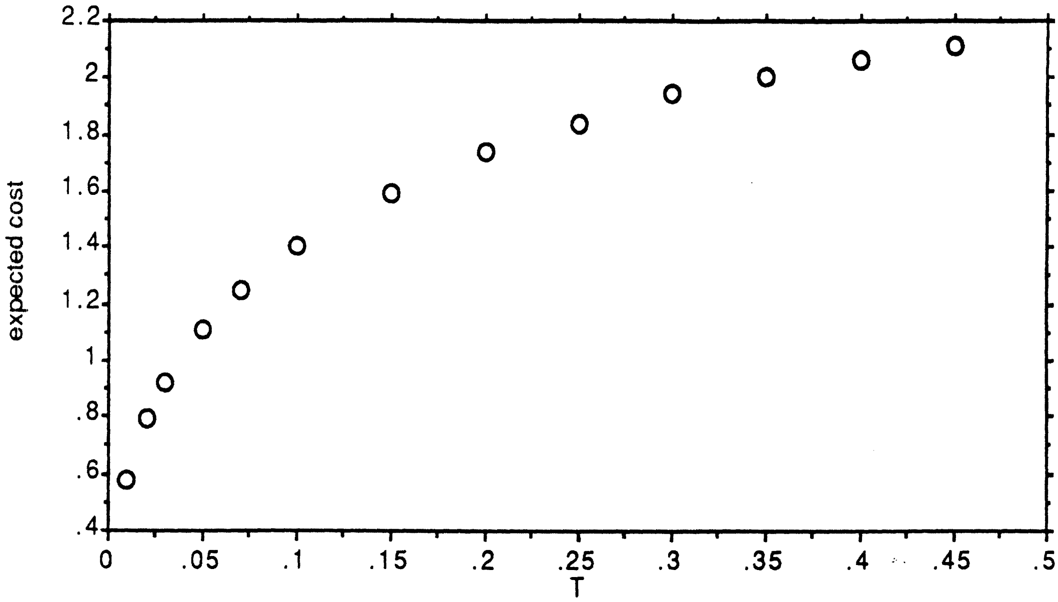

Figure 1

Example Plot of G(T) for Brownian Motion

$\mu = 10, \sigma^2 = 4, h = 1, \pi = 2$

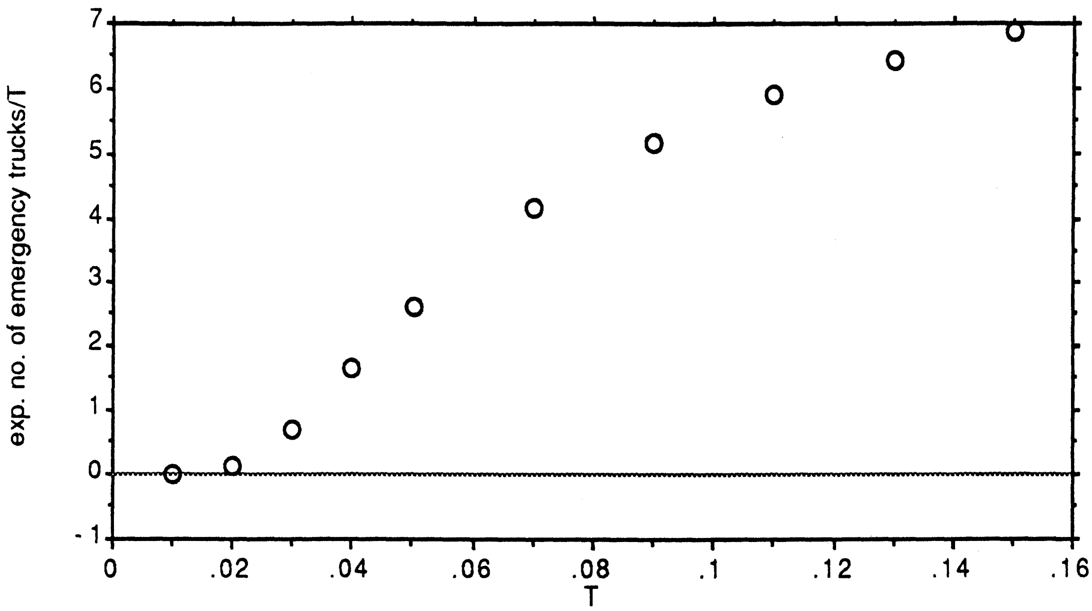

Figure 2

Example Plot of the Expected Emergency Shipment Cost Per Unit Time as a Function of T for Brownian Motion

