

MULTI-PRODUCT CAPACITATED PRODUCTION
PLANNING WITH RANDOM DEMAND

Candace Arai Yano
Department of Industrial & Operations Engineering
University of Michigan
Ann Arbor, MI 48109-2117

Technical Report 86-12

April 1986
Revised November 1986

MULTI-PRODUCT CAPACITATED PRODUCTION PLANNING
WITH RANDOM DEMAND

Candace Arai Yano
Department of Industrial and Operations Engineering
University of Michigan
Ann Arbor, MI 48109-2117

April 1986
Revised November 1986

This work was supported in part by National Science Foundation Grant
DMC 8504644.

MULTI-PRODUCT CAPACITATED PRODUCTION PLANNING
WITH RANDOM DEMAND

ABSTRACT

Traditionally, production-inventory models in which demand is random have included shortage costs or customer service requirements to ensure that production satisfies demand on a timely basis. In systems where supplier facilities feed critical or bottleneck assembly facilities, anything less than 100% satisfaction of assembly needs may result in shutting down entire production facilities. Thus, these traditional models will not provide implementable results. The supplier facilities actually deal with this need for 100% service through a combination of safety stock and rescheduling.

We investigate the problem of determining planned production intervals and safety stock quantities which provide the appropriate balance between inventory costs and setup costs in a single machine, multi-product fabrication system facing highly uncertain demand. Capacity constraints and setup times are additional complicating factors which we consider.

We develop a two-stage approach to the planning problem in this paper. First, safety stock is optimized for each candidate planned production interval. The distribution of actual production intervals resulting from the safety stock decision is considered explicitly in this stage. This provides input to another optimization problem in which jointly optimal planned production intervals are selected so as to minimize total cost subject to a capacity constraint. Computational results also are reported.

MULTI-PRODUCT CAPACITATED PRODUCTION PLANNING WITH RANDOM DEMAND

1.0 INTRODUCTION

Many manufacturing firms using Material Requirements or Manufacturing Resource Planning (MRP) systems have dual problems of frequent rescheduling and excessive inventories because of uncertain demand. In the earlier years of MRP, before the notion of safety stock became acceptable, many wondered about how to deal with frequent rescheduling (see for example, Mather, 1977). More recently, cyclic and "powers of two" fixed interval lot-timing policies (see Caie and Maxwell, 1981, and Roundy, 1984), have been proposed for situations with deterministic demand. While such policies are useful goals, there are occasions, particularly when demand and supply are highly uncertain, when rescheduling is necessary and even desirable because of the nature of the manufacturing process or customer service considerations.

Better forecasting and improved inventory data accuracy can reduce the magnitude of the rescheduling-inventory problem. However, unless demand and supply uncertainty are virtually eliminated, costs can be reduced through an appropriate balance of rescheduling and safety stock.

A typical problem that arises in make-to-order systems involves a critical facility, which (effectively) cannot be stopped, and facilities supplying parts to the critical facility. The critical facility may be a bottleneck process or an automated, paced assembly line. Production of input parts is planned according to forecasts, but forecast errors may be large. Thus, forecasts may be revised dramatically, and similarly the production schedules. It is the responsibility of the parts facilities to keep the critical facility operating in spite of their own capacity limitations and complex resource allocation problems. Overtime and extra shifts may be employed when necessary, but much of

the "real time" adjustment is done through rescheduling and safety stock. Because of the complexity of mating parts which are inputs to the critical facility, one cannot simply specify service objectives for each part. There are two possible approaches. The more desirable approach is to treat satisfaction of the dependent demand for parts as a hard constraint. The alternative is to carefully coordinate production at all component facilities, so that the critical facility can produce something, even if it is not specifically what it had planned to produce. This is feasible in a make-to-stock situation, but cannot work in a make-to-order environment. In any event, the parts facilities have a complex problem which must be solved frequently to handle dynamic changes.

Little analytical work has been done on this problem. Askin (1981) develops a heuristic approach for determining constant planned cycle lengths with consideration of the effects of these cycle lengths on safety stock. This approach permits shortages, and therefore requires no rescheduling. Graves (1981) considers a multi-product problem in which shortages are permitted and penalties are charged for backorders. Bitran and Yanasse (1984) show that some forms of the static capacitated production planning problem with stochastic demand can be solved within 10% of optimality through solution of the deterministic problem. Yano and Carlson (1986) develop a heuristic procedure to determine safety stock quantities in a simple uncapacitated assembly system when rescheduling occurs each period, but do not optimize the balance between safety stock and rescheduling.

