


BASEMENT



HD28
.M414

no. 1296-82



WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

A STUDY OF PRODUCTION SMOOTHING
IN A JOB SHOP ENVIRONMENT

Allan B. Cruickshanks
Robert D. Drescher
Stephen C. Graves

Alfred P. Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA 02139

January 1982

WP1296-82

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139

A STUDY OF PRODUCTION SMOOTHING
IN A JOB SHOP ENVIRONMENT

Allan B. Cruickshanks
Robert D. Drescher
Stephen C. Graves

Alfred P. Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA 02139

January 1982

WP1296-82

M.I.T. LIBRARIES
MAY 5 1982
RECEIVED

ABSTRACT

We consider the problem of smoothing production in a job shop in which all production is to customer order. We present an approach to production smoothing based on the concept of a planning window. A planning window is the difference between the promised delivery time and the planned production time for a product. It represents the degree of flexibility available for planning the production of committed orders. We characterize the production smoothing benefits for a range of planning windows by means of an approximate analytic model and a simulation study. These analyses show that substantial smoothing benefits result from small changes in the length of the planning window. We discuss the implementation of the production smoothing approach and illustrate this implementation with an industrial case study that was the motivation for this work.

0743903

1. Introduction

Production smoothing and planning deal with the setting of production, inventory, and work force levels to satisfy demand requirements over a medium-term planning horizon. The need for production planning results primarily from aggregate demand fluctuations, such as those that occur for seasonal product families. There has been an extensive management science effort developing decision models to support this planning process. Both Silver [5] and Hax [3] provide excellent surveys and critiques of this body of work. A typical scenario addressed by management scientists has been one in which the following strategies exist for meeting demand requirements:

- (i) varying the aggregate production level to meet anticipated fluctuations in demand while keeping inventory levels constant. The production level is varied by changing the work force level and/or by using overtime.
- (ii) varying the inventory levels to handle anticipated fluctuations in demand while keeping the aggregate production level constant and equal to the average aggregate demand rate.
- (iii) a mixture of strategies (i) and (ii).

By associating costs to the consequences of each strategy and by identifying constraints on the production planning decisions (such as satisfying forecasted demand over the planning horizon), the management scientist can then formulate a mathematical model with decision variables being the production, inventory, and work force levels, and with a criterion of production, inventory, and work force costs. The model's solution would determine the production planning decisions that minimize total costs.

In this paper we consider a new type of production smoothing problem. We consider a production operation for which it is either not possible or not

practical to stock uncommitted finished goods inventory. All production is for contracted orders. Job shops which only produce custom-made products operate in this manner by definition. The final assembly of complex products such as office or laboratory equipment, computer systems, and automobiles, is often operated without a finished goods inventory due to the multitude of options available on the final products. In addition, some manufacturers are able to maintain a policy of production only to order, based on the market conditions for their products.

Production planning in such operations, which we generically call job shop operations, clearly cannot use inventory to smooth demand fluctuations. Rather production planning must rely on other strategies. One option is strategy (i): vary production to match the demand rate. A second option corresponds roughly to strategy (ii): maintain a constant production rate while varying the delivery time of orders depending on the shop load. That is, the quoted delivery time for an order is variable and depends on the shop load, with the shop operating at a constant production level. This option is very reasonable provided that the firm's customers will tolerate the variable delivery time. However, if the customer market expects a certain delivery lead time, then this option results either in backordered demand (late delivery) or in lost sales whenever the shop load dictates an unacceptable delivery time. In this paper we explore a third production smoothing option for situations where the market requires a firm delivery lead time and does not tolerate late delivery. Here production smoothing is possible only if there is some space between the planned time to produce a product and the promised delivery time. We term this difference between the delivery lead time and the production lead time to be the planning window. Production smoothing can occur over this planning window with the extent of the production smoothing depending on the size of the

planning window. We examine this production smoothing strategy for a single stage production operation with non-seasonal demand but with significant demand fluctuations caused by the inherent randomness of product orders.

The remainder of the paper is in four sections. The next section presents a case study that was the motivation for this work. We developed a production smoothing model for a local manufacturer that operates as a job shop with production to order. In Section 3 we present an analytic analysis of an approximate model to the smoothing approach used in the case study. This analytic analysis provides a preliminary characterization of the benefits possible from production smoothing. In Section 4 we use a simulation study to compare the approximate model with the smoothing model used in the case study, and to provide a more definitive characterization of the smoothing benefits. The final section discusses the implementation of the production smoothing model both at the local manufacturer and in general.