$\mu = 10, \sigma^2 = 4$

Table 1

Problem Parameters for First Experimental Study

Parameter	Values
A	125, 500
C	1.25, 2.5, 3.75 times A
μ	16, 40, 100
σ	0.05, 0.10,.0.15 times μ
h	4000, 9000
π	0.5, 1.0 times h
n	1, 2, 4

Table 2

Problem Parameters and Optimal Solutions for Second Experimental Study

A	C	h	π	μ	σ	n*	T*	avg. vehicle utilization	service level
125	312.5	4000	2000	40	4.0	2	.0335	.67	.937
125	312.5	4000	4000	40	4.0	2	.0335	.67	.968
125	312.5	4000	2000	40	6.0	2	.0307	.61	.942
125	312.5	4000	4000	40	6.0	2	.0307	.61	.970
125	312.5	4000	2000	100	10.0	3	.0238	.79	.954
125	312.5	4000	4000	100	10.0	3	.0217	.72	.979
125	312.5	4000	2000	100	15.0	3	.0217	.72	.958
125	312.5	4000	4000	100	15.0	3	.0217	.72	.979
125	468.75	4000	2000	40	4.0	2	.0307	.61	.942
125	468.75	4000	4000	40	4.0	2	.0307	.61	.970
125	468.75	4000	2000	40	6.0	2	.0280	.56	.947
125	468.75	4000	4000	40	6.0	2	.0280	.56	.973
125	468.75	4000	2000	100	10.0	3	.0195	.65	.962
125	468.75	4000	4000	100	10.0	3	.0195	.65	.981
125	468.75	4000	2000	100	15.0	4	.0250	.63	.952
125	468.75	4000	4000	100	15.0	4	.0225	.56	.978
500	1250	4000	2000	40	4.0	4	.0870	.87	.852
500	1250	4000	4000	40	4.0	3	.0616	.82	.942
500	1250	4000	2000	40	6.0	4	.0870	.87	.852
500	1250	4000	4000	40	6.0	4	.0791	.79	.927
500	1250	4000	2000	100	10.0	6	.0551	.92	.900
500	1250	4000	4000	100	10.0	6	.0551	.92	.948
500	1250	4000	2000	100	15.0	8	.0707	.88	.876
500	1250	4000	4000	100	15.0	7	.0661	.94	.938
500	1875	4000	2000	40	4.0	4	.0791	.79	.863
500	1875	4000	4000	40	4.0	4	.0791	.79	.927
500	1875	4000	2000	40	6.0	5	.0972	.78	.837
500	1875	4000	4000	40	6.0	5	.0972	.78	.911
500	1875	4000	2000	100	10.0	9	.0750	.83	.870
500	1875	4000	4000	100	10.0	7	.0595	.85	.944
500	1875	4000	2000	100	15.0	10	.0791	.79	.863
500	1875	4000	4000	100	15.0	9	.0675	.75	.937

REFERENCES

Cook, T.M. and R.A. Russell (1978), "A Simulation and Statistical Analysis of Stochastic Vehicle Routing with Timing Constraints," *Decision Sciences* 9(4), 673-687.

Federgruen, A. and P. Zipkin (1984), "A Combined Vehicle Routing and Inventory Allocation Problem," *Operations Research* 32(5), 1019-1037.

Golden, B.L. and Stewart, Jr., W. (1978), "Vehicle Routing with Probabilistic Demands," in Proceedings of Computer Science and Statistics: Tenth Annual Symposium on the Interface (D. Hogben and D.W. Fife, eds.), National Bureau of Standards, Washington, D.C.

Golden, B.L. and J.R. Yee (1979), " A Framework for Probabilistic Vehicle Routing," *AIIE Transactions* 11(2), 109-112.

Herron, D.P. (1979), "Managing Physical Distribution for Profit," *Harvard Business Review* 57(3), 121-132.

Hill, A.V. and T.E. Vollman (1986), "Reducing Vendor Delivery Uncertainties in a JIT Environment," *Journal of Operations Management* 6 (4), 381-392.

Jackson, G. (1983), "Just-in-Time Production: Implications for Logistics Managers," *Journal of Business Logistics* 4, 1-19.

Johnson, L.A. and D.C. Montgomery (1974), Operations Research in Production Planning, Scheduling, and Inventory Control, John Wiley & Sons, New York.

Lieb, R. and R.A. Miller (1988), "JIT and Corporate Transportation Requirements," *Transportation Journal* 27, 5-10.

Tillman, F.A. (1969), "The Multiple Terminal Delivery Problem with Probabilistic Demands," Transportation Science 3 (3), 192-204.