INTERACTION BETWEEN FREQUENCY OF RESCHEDULING AND THE ROLE OF SAFETY STOCK IN MATERIAL REQUIREMENTS PLANNING SYSTEMS

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

Robert C. Carlson
Department of Industrial Engineering and
Engineering Management
Stanford University

Technical Report 80-18

Revised August 1985
and March 1986

INTERACTION BETWEEN FREQUENCY OF RESCHEDULING AND THE ROLE OF SAFETY STOCK IN MATERIAL REQUIREMENTS PLANNING SYSTEMS

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

and

Robert C. Carlson
Department of Industrial Engineering and
Engineering Management
Stanford University
Stanford, California 94305

Revised August 1985
and March 1986

ABSTRACT

We describe how the frequency of rescheduling permitted in Material Requirements Planning (MRP) systems influences the impact of safety stock on system costs, and its effectiveness in maintaining desired levels of customer service. The results are obtained from approximate analytical models and simulation studies of a single product with stochastic demand and a two-level product structure. The results indicate that in some cases it may be more economical to reschedule infrequently and use safety stock as protection against demand variations. They also indicate that the effect of changing safety stock levels is much more predictable when rescheduling is infrequent, and that increasing safety stock may actually result in degraded performance when rescheduling is frequent.

INTERACTION BETWEEN FREQUENCY OF RESCHEDULING AND THE ROLE OF SAFETY STOCK IN MATERIAL REQUIREMENTS PLANNING SYSTEMS

1. INTRODUCTION

The issues of scheduling policy and safety stock in Material Requirements Planning (MRP) systems have been topics of debate in recent years. Many researchers have proposed approaches for determining scheduling policy or safety stock but little work has been done in examining the interaction between them.

Mather (1977) recommends the use of safety stock in conjunction with firm planned orders but does not address the issue of determining appropriate quantities of safety stock for such circumstances. Bannerjee (1979) reports on a study which indicates that different scheduling policies (lot-sizing and sequencing) result in significantly different system performance as measured by his multi-criteria scores, while most safety stock policies do not produce significantly different results when these criteria are used.

Nahmias and Schmidt (1983) show that the form of the optimal policy for a two-level assembly system with per unit per period holding and shortage costs and no setup costs is similar to a "single critical number" or "order up to" policy. Lambrecht, Muckstadt and Luyten (1984) show that the form of the optimal policy for a similar system with positive setup costs is of the (s,S) type. They also propose a heuristic solution procedure which is shown to be effective and efficient for small problems. The (s,S) policy implicitly handles both the safety stock and rescheduling issues simultaneously. However, the policy explicitly <u>permits</u> rescheduling in each period. The single critical number policy requires an order in each period with positive demand.

In the next section we discuss the problem assumptions. Section 3 describes the simulation model. Section 4 briefly describes circumstances in which component safety stock is unnecessary. In the subsequent four sections, we present results which characterize the effect of finished product and component safety stock under each of two different rescheduling policies. We conclude with a summary in Section 9.

2. ELEMENTS OF THE PROBLEM

We are concerned with the interaction between frequency of rescheduling and the role of safety stock in terms of impact on system performance. Two measures of system performance are used here. The first is service level, measured as the percent of demand filled immediately from stock, often referred to as "fill-rate." The second measure is total cost, which is the sum of production/order setup costs and inventory holding costs.

We study this interaction in a rolling horizon environment. In a rolling production schedule, one implements only the initial decision in each of a series of finite horizon plans. Most MRP systems operate using rolling schedules.

We examine systems in which only demand is uncertain, so that leadtimes, yields, and supply timing and quantity are certain. We also assume that there are no capacity constraints which limit either production or purchase quantities. This examination of the role of safety stock under these conditions is consistent with the findings from a simulation study by Whybark and Williams (1976). Their results indicate that safety stock is a more effective buffer against demand quantity uncertainty than is safety leadtime.

Two different scheduling policies are used. We refer to the first policy as "fixed scheduling" and to the second as "flexible scheduling". Under fixed scheduling, the timing of planned setups is fixed far in advance of the setup. Such a policy allows for the timing of setups to be established with certainty, so that "nervous" schedules are avoided. A "nervous" schedule is one in which the timing of planned setups changes as the schedule rolls forward. In a fixed-scheduling environment, emergency orders and expediting are not allowed, and thus "rescheduling" does not occur. However, the size of each planned production or order quantity is allowed to vary until the production run is begun or the order is placed, at which time the quantity is fixed.

In the flexible-scheduling environment, replanning occurs each period. Actual demand and updated forecasts are reflected in net requirements, and both the timing of setups and the quantity produced in each batch may change. Therefore, emergency setups (setups which occur earlier than previously planned) may occur. We assume, however, that emergency setups are never scheduled solely for the purpose of replenishing safety stock. That is, an order is placed only if forecasted demand during the leadtime ahead exceeds the inventory position.

These two policies are extremes. Most currently implemented MRP systems employ policies which lie somewhere between these extremes. However, they tend to be more similar to flexible scheduling than to fixed scheduling since most MRP schedules are updated on a regular basis, and production runs are scheduled to satisfy any net requirements.

We examine a single product with a two-level product structure in which two components purchased from outside vendors which are assembled into the finished product. This product structure is diagrammed in Figure 1. Although this product structure is simple, it is complex enough to provide insight into the effects of independent demand versus dependent demand and mating of components. However, it does not include multiple end-items, common components, or multiple levels, all of which typically occur in industrial situations. Thus, we do not expect the results to be completely general. The reader is referred to Baker (1985), Baker et al. (1985), and McClelland and Wagner (1985) for some preliminary results on systems with common components.

FIGURE 1

Our ultimate objectives were to determine cost-effective safety stock levels for both the finished product and components and to ascertain whether or not frequent rescheduling is desirable. These results are described briefly in sections 7, 8, and 9, and in more detail elsewhere by Yano and Carlson (1984, 1985, 1986). However, in order to address these optimization-related issues, it was first necessary to understand much more fundamental relationships and interactions; and it is these results that are reported here. In some cases, commonly held beliefs are confirmed; however, in other cases, the results indicate strongly non-monotonic relationships when we might expect monotonicity.