We investigate the setup cost-inventory holding cost tradeoff in capacitated, dynamically changing single-machine systems with the objective of understanding how costs and system characteristics affect this tradeoff. In this paper we address the planning problem, that is, determining planned

production intervals and safety stock quantities consistent with them. This must be done in view of capacity constraints and the necessity of rescheduling. Detailed scheduling and rescheduling in a dynamic environment is a much more complex problem. We are currently investigating it and will report on it in a later paper. The reason for solving the planning problem is to take into account aggregate tradeoffs and constraints first, so that there is opportunity for rational decision-making in the dynamic problem. Without such planning, the dynamic decisions primarily involve "putting out fires," which, in general, is an extremely costly approach. In addition, we believe that the results of the single-machine investigation will provide insight into multi-machine problems. We note, however, that in the scenario described in the next section, there are several facilities which have only one dominant or bottleneck machine, so the results have many direct applications.

We describe the problem scenario which is the motivation for this work in section 2. Model assumptions are detailed in section 3. An approach to the steady-state problem is discussed in section 4, and related computational results are included in section 5. We conclude with a summary and discussion in section 6.

2.0 PROBLEM SCENARIO

The motivation for this research is a manufacturing system consisting of an automated, paced assembly facility and several component fabrication facilities. The finished product is semi-custom, with some standard parts and some options. Demand for the finished product is uncertain, both in total volume and in option mix.

Detailed component demand forecasts are transmitted from the assembly facility to the component facilities weekly in a fashion similar to MRP product explosions. Component demand quantities are firm for a very short horizon

(i.e., a day to a week), but the remaining information constitutes forecasts only. Analysis of forecasts and actual demands for specific parts indicates that forecasts for imminent periods are not much more accurate than forecasts for periods farther out in the horizon, and that both have large (absolute and relative) forecast errors.

Scheduling of component fabrication facilities is difficult. Some of the facilities have long setup or changeover times, and in several cases, scrap losses during setups and learning curve effects result in "out of pocket" setup costs which are high. Thus, relatively long production runs are desired. However, components are "ordered" by the assembly facility periodically, and the frequency of shipment from each component facility is determined by transportation and inventory considerations. Nearly all shipments are composed of mixed parts. The combination of high setup costs and highly uncertain component demands make adaptation by safety stock or rescheduling alone uneconomical, if not infeasible.

The model developed in the next section can be viewed as relating to planning of setups of one type (either major (family) or minor (item)) which are not sequence dependent.

3.0 MODEL ASSUMPTIONS

To focus on the effects of uncertain demand, we make some assumptions in this study which we hope to drop in future research. In this paper, we assume that each item must be processed by only one (dominant) machine. We also assume that the component facility has an infinite supply of input materials, so supply uncertainty is not an issue. In reality, supply uncertainty appears to be a serious problem in the scenario described above. However, this "supply uncertainty" arises primarily because orders made by the component facility for

incoming parts and materials are no longer consistent with the revised schedule at the assembly facility. Thus, supply is indeed a problem, but the source of the problem is uncertain demand.

We assume that each component facility receives an "order" from the assembly facility periodically and that it must supply those parts by the end of the same period (i.e., with the next shipment) from current production or from inventory. Thus, if inventory is insufficient, a production run is mandatory (i.e., backorders are not permitted).

Since it is possible to produce many products (and in reasonable quantities) in a shipping interval, the production plus transport leadtimes can be treated as a constant for each component facility. In fact, we assume that the production plus transport leadtime is less than one period, which is almost essential to ensure 100% service. This is a strong assumption, but it is satisfied by all products in our application, and is true in a variety of manufacturing environments.

We assume that setup times are positive and therefore affect capacity and labor utilization. If setup times are very small, as in some highly automated systems, the effects of rescheduling on capacity utilization are much less severe, so that frequent rescheduling is not as problematic. We also assume that setup times and setup costs are sequence independent.

Costs to be included in the model are: (i) setup costs, (ii) inventory holding costs charged on average inventory, and (iii) overtime premium costs (when overtime is permitted). Regular time costs and the non-premium portion of overtime are treated as sunk costs since all demand must be satisfied.

We assume that when overtime is permitted, an overtime premium fairly reflects the penalty for exceeding normal capacity limits. Over the short term, overtime and weekend work are the only practical workforce adjustments. There

is, of course, a limit on overtime capacity. Yet, addition of two hours to each eight-hour shift on weekdays and extra shifts on both weekend days can increase capacity of a two-shift operation by 65%. Thus, within the usual limits of demand fluctuations, overtime is effectively unconstrained. Nevertheless, the actual schedule must satisfy this constraint in each period as well as satisfying it on average.