2. Case Study

The production smoothing model originated from a case study at a large manufacturing firm. This firm produces a variety of electromechanical instruments used in extremely high value applications. The firm operates in a business environment characterized by long lead times, and with a small number of powerful buyers who view the product as a minor but important component in the assembly of their final products. Demands for the firm's products are unpredictable and fluctuate dramatically, depending on the sales of the customers' final products. The restrictions of no backorders and no uncommitted inventory are very much a reality for this firm. The firm's management feels that the irregular and uncertain nature of the marketplace, along with their products' high value and customer-specific nature, dictate a policy of production only to firm orders. As such, holding inventory for the purpose of smoothing production is not a viable option. Backordering of demand is also not an option for production smoothing since the firm's customers cannot tolerate a late delivery due to their own tight production schedules.

We examined the production planning function in one department within the firm. This department is typical of other departments, and manufactures approximately fifty products, most of which perform a similar general function. These fifty products are grouped into twelve families based on similarities in design and on the sharing of production equipment and labor. Since each item in a family shares the same production resources, aggregate production planning takes place at the family level.

The manufacturing process is primarily an assembly process. Orders are received from customers for standard products that are modified to their specifications. Long production lead times for the firm are mainly a result of long lead times for purchased component parts and subassemblies. The

lengthy production lead times are acceptable to the customers, given on-time delivery, since their final production has a much longer production cycle. However, the market is competitive, with delivery lead times as one marketing factor along with product quality and the certainty of delivery within the promised lead time.

Production planning takes place on a monthly basis. Traditional management practice is to promise delivery based on the production lead time. Thus production of a particular order is planned to start within the month of receipt of the order. This practice does not allow the management much flexibility, and can create a very irregular production schedule. A family's monthly production schedule often fluctuates by 50% or more in the volume of orders. Figure 1 gives an example of the demand history, and subsequent production schedule, for a typical product family.

Management is concerned about irregular production schedules and the resulting costs incurred in changing production level from month to month. Shifts in monthly production level cannot be accommodated by hiring and firing workers due to union agreements and the effect on worker morale. Since the work force is effectively fixed, peaks in the production schedule are met by using overtime whereas periods of low demand mean slack time for the workers.

Thus the firm has been interested in new approaches to production planning that would reduce these costs. However, any proposed solution has to recognize the restrictions on the implementation of any new policy in the firm's environment. Production planning is done according to very specific and simplistic rules by union employees who hold their positions by virtue of seniority. Any new method cannot involve a complicated procedure, but has to be extremely simple to understand and use. In addition, there are data processing limitations. Neither hardware nor systems



Figure 1: Monthly Demand for Product Family "X"

personnel are available for any sophisticated manufacturing applications. This also necessitates a straightforward approach.

Since the firm will neither hold uncommitted inventory nor backorder demand, the only immediate option for production smoothing seems to be to create a planning window by increasing the delivery lead time quoted to customers or by decreasing the production lead time or both. The difference between the delivery lead time and the production lead time is a planning window over which a production planner can smooth production. For $L_d(L_p)$ being the delivery (production) lead time, we say that there is a planning window of length $L_d - L_p + 1$. In any month τ the production planner sets the production level for completion in month $\tau + L_p$ based on knowledge of all unfilled contracted demands for delivery in months $\tau + L_p, \dots, \tau + L_d$. This production level must satisfy all unfilled demands for delivery in month $\tau + L_p$ to ensure no backorders. Otherwise, the production planner may set this production level to smooth production as much as possible over the planning window that ranges from month $\tau + L_p$ to month $\tau + L_d$. This may result in the early production of some orders and consequently an inventory of finished goods; however, this inventory is not speculative but is committed to firm orders.