3. SIMULATION MODEL

We first introduce notation used throughout the remainder of the paper.

μ : average demand per period

σ : standard deviation of the demand process

S_i : setup cost for item i

h_i : inventory holding cost charged on end-of-period inventory of item i

L_i : leadtime for item i

T_i : natural cycle of item i

$= \sqrt{2S_i/D_i h_i}$

k_i : safety stock multiplier for item i, where safety stock quantity equals $k_i \sqrt{T_i + L_i}\, \sigma$

We developed a simulation model to study the forementioned relationships and interactions over a range of values of parameters which influence the system, including holding costs, natural cycle length, and demand variability. We normalized h_1 at 1.0, and let h_i = 0.10, 0.25, or 0.40, for i= 2, 3. These alternatives encompass the range of many value-added structures.

Another factor which we vary is the natural cycle. Coined by Baker (1977), it can be viewed as the mean number of periods of demand in a Period Order

Quantity (POQ), or the average time between order/production setups measured in number of periods. We let $T_1 = 2$ or 4, and $T_i = T_1$, $2T_1$, or $3T_1$ for i = 2, 3. The values of L_i were 1 or 5 for all i.

Demand is distributed normally with $\mu = 200$ per period and $\sigma = 10$, 30, or 50. We consider this range of standard deviations sufficiently large to cover many reasonable stationary demand situations. The demand forecast is set equal to mean demand; hence, the standard deviation of demand is equivalent to the standard deviation of the forecast error.

The fixed schedule is achieved by fixing the timing of all orders for a period of time equal to the largest integer multiple of the natural cycle length less than the length of the planning window. This technique limits schedule changes to the end of the horizon, thereby essentially eliminating nervousness in the system. For example, an item with $T_i = 4$ would have its production schedule fixed for 20 periods if the planning horizon length is 24 periods (since 24/4 - 1 = 20). This provides for some flexibility at the end of the horizon which is required in order to avoid scheduling setups whose timing would be different if more demand information were available. The timing of production runs later in the horizon become fixed as the horizon rolls forward. Throughout the study, we use a planning window of 24 periods, which in all cases is at least twice the length of the largest natural cycle.

Throughout the study we use simple ordering policies which could be implemented in practice, yet might be expected to provide good results. In the fixed scheduling scenario we use an (R_i, T_i) policy, where item i orders up to R_i every T_i periods, and $R_i = (T_i + L_i + 1)\mu + k_i \sqrt{T_i + L_i}\,\sigma$. Since the setups are fixed and the demand process is stationary, the (R_i,T_i) policy is optimal. The critical issues with regard to this policy are the optimal safety stock quantities and the effect of fixing the production interval on system cost.

In the flexible scheduling environment, we use the same formula for R_i and optimize k_i for the specified value of T_i. An order is triggered using the standard MRP computations (i.e., an order would be triggered if there is a positive net requirement L_i periods hence). This ordering policy is not necessary optimal. Our reasons for using this policy are the relative ease of optimizing k_i using an algorithm briefly discussed later and the fact that standard MRP software can implement the order triggering.

There is also a behavioral reason for not using a positive inventory value for an order trigger. When there is a positive order trigger, the safety stock tends not to be used as it should be (to handle variations in demand), but instead tends to be treated as "unusable" stock. On the other hand, adding extra units to the order to avoid shortages (in the case of the fixed schedule) or to avoid early setups (in the case of the flexible schedule) is easily accepted without the usual safety stock connotations.

The bicriteria performance measure (total cost and fill-rate) made it difficult to perform statistical analyses in the usual manner. It would have been impossible to calibrate the parameters so as to achieve the same cost (but different fill-rates) or to achieve the same fill-rates (but different costs). Therefore, we chose to use a large number of simulation runs (typically 50, but always at least 25) and used common random numbers (i.e., each set of safety stock parameters faced the same random demand pattern) so as to reduce the variance of the differences among the safety stock and rescheduling policies. The combination of the large number of simulations runs and variance reduction techniques permits us to place a reasonable degree of confidence in the results.

4. CIRCUMSTANCES WHERE COMPONENT SAFETY STOCK IS UNNECESSARY

We note here that component safety stock has no impact on the end-item service level if the component natural cycle is the same as its parent (successor), and the production schedule and quantity of each item in the product structure is fixed for at least its own cumulative leadtime. The cumulative leadtime is the total flow time for that item beginning from the earliest procurement in the product structure until the item is manufactured/assembled. One special case of the above is the situation in which all leadtimes are zero, in which case no component safety stock is necessary.

A simple example of a situation in which the schedule of each item is fixed for its own cumulative leadtime is a two-level serial product structure in which

(1) both levels have $T_i = 2$,

(2) level 2 (component) has $L_i = 3$, and

(3) level 1 (end-item) has $L_i = 1$.

If the production schedule of level 1 is fixed for at least 4 periods and the production schedule of level 2 is fixed for at least 3 periods, level 2 will provide all units requested.

Under the conditions discussed above, an item with the same order/production frequency as its parent always will "know" exactly how many units are required by the next production stage. Hence, safety stock cannot increase service levels in such a situation.

In the following sections we examine the effect of end-item and component safety stock on costs and service levels in various scheduling environments where safety stock will affect service levels as well as costs.

5. END-ITEM SAFETY STOCK UNDER FIXED SCHEDULING

The simulation studies indicate that if fixed scheduling is used on level 1, then increasing end-item safety stock always increases both the fill-rate and costs. Further, there are decreasing marginal returns for safety stock. These results are intuitively evident, so we will not discuss them further.

More important, however, is the fact that the average fill-rate observed in the simulations is a linear function of the (theoretical) service level that would be achieved by a hypothetical single stage system with a cumulative leadtime equal to the actual cumulative leadtime for the assembly system.