The objective is to minimize expected total costs. There may, of course, be rescheduling costs incurred beyond the necessary setup costs resulting from a schedule change. On occasion, there is added inconvenience, but no true additional costs. In other circumstances, it may be necessary to expedite material inputs, which may then cause rescheduling at other supply facilities. While it would be desirable to include these costs, they are usually so difficult to quantify that we believe a better understanding of them is needed before they can be included appropriately in this model.

4.0 THE CAPACITATED STEADY-STATE PROBLEM

We will focus on stationary, but highly variable demand processes in this study. The primary reason for taking this approach is that it is not the non-stationarity that is responsible for the scheduling difficulties, but the magnitude of the forecast errors. In fact, the one-week-ahead forecast errors for some components are so high that it is difficult to distinguish between forecast errors and non-stationarity. Thus, it is reasonable, at least initially, to model the process as if it were stationary but highly variable. We will assume here that demand has a normal distribution, but the approach can be generalized to other distributions.

To facilitate scheduling at each component facility, it has been suggested (see, for example, Elmaghraby, 1978 and Caie and Maxwell, 1981) that each

production run represent some constant time supply. Such a policy simplifies the capacitated production planning problem to one of coordinating a number of cyclic production schedules, one for each item. Even when demand is stochastic and each production run may vary both in quantity and occasionally in terms of cycle length (production interval), having a target cycle length can help in production smoothing and in procurement of input materials. In addition, this concept of a target cycle length is extremely helpful in analyzing the tradeoff between setup costs and inventory holding costs.

Problem Formulation

Notation:

- S_i = setup cost for item i
- h_i = holding cost charged on end-of-period inventory of item i
- λ = overtime premium for equivalent of regular capacity
- τ_i = setup time for item i
- P_i = processing time per unit of item i
- I_{it} = inventory of item i at end of period t
- D_i = average demand per period for item i
- σ_i = standard deviation of demand for item i for one period
- k_i = safety stock multiplier for item i , where the safety stock quantity is $k_i \sqrt{n_i - 1} \sigma$
- n_i = planned (or target) production cycle for item i (in periods)
- η_i = actual production cycle length (in periods)

For notational simplicity, we drop the "i" subscripts in the following development and re-introduce them when needed later in the paper.

The objective criterion in this problem is minimization of average cost per period, which we represent as average cost per cycle/average cycle length. Since each production cycle for an item does not correspond to time between

adjacent regeneration points, this representation is not exact. However, since the system regenerates infrequently, there are few practical alternatives. Nevertheless, this approximate representation has been used successfully in many other similar applications.

The system operates as follows. At the beginning of each period a random demand d_t is observed. If $I_{t-1} < d_t$, a production run is made to raise the inventory position to $d_t + (n-1)D + k\sqrt{n-1} \sigma$. In other words, if there is a positive net requirement in period t , a production run is made for an n -period supply plus safety stock. Observe that demand in the current period is known before production is scheduled, whereas demand in the next $n-1$ periods is still uncertain. Safety stock is represented as $k\sqrt{n-1} \sigma$ because there are $n-1$ periods for which safety stock must provide protection.

We first derive the expected cost per cycle and the distribution of cycle lengths for a single item as a function of the planned production cycle and the safety stock quantity. First consider the average cycle stock as a function of the planned production cycle and the actual cycle length. If we assume that when an order is completed, the cycle stock on hand is brought up to a value of nD (on average) and that demand occurs at a constant rate during the (actual) cycle, then average cycle stock inventory is clearly $nD/2$ for any actual cycle length. Similarly, the safety stock quantity does not depend upon the actual cycle length since it is "optimized" for the planned cycle length. We can write the cost per cycle as:

$$S + E(\eta)h[Dn/2 + k\sqrt{n-1} \sigma],$$

where $E(\eta)$ is the expected cycle length.

We next derive the probability that the actual cycle length (denoted η) is equal to T for some $0 < T < \infty$, and use this information to derive $E(\eta)$.