Within this planning framework, which is consistent with the firm's planning environment, there are two design questions: (i) what is the best length for the planning window?; and (ii) how should production be smoothed over the planning window? In order to address the first question we need an answer to the second question. We propose the simple smoothing model S1 as given by

$$S1 \quad P_t = \max_{j=1,2,\dots,n} \left[\frac{\sum_{k=0}^{j-1} D_{t+k} - I_{t-1}}{j} \right] \quad (1)$$

where P_t is production for completion in month t , D_t is contracted demand for

delivery in month t , I_t is planned ending inventory in month t , and $n = L_d - L_p + 1$ is the window length. This procedure first determines the average net production requirements over the first j periods in the planning window, for $j=1,2,\dots,n$. The maximum of these average net production requirements is then taken to be P_t , the production for completion in month t . Suppose that this maximum occurs for $j = j^*$; then by producing at the constant rate P_t for the first j^* periods in the planning window, we have no backorders over these periods and the ending inventory after j^* periods is zero. Furthermore, j^* is the largest value of the index j ($1 \leq j \leq n$) such that this is true. This procedure S1 attempts to smooth the demand over the planning window as much as possible subject to no backorders. We believe this smoothing procedure to be a reasonable and implementable method chosen from among a set of candidate procedures.

The smoothing model S1 bears some resemblance to the single-machine model studied by Baker and Bertrand [1]. They use an allowance factor, which is analogous to the planning window in S1, for setting the due dates of jobs arriving to the single-machine system. They assume that the production level is fixed and measure the resulting job tardiness from their allowance factor and their scheduling rules. In contrast, model S1 constrains job tardiness (i.e. backorders) to be zero, while allowing the production level to vary. We will see that our results that relate the variability of production levels to the length of the planning window, are very similar in form to the results of Baker and Bertrand [1] that relate job tardiness to the size of the allowance factor.

To examine the quality of smoothing procedure S1 as well as the benefits from longer planning windows, we simulate the use of the smoothing method on several product families using actual demand histories. Figure 2 depicts the type of results obtained for a typical product family. This




Figure 2: Trade-off Between Smoothness and Inventory for Product Family "X"

figure shows how the smoothness of the production plan and the average inventory level vary with the length of the planning window. In this figure we measure production smoothness by the average absolute change in production level, (i.e. $E(|P_t - P_{t+1}|)$) where $E(\)$ denotes expectation). The average monthly change in the production level decreases dramatically for a small increase in the window length. As the window length gets larger, however, the rate of change decreases. Conversely, the inventory level increases rapidly for the initial increases in window length. As the window length gets larger, however, the marginal increase in the inventory level decreases. Similar results occur when we measure production smoothness by the variance of the production level and by the expected value of the square of the change in production level ($E[(P_t - P_{t+1})^2]$).

The results of the simulations were very encouraging in that they indicate that substantial production smoothing could be obtained with smoothing model S1 and with a fairly short planning window. However, any increase in the size of the planning window results in increased inventory. To understand better the behavior of the smoothing model S1 as well as the tradeoff between smoothness and inventory, we analyze the smoothing method in the next section.

3. An Approximate Analytic Model

The case study provides some evidence of the potential benefits that the smoothing model S1 can provide. The next two sections attempt to provide some insight into the behavior of this production smoothing model. This section presents an analytic description of the behavior of a model that is closely related to the production smoothing model S1. The following section uses a Monte Carlo simulation to gain further understanding of the performance of the production smoothing model S1.

We have been unsuccessful at deriving any general analytic results for the production smoothing model S1. The difficulty in analysis seems to stem from the requirement of no backorders, which is enforced by the maximization operation in (1). However, by relaxing this requirement, we obtain a model that is quite tractable. Consider the following model S2 as an approximation to S1:

$$\begin{array}{l}
 \text{S2} \quad P_t^* = \frac{\sum_{k=0}^{n-1} D_{t+k} - I_{t-1}^*}{n} \quad (2)
 \end{array}$$

Here the production level for completion in month t is just the average demand, net of planned inventory (I_t^*), over the planning window of length n. We note, though, that this model permits backorders. We hypothesize that the behavior of the model S2 is indicative of that for model S1. The back-order restriction, as implemented in (1), just results in production smoothing over a shorter planning window. For instance, if the maximization in (1) occurs at j^* ($j^* \leq n$), then in effect the smoothing model S1 acts as if the planning window is for j^* periods rather than n periods, and consequently sets the production level equal to the average demand level, net of inventory, over the planning window of length j^* .