We used regression analysis to determine the best fit of observed service levels as a linear function of these "theoretical" service levels. Each data point for the regression analyses is the mean of 50 problems, each with a 24-period horizon. We vary the end-item safety stock multiplier, k, from zero to a value which yields a theoretical service level of 0.98, in steps of 0.2. The only exception to this rule is for a standard deviation = 10. There we use a maximum value of $k_1 = 0.6$, (yielding service levels in excess of .98) in order to establish a set of data large enough for analysis. Component safety stock is set equal to zero.

Regression analyses of these relationships for a number of combinations of natural cycles, and for two different scheduling policies for the components indicated highly linear relationships. In fact, in cases where the relationship had the form $y = a + bx$, the smallest value of R^2 was .995. In cases where the relationship had the for $y = bx$, the slope estimate was significant at the $\alpha =$ 0.00001 level. The high values of R^2 and the high levels of significance of the slope estimates indicate that actual service levels can be predicted accurately from the theoretical service level when either scheduling policy is used for the components. The results are not unexpected, but they suggest that when fixed scheduling is used on level 1, finished product safety stock can be set with predictable results. This is not the case when flexible scheduling is used, as we shall see in the next section.

6. END-ITEM SAFETY STOCK UNDER FLEXIBLE SCHEDULING

When flexible scheduling is used on level 1, one of four events has been observed to occur when end-item safety stock is increased:

(1) costs and fill-rate increase,

(2) costs and fill-rate decrease,

(3) costs decrease and fill-rate increases, or

(4) costs increase and the fill-rate decreases.

The reason for these rather unpredictable results is that safety stock serves two purposes in a flexible scheduling environment. It serves to reduce shortages, as one would expect; it also may serve to prevent an emergency setup which would have occurred otherwise. Practitioners are aware of the latter phenomenon, but many procedures designed to determine safety stock levels do not consider this effect.

We found that (1) occurs when the safety stock behaves as expected; (2) occurs when an emergency setup is prevented (saving costs) but availability over the horizon decreases as a result; (3) occurs when an emergency setup is prevented and availability over the horizon is increased thereby, and (4) occurs when an emergency setup is eliminated but the additional safety stock more than offsets the savings (too much inventory at the wrong time and too little in severe shortage situations). Simulation results which depict these phenomena are shown in Figures 2 and 3. Safety stock quantities for second-level components were set to zero for these simulation studies. Therefore, the impact of emergency setups on both levels--directly for level 1 and indirectly for level 2--is reflected in the costs.

FIGURES 2 and 3

The results here indicate that there are some complex relationships to consider when determining safety stock levels when rescheduling is frequent. One may also interpret the results as suggesting that replanning needs to be done with great care, and that increasing safety stock may not lead to the desired results.

7. SECOND-LEVEL COMPONENT SAFETY STOCK UNDER FIXED SCHEDULING

We simulated MRP systems with two-level product structures, varying the following factors: (1) holding cost rates for components, (2) natural cycle lengths of the end-item and components, (3) leadtimes, (4) variability of demand, and (5) safety stock levels, in order to gain insights into the impact of these factors on system performance measured in terms of total cost and service level.

These simulation studies indicate that increasing component safety stock increases both costs and service levels; but that in most situations, it is not cost-effective, as illustrated in Figure 4. Under special conditions, however, some positive quantities of component safety stock are cost-effective. A cost versus service level curve for such conditions is depicted in Figure 5. Observe that the parameters for this case are extreme, with 80 percent of the value of the product being added in final assembly and a long natural cycle for item 3.

By examining the characteristics of these situations, we determined the elements which are most important in the tradeoff between end-item and component safety stock. They are: (1) holding cost of the component relative to that of the end-item, (2) proportion of the value of the end-item added at the last assembly/manufacturing stage, (3) frequency of setups of the component relative to that of the end-item, and (4) availability of "partner" components with which a particular component must be mated.

FIGURES 4 and 5

Having developed an understanding of these major factors and the interactions among them, we developed an algorithm to determine cost-effective safety stock levels in a two-level product structure (Yano and Carlson, 1984). Computational results indicated that several factors needed to be present simultaneously for component safety stock to be economical. The factors are: (1) the holding cost of the component relative to that of the finished product must be very low (e.g., 1 to 10 ratio), (2) the time between setups of the component cannot be much larger than 2 times that of the finished product assembly, (3) the fill-rate must be very high (i.e., in excess of 99% in most cases), and (4) the "mate" component must have high availability. High availability can be achieved in two ways. The first is a large amount of safety stock, which, in turn, requires the mate to have characteristics (1) and (2) as well. The second way to achieve high availability is a long natural cycle, providing for infrequent stockout occasions. Although circumstances with all four characteristics do occur, they are rare in practice. For detailed descriptions of the algorithm and experimental results, see Yano and Carlson (1984).

8. SECOND-LEVEL COMPONENT SAFETY STOCK UNDER FLEXIBLE SCHEDULING

Simulation studies of a two-level system indicate that when fixed scheduling is used on level 1 and flexible scheduling is used on level 2, the primary effect of additional component safety stock is to change total costs, while the impact on service level in most cases is insignificant. The reason is that safety stock serves to avert emergency setups which would have occurred in the absence of safety stock. In so doing, the fixed per cycle costs are spread over a longer time period, and the cost per unit time declines. This savings must be balanced with the cost of holding the safety stock. Figure 6 illustrates how total costs vary as a function of safety stock for a typical case. We developed an algorithm (Carlson and Yano, 1986) which determines approximately optimal safety stock levels under these conditions. For our test cases, we found that it was desirable to hold a moderate to high level of safety stock (k = 1.0 or greater) to obtain the best performance.