Clearly, the actual cycle length will be T if there was sufficient stock to

cover demand through the first T periods of the cycle, but not enough to cover demand in the $(T+1)$ st period in the cycle. This is the same as the first passage time of Brownian motion with a drift (with drift parameter $\mu=D$ and diffusion coefficient σ^2) occurring between time $T-1$ and T (see Karlin and Taylor, 1975, p. 355). Since demand in the current period is known, we start the next period with $(n-1)D + k\sqrt{n-1} \sigma$ units of inventory. The actual cycle length is the time required to use up this inventory plus one period (i.e., the first period for which demand is known). Recall that demand is observed at the beginning of the period, so that demand in the subsequent period can be viewed as the difference between the state (inventory position) now and the state at the beginning of the next period. Let $z = (n-1)D + k\sqrt{n-1} \sigma$. Then we can write

$$P[\eta = T] = \int_{T-1}^T \frac{z}{\sigma \sqrt{2\pi t^3}} e^{-(z-Dt)^2 / 2\sigma^2 t} dt \quad (1)$$

$$E(\eta) = \sum_{j=0}^{\infty} j P[\eta = j] \quad (2)$$

$$E(\text{cycle cost}) = S + E(\eta)h[Dn/2 + k\sqrt{n-1} \sigma] \quad (3)$$

The objective function is simply the ratio of equations (3) and (2). The expression in (1) must be evaluated numerically because there is no closed form representation of it.

For any value of n , we can attempt to optimize k using a one-dimensional search procedure. In the Appendix we show that the average cost per period is likely to be convex in k for fixed n . We also expect that for a given i the cost per period (using the optimal safety stock quantity) is discretely convex in n , and present numerical evidence of this relationship in the next section. For the moment, let us assume that we can obtain a near-optimal value of k for each n , and can therefore express the expected cost per period as a function of

the planned cycle length alone.

We can now use this information to select jointly optimal planned cycle lengths which satisfy the capacity constraint. Before continuing, it is important to note that it is desirable from an administrative perspective to make n an integer, since non-integral planned cycle lengths are often extremely difficult to implement. (It is difficult to tell a worker to produce a 3.7 week supply). Moreover, it is difficult to coordinate production schedules for products having non-integer production cycles. In the remainder of the paper, we will assume that n is integer. Nevertheless, the expected cycle length, in general, will not be an integer. For purposes of capacity planning, the expected cycle lengths provide much more useful information than the planned cycle lengths, since they indicate how often, on average, a setup will occur.

Finding Optimal Planned Production Intervals

To find the jointly optimal planned cycle lengths for each item given an aggregate capacity constraint, we might formulate the problem as:

$$\begin{aligned}
 \text{(P1)} \quad & \text{minimize } \sum_i \sum_n c_{in} y_{in} \\
 & \text{subject to } \sum_i \sum_n a_{in} y_{in} \leq 1 \\
 & \sum_n y_{in} = 1, \quad \forall i \\
 & y_{in} \in \{0,1\}, \quad \forall i,n
 \end{aligned}$$

where c_{in} = expected cost per period of item i using planned cycle length n

T_{in} = expected cycle length of item i if it uses planned cycle length n

$r_{in} = \tau_i + P_i T_{in} D_i$ = expected processing and setup time for one average cycle of item i if it uses planned cycle length n

$a_{in} = r_{in}/T_{in}$ = fraction of capacity devoted to item i if it uses planned cycle length n

$$y_{in} = \begin{cases} 1 & \text{if item } i \text{ uses planned cycle length } n \\ 0 & \text{otherwise} \end{cases}$$

The objective is to minimize expected total system cost per period subject to (i) an aggregate capacity constraint, (ii) assignment constraints (one cycle length for each item), and (iii) binary constraints on the assignments.

Observe that there is only one aggregate capacity constraint. Therefore, it would be quite easy to relax this constraint using a Lagrange multiplier. The problem then becomes (after rearrangement of terms):

$$(P2) \quad \text{minimize} \quad \sum_i \sum_n (c_{in} + \lambda a_{in}) y_{in} - \lambda$$

$$\text{subject to} \quad \sum_n y_{in} = 1, \quad \forall i$$

$$y_{in} \in \{0,1\}, \quad \forall i,n$$

where $\lambda =$ Lagrange multiplier.

Given any value of λ , we only need to find the cheapest n for each i since there are no additional constraints on the y_{in} s. Fortunately, in this application we know the value of λ -- it is the cost of overtime for a one unit change in capacity. Since capacity is measured in percentages of total capacity, λ is the equivalent cost of overtime for the entire regular time operation. This may be a large value, but is reasonably easy to estimate. Thus, this problem can be solved by inspection. Some computational effort is required to determine the a_{in} s. However, the primary difficulty involves finding the optimal value of k for each i and n . If one would like to satisfy the capacity constraint with no overtime, the problem can be solved quite easily by searching for a value of λ such that the constraint is satisfied at equality. Any reasonable procedure to do this (e.g., branch and bound in conjunction with subgradient optimization)

would be acceptable since for a given λ the problem can be solved by inspection.