From (2) we can express the change in the product level from period to period for the approximate model as

$$P_t^* - P_{t-1}^* = \frac{D_{t+n-1} - D_{t-1} + I_{t-2}^* - I_{t-1}^*}{n} \quad (3)$$

By substituting into (3) the inventory balance equation

$$I_t^* = I_{t-1}^* + P_t^* - D_t, \quad (4)$$

we obtain after rearrangement

$$P_t^* = (1/n)D_{t+n-1} + (1 - 1/n)P_{t-1}^*. \quad (5)$$

But this equation is of the form of an exponential smoothing model with a smoothing parameter of $1/n$. The equation states that the production level for completion in period t is a convex combination of the production level set in the previous month with the demand orders received in the most recent period for delivery in period $t+n-1$; the weights for the convex combination depend upon the length of the planning window.

The smoothing model, as stated in (5), resembles the control numbers approach to production planning (Magee and Boodman [4], pp. 199-207). The control numbers approach sets the current period's production level to be the last period's production level plus an adjustment factor. The adjustment factor is a fraction (prespecified as the control number) of the deviation between planned and actual cumulative production. We can interpret (5) as a control numbers approach where the control number is $1/n$ and where the deviation between planned and actual cumulative production is given by $(D_{t+n-1} - P_{t-1}^*)$.

By recursive substitution in (5), we obtain the following expression for the production level

$$P_t^* = (1/n) \sum_{k=0}^{\infty} \left(\frac{n-1}{n}\right)^k D_{t+n-1-k} \quad (6)$$

where we have assumed that an infinite demand history exists, i.e. $\{D_t\}$

for $-\infty < \tau \leq t+n-1$. If the demands $\{D_\tau\}$ are independent and identically distributed (i.i.d.) random variables, each with mean D and variance σ_D^2 , then we obtain from (6)

$$E(P_t^*) = D \quad (7a)$$

$$\text{Var}(P_t^*) = \sigma_D^2 / (2n-1). \quad (7b)$$

Normalizing these results with respect to the average demand and variance of the demand gives:

$$\bar{P}^* = E(P_t^*) / D = 1 \quad (8a)$$

$$\sigma_{P^*}^2 = \text{Var}(P_t^*) / \sigma_D^2 = 1 / (2n-1) \quad (8b)$$

This result shows that, for the approximate model S2, the magnitude of the smoothing effect as measured by the normalized production variance depends only on the length of the planning window. For example increasing the window length from 1 period to 2 periods reduces the production variance from equaling the demand variance to 33% of the demand variance, regardless of the parameters or distribution of the demand process.

In addition to production variance, a second measure of production smoothness is the period-by-period change in the production level. From (6) we find that

$$\Delta P_t^* = P_{t+1}^* - P_t^* = (1/n)D_{t+n} - (1/n)^2 \sum_{k=0}^{\infty} \left(\frac{n-1}{n}\right)^k D_{t+n-1-k}. \quad (9)$$

Then, if the demands are distributed as i.i.d. random variables, we have

$$E(\Delta P_t^*) = 0 \quad (10a)$$

$$\text{Var}(\Delta P_t^*) = 2\sigma_D^2 / n(2n-1). \quad (10b)$$

Again, we see that the normalized variance of the production level change only depends on the length of the planning window:

$$\sigma_{\Delta P}^2 * = \text{Var}(\Delta P_t^*) / h_D^2 = 2/n(2n-1) \quad (11)$$

Figure 3 plots both the normalized production variance and normalized production change variance as a function of the window length. In both instances there is a dramatic reduction in the variance from increasing the window length from 1 to 2 periods, yet little benefit from increasing the window length beyond 3 periods. Increasing the window length from 1 to 2 periods reduces the normalized production variance from 1.00 to .33, and reduces the normalized production change variance from 2.00 to .33. At a window length of three periods, the normalized production variance is .20 whereas the normalized production change variance is .13. The nature of this variance reduction for the approximate model is similar to that observed in the case study.