FIGURE 6

Evidence from our simulation studies indicates that the effect of component safety stock in a system in which flexible scheduling is used on both levels is highly unpredictable, and we therefore do not report the details here. The primary reason for this is that there may be emergency setups for the parent item which necessitate emergency setups of the component as well. This type of emergency setup for a component cannot be averted by normal quantities of component safety stock. However, the emergency setups which result from attempts to avoid anticipated component shortages may be averted by component safety stock. The savings from reducing the number of emergency setups of the second type must be balanced with the cost of safety stock.

The tradeoffs discussed in this section are relevant only for components whose natural cycles are greater than that of the end-item. When the natural cycle of a component is equal to that of its parent, there is no benefit to scheduling an emergency setup. All of its setups are "synchronized" with those of its parent (successor), and any orders placed earlier than planned will increase setup costs and holding costs, with no impact on customer service levels.

9. DISCUSSION AND CONCLUSIONS

Having developed one algorithm to determine cost effective amount of safety stock for each component rescheduling policy, we were able to evaluate the desirability of frequent component rescheduling. We selected a set of parameters which would give "flexibility" considerable advantage, determined appropriate safety stock levels for each scenario, and then compared the costs using simulation. The results indicated that even when the parameters would make flexibility advantageous, the fixed scheduling policy was more economical (see Yano and Carlson (1985) for details). The situation which we examined had a stationary demand process with moderate forecast errors. Futher research is required to determine whether similar results would be obtained in a situation with an erratic or non-stationary demand pattern. Additional research is also needed to develop <u>simple</u> procedures to obtain ordering policies for semi-flexible scheduling policies in which the rescheduling decision is "optimized," as well as methods to incorporate such policies in existing MRP systems. Nevertheless, the results suggest that frequent rescheduling should be done with caution.

Currently, many MRP systems operate with schedules being replanned periodically, typically once a week. In such situations, comparable to our flexible scheduling policy, the implications of the findings here are significant. First, increasing end-item safety stock may not provide the anticipated increase in service level, and may actually decrease service level in the short run. Second components safety stock will not lead to significant increases in the service level, but may reduce costs. Finally, it may not be economical to replan too frequently. The coordination provided by fixing the timing of production runs may yield excellent, predictable customer service levels quite economically with appropriate levels of safety stock.