Determining Candidate Planned Cycle Lengths

Although equations (1) - (3) are simple conceptually, it is not easy to specify in advance the values of n for which the corresponding optimal k s should be found. However, it would be reasonable to expect that values of n near the optimal unconstrained deterministic problem (i.e., cycle implied by the economic production quantity) would be advantageous. Since the system is assumed to be capacity constrained, it is unlikely that smaller values of n would be chosen. On the other hand, it is quite conceivable that an order cycle up to 2 or more times the EPQ solution might be used since the relative cost penalty would be small, but capacity utilization might be eased considerably. The choice of candidate planned cycles lengths could vary from application to application and with other administrative considerations, such as ease of implementation.

Comparison with Deterministic Formulation

A deterministic problem can be formulated as P1 with appropriate definitions modified. The values of c_{in} in the deterministic problem will tend to be smaller than those of the stochastic problem for two reasons. First, $E(\eta)$ in the stochastic problem will tend to be less than n . This phenomenon is illustrated in the next section. Secondly, there is no safety stock in the deterministic problem.

In addition, since the expected cycle lengths are smaller than the planned cycle lengths, the values of a_{in} in the deterministic problem are smaller than the corresponding values in the stochastic problem. Thus, a solution which is feasible for the stochastic problem is feasible for the deterministic problem, but not vice versa.

5.0 CONVEXITY AND STABILITY ISSUES

We mentioned two conjectures in the last section that we discuss more fully in this section:

- (1) For fixed i and n , the expected cost per period is convex in k ; and
- (2) For fixed i , the values of c_{in} are discretely convex in n .

In the Appendix, we show that under reasonable assumptions, there is strong analytical support for the first conjecture. The second conjecture cannot be proved rigorously because the optimal value of k (for each i and n) cannot be expressed in closed form. We performed an empirical study of this relationship using parameters $D = \mu = 200$, $h = 1.0$, $S = 100, 400, 900$, or 1600 , and $\sigma = 10, 30$, or 50 . The chosen values of S correspond to cycles of 1, 2, 3, or 4 periods, respectively, for the deterministic problem. Some representative results appear in Figure 1. It appears reasonable to assume that the c_{in} values are discretely convex in n for integer values of n . This helps to limit candidate planned cycle lengths to those near to or greater than the minimum-cost planned cycle length.

FIGURE 1

The results illustrated in Figure 1 also indicate that the c_{in} values are much less sensitive to n (particularly to the right of the minimum point) than would be predicted by the deterministic model. This has an important implication: the penalties for deviating from the unconstrained solution (i.e., increasing the cycle times) to satisfy a capacity constraint are smaller than in the deterministic problem. We also found that the unconstrained minimand was at least as large as the solution indicated by the corresponding (deterministic) EPQ, making the unconstrained optimal solution closer to capacity-feasibility than the solution to the deterministic problem.

The results have some other interesting characteristics. We observed that

some safety stock was desirable for situations with low demand variability and with planned cycles lengths not much larger than the EPQ would indicate. A small quantity of safety stock can help to reduce the number of "short" cycles. For situations with high demand variability and/or long planned cycles lengths (relative to the EPQ), it is optimal to carry no safety stock. Unfortunately, this leads to expected cycle lengths of an integer plus a half when the distributions of the forecast errors are symmetric. This is possibly the most difficult situation to deal with from a scheduling perspective since this roughly means that it is equally likely that a "short" or a "normal" cycle will occur. If this is true, then there is a 50% chance that the schedule for each item may change, resulting in only a $(0.5)^N$ chance that an entire rotation can be executed as planned.

This phenomenon led us to consider planned cycle lengths with the "integer plus one half" characteristic. These alternatives have safety stock which is disguised as additional cycle stock. Experimental results indicate that for these planned cycle lengths, it is usually optimal to have no additional safety stock, leading to an integer valued expected cycle length--and a relatively high probability that the actual cycle length will be equal to the expected cycle length. This makes scheduling easier, as well as making the schedule much easier to adhere to. Of course, some planned cycle lengths of this type (e.g., 1.5) are uneconomical and can be excluded from consideration. There are also administrative and implementation considerations to be addressed.

Companies concerned about schedule stability might consider using planned cycle lengths only of this type in the optimization problem in section 5. The economic cost of the additional stability can be determined quite easily by solving the problem in both ways.