From (2) and (6), we can express the inventory at the end of period t in terms of the contracted demands:

$$\begin{aligned} I_t^* &= \sum_{k=1}^n D_{t+k} - \sum_{k=0}^{\infty} \left(\frac{n-1}{n}\right)^k D_{t+n-k} \\ &= \sum_{k=0}^{n-1} \left[1 - \left(\frac{n-1}{n}\right)^k\right] D_{t+n-k} - \sum_{k=n}^{\infty} \left(\frac{n-1}{n}\right)^k D_{t+n-k} \end{aligned} \quad (12)$$

If the demands are i.i.d., the mean and variance of the inventory level are

$$E(I_t^*) = 0 \quad (13a)$$

$$\text{Var}(I_t^*) = \sigma_D^2 \left[2n \left(\frac{n-1}{n}\right)^n - \frac{n(n-1)}{2n-1} \right]. \quad (13b)$$

The normalized variance of the inventory level again depends solely on the window length:




Figure 3: Normalized Production Variance and Normalized Production Change Variance for Smoothing Model S2

$$\sigma_{I^*}^2 = \frac{\text{Var}(I_t^*)}{\sigma_D^2} = 2n\left(\frac{n-1}{n}\right)^n - \frac{n(n-1)}{2n-1} \quad (14)$$

The inventory level fluctuates around zero with negative inventory being backordered demand. Consequently we do not expect any correspondence between the inventory behavior of the approximate model S2 and that of the model S1, since the model S1 does not permit backorders. We note, though, that if uncommitted finished goods inventory could be held, then the approximate model S2 could be used in practice by holding a safety stock of inventory. For instance, if demand is normally distributed, by holding a safety stock of two standard deviations of I_t^* [as determined from (13b)], we have a service level of 98%. Hence, σ_{I^*} as determined from (14) is indicative of the amount of safety stock required to make the approximate model S2 operational, provided that uncommitted inventory can be stocked. Figure 4 shows the relationship between σ_{I^*} and window length.

The analysis given for the approximate model S2 is quite simple. However, our ultimate goal is to understand the behavior of the smoothing model S1. Earlier we provided some limited justification for the analysis of the approximate model S2 as a surrogate for the analysis of S1. We now present some additional motivation for the comparability of the two models. One mechanism for studying the behavior of a smoothing model is to characterize the model's response to an extreme input stream. One such extreme demand (input) stream is a pulse function given by

$$D_t = D \quad t \neq 0$$

$$D_0 = D + \delta$$

where δ is a positive or negative pulse. We consider the behavior of models S1 and S1 for two cases, depending upon the sign of the pulse.




Figure 4: Normalized Standard Deviation for Smoothing Model S2

Case 1: When δ is negative, the smoothing model S1 with window length n gives the following response:

$$\begin{aligned} P_t &= D && \text{for } t < 0 \\ P_t &= D + \frac{\delta}{n} \left(\frac{n-1}{n}\right)^t && \text{for } t \geq 0 \end{aligned} \tag{15}$$

The approximate model S2 has a nearly identical response:

$$\begin{aligned} P_t^* &= D && \text{for } t \leq -n \\ P_t^* &= D + \frac{\delta}{n} \left(\frac{n-1}{n}\right)^{t+1-n} && \text{for } t > -n \end{aligned} \tag{16}$$

The responses given by (15) and by (16) are exactly the same except that one is offset by $n-1$ periods from the other.

Case 2: When δ is positive, the response of S1 is given by:

$$\begin{aligned} P_t &= D && \text{for } t \leq -n, \\ P_t &= D + \frac{\delta}{n} && \text{for } -n < t \leq 0, \\ P_t &= D && \text{for } t > 0. \end{aligned} \tag{17}$$

The response of S2 to a positive pulse is the same as for a negative pulse:

$$\begin{aligned} P_t^* &= D && \text{for } t \leq -n \\ P_t^* &= D + \frac{\delta}{n} \left(\frac{n-1}{n}\right)^{t+1-n} && \text{for } t > -n \end{aligned} \tag{18}$$

These responses (17) and (18) seem quite different. Whereas the smoothing model S1 averages the pulse over n periods, the model S2 smooths the pulse at a geometric decay rate. However, the two responses are very similar with regard to two key statistical measures. For any time series $\{X_t\}$, define $\text{Var}(X_t, a, b)$ to be the variance of the time series over the time interval $[a, b]$. Then for

$\{P_t\}$ and $\{P_t^*\}$ given by (17) and (18) respectively, we can show that

$$\lim_{s \rightarrow \infty} \frac{\text{Var}(P_t, -n, s)}{\text{Var}(P_t^*, -n, s)} = \frac{2n-1}{n} \quad (19)$$

and

$$\lim_{s \rightarrow \infty} \frac{\text{Var}(P_t - P_{t+1}, -n, s)}{\text{Var}(P_t^* - P_{t+1}^*, -n, s)} = \frac{2n-1}{n} \quad (20)$$

That is, both the ratio of production variances and the ratio of the production change variances for the two model responses have the same limit.