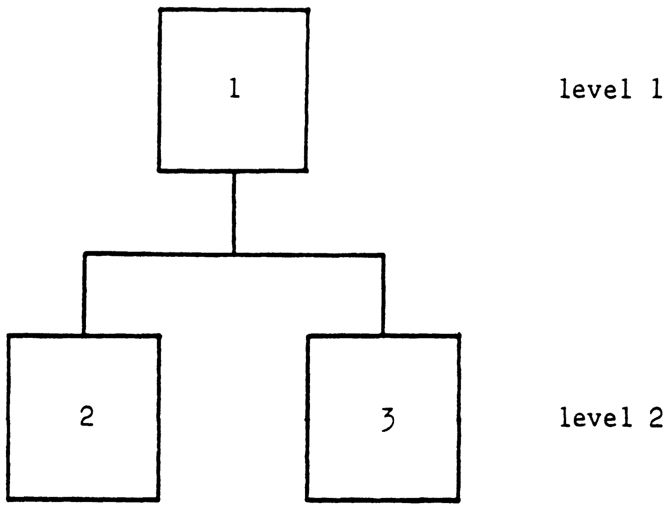

Figure 1

Two-Level Product Structure With Component Numbers

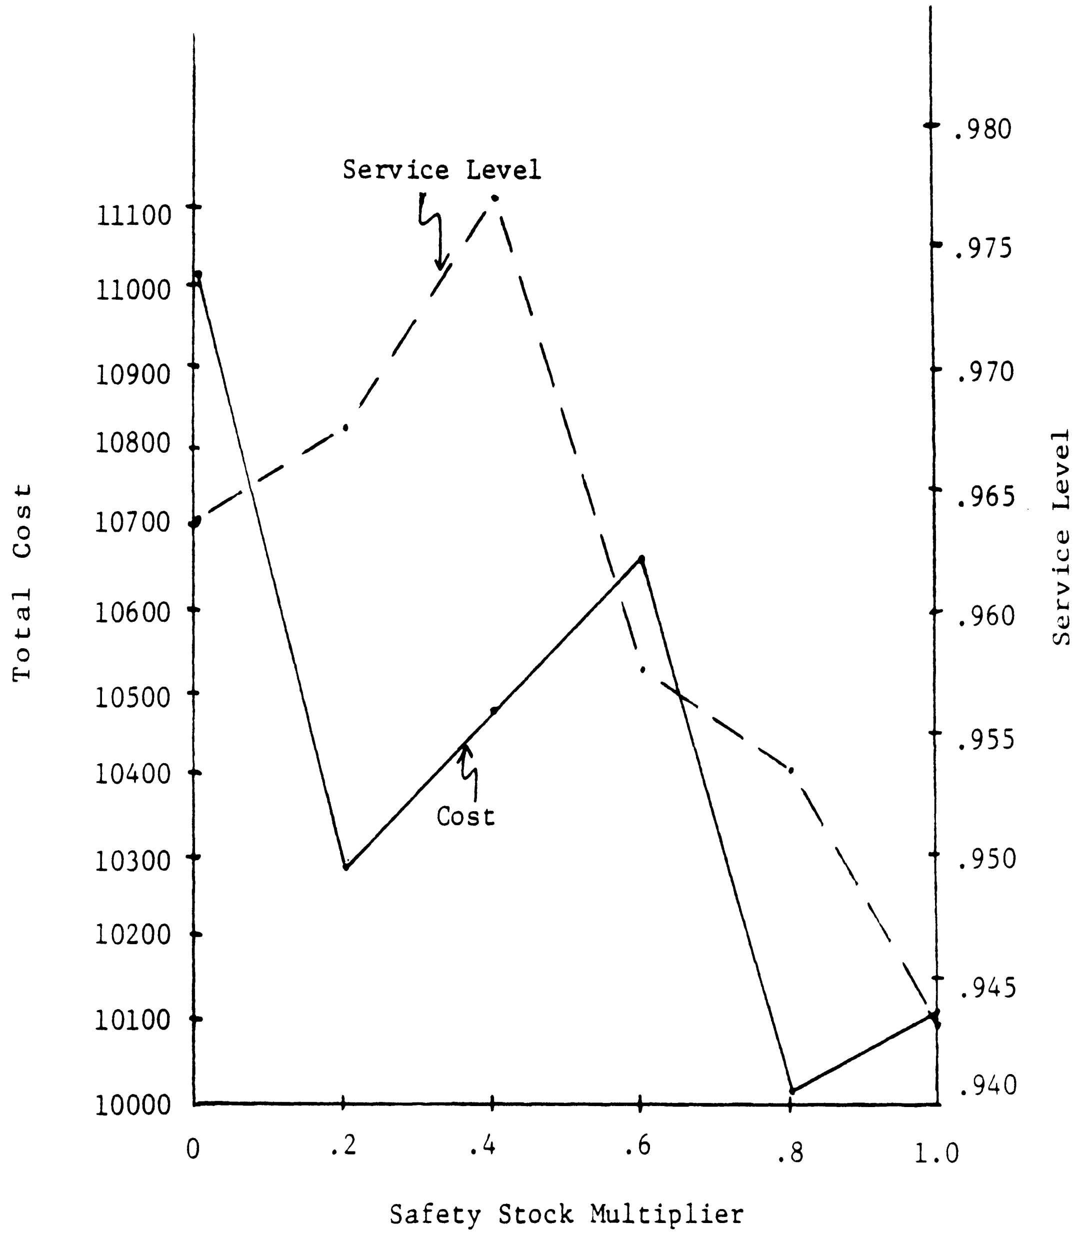

Figure 2

Total Cost and Service Level as a Function
of Safety Stock Quantity with
Flexible Scheduling on Both Levels
$\underset{\sim}{T} = (2,2,2)$, $\underset{\sim}{L} = (1,1,1)$, $\sigma=30$

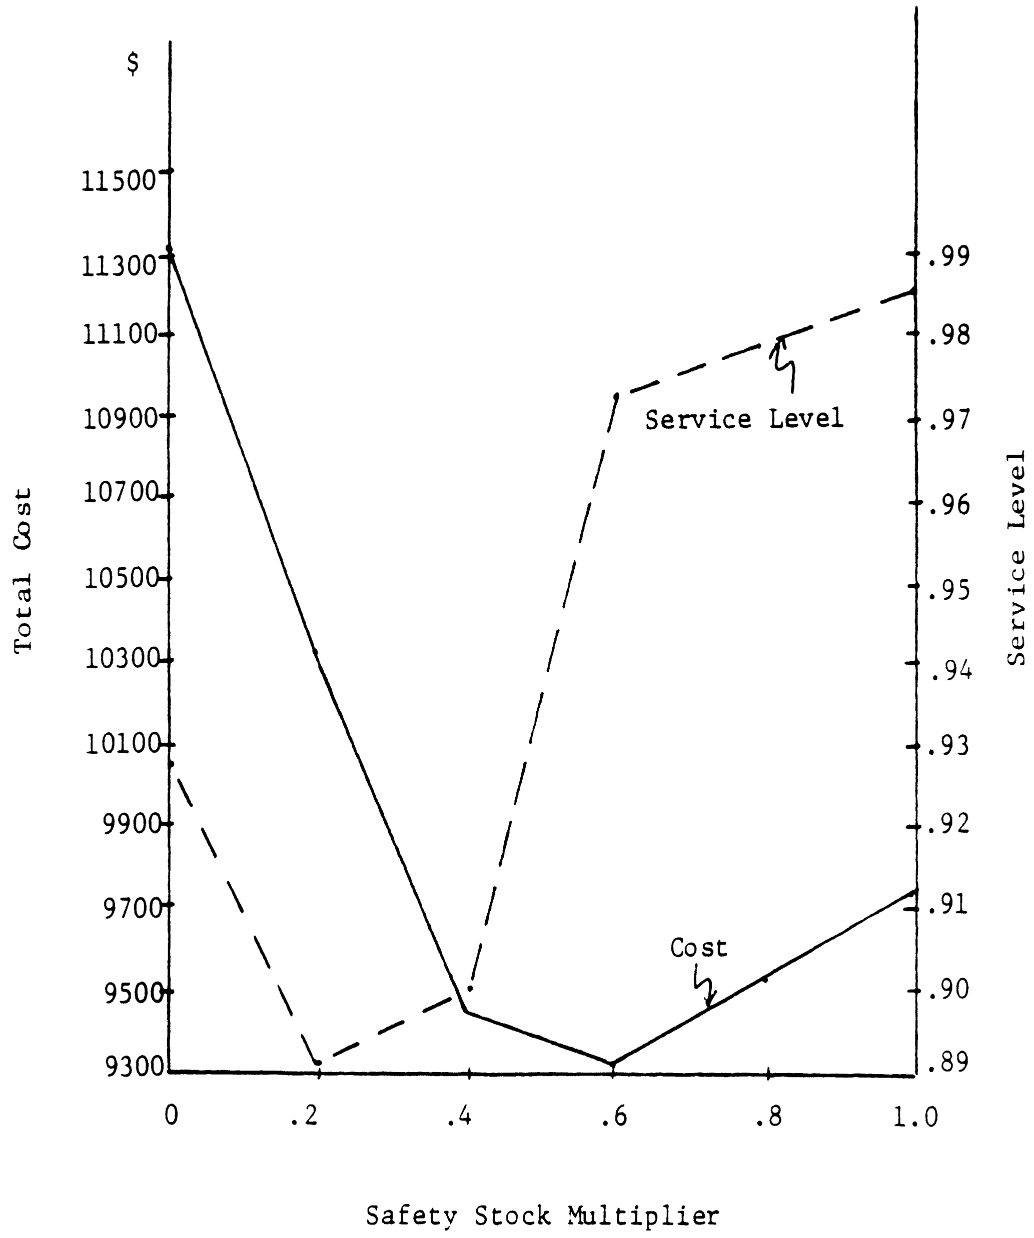

Figure 3

Total Cost and Service Level as a Function
of Safety Stock Quantity with
Flexible Scheduling on Both Levels
T = (2,2,2), L = (1,5,5), σ=30