6.0 COMPUTATIONAL RESULTS

We attempted to gather usable data from the facility described earlier in the paper. Unfortunately, the only data readily available were for production quantities, not "demand" quantities. The data collection effort was further complicated by variable yield rates and stochastic setup times. (The time to set up a machine increases if adjacent machines are also being set up at the same time). While we intend to pursue modeling of and development of a solution approach for the real problem, it seemed reasonable to "prove out" the approach described in this paper as a rational first step in better planning.

We adapted some data from an example in Rogers (1958) by increasing the demands and adding setup times so that the machine capacity would be a constraint (see Table 1). For coefficients of variation of 0.1 and 0.3, we determined optimal safety stock levels and the resulting c_{in} and a_{in} values. We then used a simple trial-and-error approach to finding a value of λ which would make the solution capacity-feasible for each coefficient of variation. The optimal solution for both problems is $n^* = (6,12,3,7,7)$. An almost-feasible solution (using .01% more capacity than permitted) is $n^* = (6,11,3,7,7)$ and the best powers-of-two policy is $n^* = (8,8,4,4,8)$ which incurs a penalty of approximately 1.5% of total costs and uses 1% less capacity.

TABLE 1

7.0 SUMMARY AND CONCLUSIONS

We have developed a procedure to determine jointly optimal planned cycle lengths and safety stock quantities for several products which are produced on the same machine. The products have uncertain demand which must be satisfied within a very short period of time. The solution procedure has two parts: optimizing safety stock for each candidate planned cycle length for each product; and selecting the minimum cost set of planned cycle lengths consistent

with the capacity constraint.

Computational results indicate that the expected cycle lengths are shorter than planned, even when optimal safety stock quantities are included. Thus, the (unconstrained) optimal planned cycle lengths are generally longer than those from a solution to the deterministic problem. That is, the optimal solution compensates for the fact that the expected cycle lengths are shorter than the planned cycle lengths.

Results also indicate that positive safety stock is optimal in situations with low demand variability and if one chooses to use planned cycle lengths which are small compared with the EPQ solution. In such situations a small amount of safety stock can significantly reduce the number of "short" cycles. In other cases, it is more economical to simply reschedule as needed. It was also suggested that planned cycle lengths of the "integer plus one half" type would lead to integer-valued expected cycle lengths and a greater likelihood that the actual cycle length would be equal to the expected cycle length. This approach in which extra cycle stock serves as safety stock could lead to much more stable schedules.

Further research is needed on the short-term scheduling aspect of the problem. Although the solution to the production planning problem may satisfy aggregate annual capacity constraints, there is no guarantee that the various production cycles will fit together well (even if everything occurs as planned). Moreover, since the production cycles are random, it may occur by chance that production requirements in any period far exceed capacity. Therefore, it may be necessary to produce some products earlier than planned. We are currently investigating methods for dealing with the short-term problem while taking into account long-term capacity considerations which are reflected in the analyses here.

Additional research is also needed for multi-product systems where demand is uncertain but backorders are permitted, and for multi-machine problems.

7.0 ACKNOWLEDGEMENT

The author would like to express appreciation to the National Science Foundation for support under grant DMC-8504644 and to two anonymous referees for their helpful comments on an earlier version of this paper.