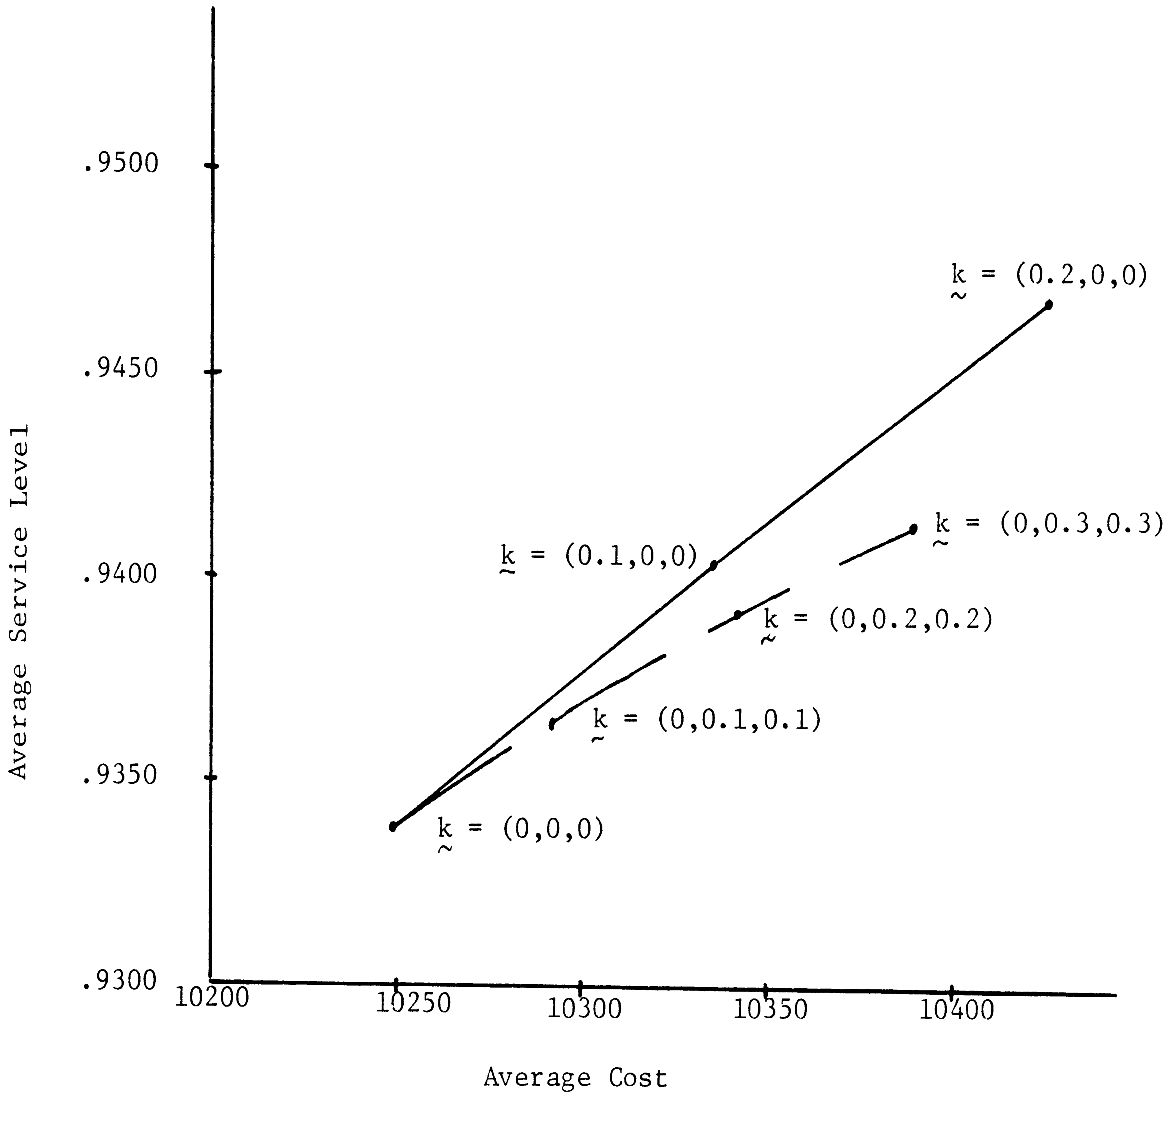

Figure 4

Results for $\underset{\sim}{T} = (2,4,4)$, $\underset{\sim}{h} = (1,0.1,0.1)$, $\underset{\sim}{L} = (1,1,1)$, $\sigma = 10$

Under Fixed Scheduling

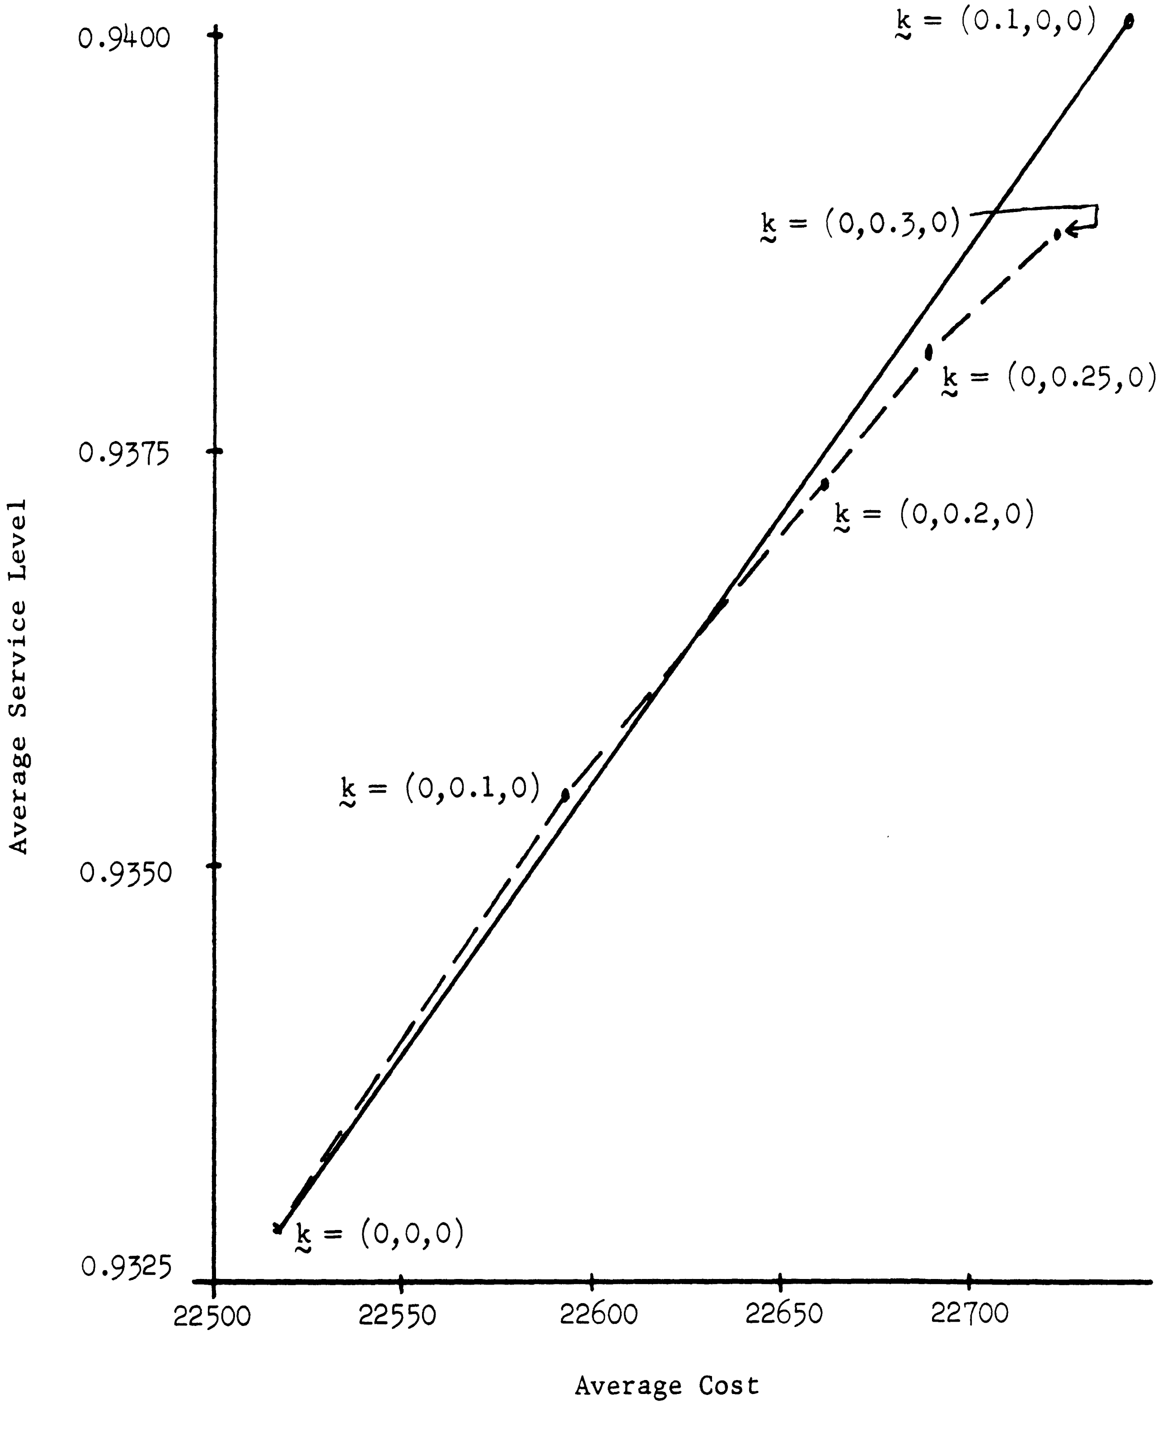

Figure 5

Results for $\underset{\sim}{T} = (4,4,12)$, $\underset{\sim}{h} = (1,0.1,0.1)$, $\underset{\sim}{L} = (1,1,1)$
Under Fixed Scheduling