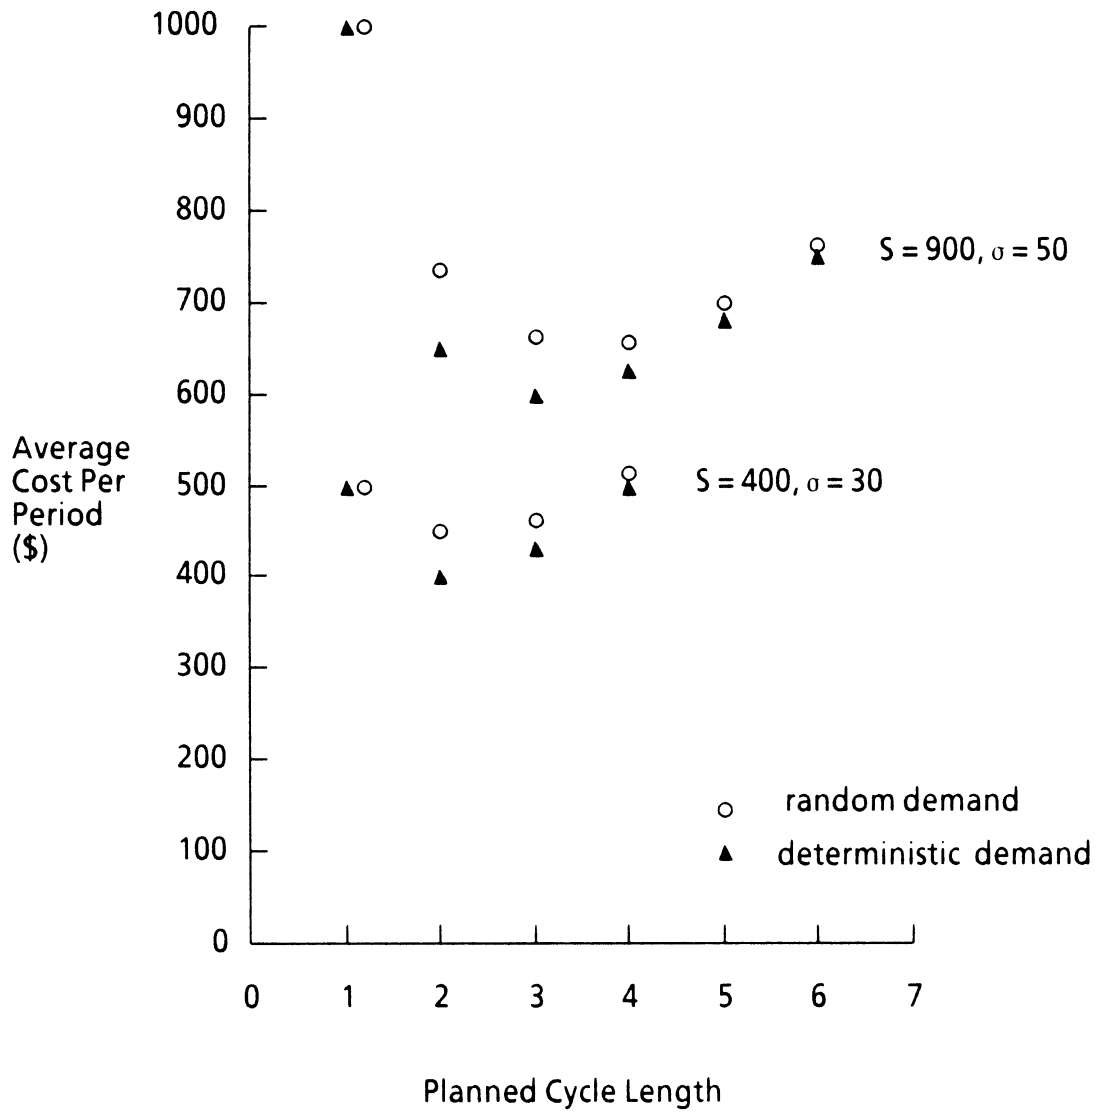


Figure 1
Examples of c_{in} Values

TABLE 1

Problem Data

Item	S_i	h_i	d_i (per day)	Prod. rate (per day)	Setup time (days)
1	50	.05	80	500	.15
2	75	.02	160	1000	.225
3	50	.10	320	800	.15
4	100	.20	40	2500	.30
5	150	.05	240	3750	.45

REFERENCES

- Askin, R.G., (1981), "A Procedure for Production Lot Sizing with Probabilistic Dynamic Demand," AIIE Transactions 13(2) 132-137.
- Bitran, G.R. and H.H. Yanasse, (1984), "Deterministic Approximations to Stochastic Production Problems," Operations Research 32(5), 999-1018.
- Caie, J. and W.L. Maxwell, (1981), "Hierarchical Machine Load Planning," Studies in the Management Sciences 16, North Holland.
- Carlson, R.C. and C.A. Yano, (1986), "Safety Stock in MRP Systems-Systems with Emergency Setups for Components," Management Science 32(4), 403-412.
- Elmaghraby, S.E., (1978), "The Economic Lot Scheduling Problem (ELSP): Review and Extensions," Management Science 24(6), 587-598.
- Graves, S.C., (1981), "The Multi-Product Production Cycling Problem," AIIE Transactions 12(3), 233-240.
- Karlin, S. and H.M. Taylor (1975), A First Course in Stochastic Processes, New York: Academic Press, Inc.
- Mather, H., (1977), "Reschedule the Schedules You Just Rescheduled - A Way of Life for MRP?", Production and Inventory Management (1), 69-79.
- Maxwell, W.L. and H. Singh, (1983), "The Effect of Restricting Cycle Times in the Economic Lot Scheduling Problem," IIE Transactions 15(3), 235-241.
- Rogers, J. (1958), "A Computation Approach to the Economic Lot Scheduling Problem," Management Science 4(3), 264-291.
- Roundy, R.O., (1984), "98% Effective Power-of-Two Lot Sizing for Multi-Product Multi-Stage Production/Inventory Systems," Technical Report No. 642, School of Operations Research and Industrial Engineering, Cornell University.

APPENDIX

In this appendix, we show that under reasonable assumptions and for fixed n , the cost per unit time is likely to be a convex function of k . Observe that for fixed i and n , the expected cost per period is

$$h[nD/2 + k\sqrt{n-1} \sigma] + S/E(\eta). \quad (A-1)$$

The first partial derivative of this with respect to k is

$$h\sqrt{n-1} \sigma - S \partial E(\eta) / \partial k / [E(\eta)]^2$$

and the second partial derivative with respect to k is

$$- S \partial^2 E(\eta) / \partial k^2 / [E(\eta)]^2 + 2S [\partial E(\eta) / \partial k]^2 / [E(\eta)]^2$$

To simplify the notation, let

$$a = \sum_{T=1}^{\infty} T \int_{T-1}^T (\sqrt{2\pi} \sigma t^{5/2})^{-1} e^{-(z-Dt)^2 / 2\sigma^2 t} dt$$

$$b = \sum_{T=1}^{\infty} T \int_{T-1}^T (\sqrt{2\pi} \sigma t^{7/2})^{-1} e^{-(z-Dt)^2 / 2\sigma^2 t} dt$$

It can be shown that

$$\begin{aligned} \partial E(\eta) / \partial k &= \sqrt{n-1} E(\eta) \{D/\sigma + \sigma / [(n-1)D + k\sqrt{n-1} \sigma]\} \\ &= \sqrt{n-1} [(n-1)D + k\sqrt{n-1} \sigma]^2 a \end{aligned} \quad (A-2)$$

and

$$\begin{aligned} \partial^2 E(\eta) / \partial k^2 &= (n-1) \{ [D^2/\sigma^2 + 2D / [(n-1)D + k\sqrt{n-1} \sigma]] E(\eta) \\ &= 2[(n-1)D + k\sqrt{n-1} \sigma] \sigma a \\ &+ [(n-1)D + k\sqrt{n-1} \sigma]^3 b/\sigma \end{aligned}$$

To show that (A-1) is convex, we need to show that

$$2[\partial E(\eta) / \partial k]^2 / E(\eta) - \partial^2 E(\eta) / \partial k^2 \geq 0$$

This is extremely difficult to do because of the forms of a and b. Even if we recognize that $\partial E(\eta)/\partial k \geq 0$ (which is intuitively evident, since the runout time is stochastically increasing with k), and that

$$a > \int_0^{\infty} (\sqrt{2\pi} \sigma t^{5/2})^{-1} e^{-(z-Dt)^2/2\sigma^2 t} dt \quad (A-3)$$

$$= [(n-1)D + k\sqrt{n-1} \sigma]^{-1}$$

b still remains in the formula. The inequality above arises from the fact that each probability term (i.e., definite integral) in the expression for a is multiplied by the upper limit of the integral, whereas in the right hand of (A-3), t takes on values between the lower and upper limits.

The results so far with regard to convexity are inconclusive. If, however, we make a few reasonable approximations to simplify the formulas, we can show that $\partial^2 E(\eta)/\partial k^2$ is likely to be non-positive. Using the results of (A-2) and (A-3), we have

$$\partial E(\eta)/\partial k < \sqrt{n-1} \{E(\eta)[D/\sigma + \sigma/[(n-1)D + k\sqrt{n-1} \sigma]] - [(n-1)D + k\sqrt{n-1} \sigma]\} \quad (A-4)$$

Now

$$[(n-1)D + k\sqrt{n-1} \sigma]/D \cong E(\eta) = c(k)$$

where c(k) is a decreasing function of k. (Observe that when k = 0, the expression on the left hand side is equal to E(η) = 0.5, so c(0) = 0.5). Substituting this into (A-4) we get

$$\partial E(\eta)/\partial k < \sigma E(\eta)/D[E(\eta) - c(k)] + c(k)D - (1-1/\sigma)DE(\eta)$$

Note that the first term on the right hand side is decreasing as k increases since the ratio $E(\eta)/[E(\eta) - c(k)]$ decreases as k increases. (The numerator of the ratio increases, but the denominator increases more rapidly). Clearly, the second term is decreasing in k since c(k) is decreasing in k. Also, the

absolute value of the third term increases as k increases for $\sigma > 1$. Therefore, the bound on the right hand side is decreasing in k . This fact, in conjunction with the (reasonable) assumption that $\partial E(\eta)/\partial k \geq 0$ make it likely that the cost per unit time is convex in k .