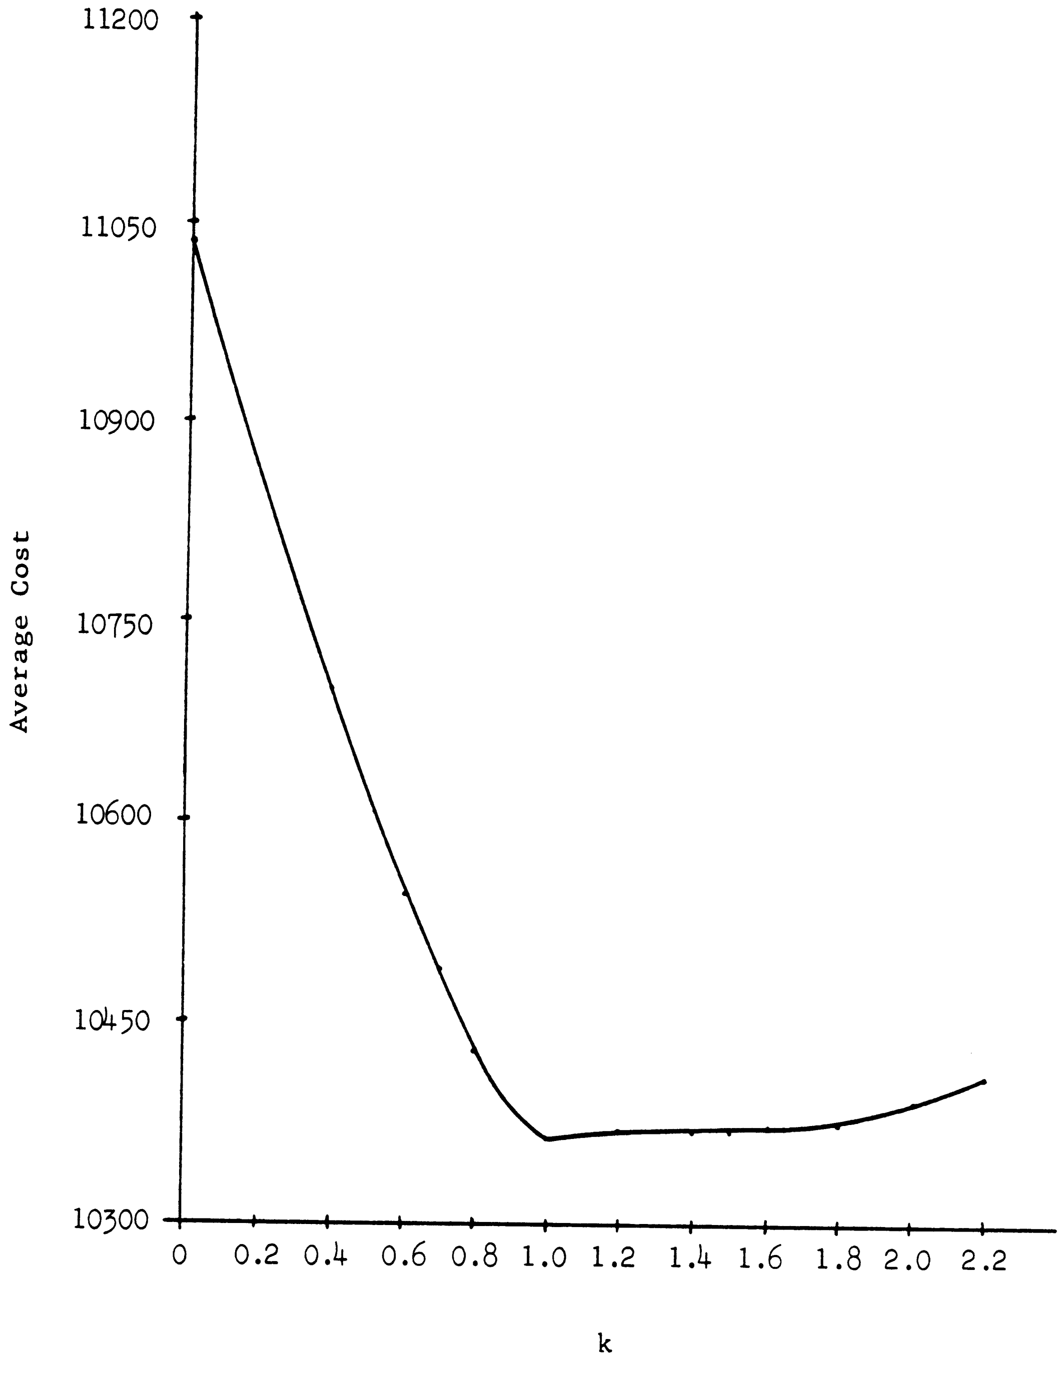

Figure 6

Results for $\underset{\sim}{T} = (2,4,4)$, $\underset{\sim}{h} = (1,0.1,0.1)$, $\underset{\sim}{L} = (1,1,1)$, $\sigma = 10$
With Fixed Scheduling on Level 1
and Flexible Scheduling on Level 2

REFERENCES

Baker, K. R., "An Experimental Study of the Effectiveness of Rolling Schedules in Production Planning." Decision Sciences, Vol. 8 (1977), pp. 17-27.

Baker, K.R., "Safety Stocks and Component Commonality." Working Paper No. 159, Amos Tuck School of Business Administration, Dartmouth College, October 1984.

Baker, K.R., M.J. Magazine, H.L.W. Nuttle, "The Effect of Commonality on Safety Stock in a Simple Inventory Model." Working Paper No. 165, Amos Tuck School of Business Administration, Dartmouth College, January 1985.

Bannerjee, A., "Relative Importance of Scheduling and Safety Stocks in Material Requirements Planning Systems." AIDS Proceedings, New Orleans (1979), 251.

Carlson, R. C. and C. A. Yano, "Safety Stocks in MRP Systems with Emergency Setups for Components." Management Science, Volume 32 (1986).

Lambrecht, M. R., J. A. Muckstadt, and R. Luyten, "Protective Stocks in Multi-Stage Production Systems." International Journal of Production Research, Vol. 22 (1984), pp. 1001-1027.

Mather, H., "Reschedule the Reschedules You Just Rescheduled - A Way of Life for MRP?" Production and Inventory Management, Vol. 18 (1977), pp. 69-79.

McClelland, M. and H.M. Wagner, "The Impact of the Product Structure and Increased Part Commonality on Safety Stock Effectiveness." Working Paper, School of Business Administration, North Carolina State University, 1985.

Schmidt, C. and S. Nahmias, "Optimal Policy for a Two Stage Assembly System Under Random Demand." Operations Research, Vol. 33 (1985), pp. 1130-1145.

Whybark, D. C. and J. G. Williams, "Material Requirements Planning Under Uncertainty." Decision Sciences, Vol. 7 (1976), pp. 595-606.

Yano, C.A. and R. C. Carlson, "An Analysis of Scheduling Policies in Multi-echelon Production Systems." IIE Transactions, Vol. 17 (1985), pp. 370-377.

Yano, C. A. and R. C. Carlson, "Buffering Against Demand Uncertainty in Material Requirements Planning Systems - System with No Emergency Setups." Technical Report 84-13, Department of Industrial and Operations Engineering, University of Michigan, 1984.