Furthermore, the limit of these ratios depends only on n , the length of the planning window. Hence, although the two models smooth the positive pulse differently, the smoothness of the two responses, as measured by the production variance and by the production change variance, are quite similar and differ only by a scale factor that depends on the window length.

The examination of the two models' responses to a positive and a negative pulse provide some evidence of how the approximate model S2 relates to the model S1. This analysis indicates that the smoothing behavior of model S1 may differ from that for S2 only by a scale factor that depends only on the window length. In the next section we empirically explore this correspondence via a Monte Carlo simulation of model S1.

4. Simulation Study

The analytic derivation in the previous section characterizes the production smoothing model S2 for which backorders are allowed. The question remains as to the behavior of the production smoothing model S1 in which no backorders are allowed. Furthermore, we desire to see what correspondence, if any, exists between the two smoothing models. To address these questions and to better understand the observed benefits of smoothing, we developed a simulation model to study model S1.

We simulated the smoothing model S1 with six independently-generated demand streams. Each demand stream consists of i.i.d. normally-distributed demands with a mean demand for each stream of 200. The demand streams differ only in their standard deviations, with the standard deviation taking one of six values: $\sigma_D = 5, 10, 20, 30, 40, 50$. For each demand stream we simulate over 10,000 periods the smoothing model S1 with window length n for each $n=1,2,\dots,9$. The simulation results are given in Figures 5, 6, and 7, and in Table 1. Results presented without all six cases (i.e. $\sigma_D = 5,10,20, 30,40,50$) being plotted are considered to summarize the entire set of available data. In such cases the data not shown would either be superimposed on other data or could be easily predicted, despite its absence, from the observed trends.

There are two striking observations from the simulation results. First, the smoothing behavior of S1 for normal i.i.d. demand depends on the parameters of the demand distribution only as scale factors. Consequently, for any normal i.i.d. demand stream we would use the factors given in Table 1 to predict the production smoothness, as given by the production variance and production change variance, and the inventory consequences of model S1. Second, the smoothing behavior of S1, as given by the production variance and by the production change variance, is remarkably similar to that of the approximate model S2. The production change variances are virtually identical for the




Figure 5: Smoothing Models S1 and S2 - Normalized Production Variance




Figure 6: Comparison Between Smoothing Models S1 and S2 - Normalized Production Change Variance




Figure 7: Normalized Average Inventory Level for Smoothing Model ST

NORMALIZED AVERAGE INVENTORY

Table 1: Summary of Normalized Results for Models S1 and S2 for Normally-Distributed Demand

Window Length	MODEL S1			MODEL S2	
	Production Variance σ_P^2	Production Change Variance $\sigma_{\Delta P}^2$	Average Inventory \bar{I}/σ_D	Production Variance σ_P^{2*}	Production Change Variance $\sigma_{\Delta P}^{2*}$
1	1.00 (1.00,1.00)	2.00 (2.00,2.00)	0.00 (0.00,0.00)	1.00	2.00
2	.39 † (.392,.387)	.35 (.352,.347)	.34 (.343,.342)	.33	.33
3	.24 (.246,.242)	.14 (.140,.135)	.58 (.585,.581)	.20	.13
4	.18 (.178,.176)	.07 (.076,.072)	.78 (.786,.777)	.14	.07
5	.14 (.139,.139)	.05 (.047,.044)	.95 (.961,.946)	.11	.04
6	.12 (.114,.116)	.03 (.031,.031)	1.11 (1.117,1.096)	.09	.03
7	.10 (.097,.099)	.02 (.023,.023)	1.24 (1.259,1.234)	.08	.02
8	.09 (.084,.087)	.02 (.017,.018)	1.37 (1.388,1.362)	.07	.02
9	.08 (.074,.077)	.01 (.014,.014)	1.49 (1.507,1.482)	.06	.01

†

$\begin{matrix} x \\ (y,z) \end{matrix}$ = x is average result over six cases, $\sigma_D = 5,10,20,30,40,50$.
y is result for $\sigma_D = 5$.
z is result for $\sigma_D = 50$.

two models, except at a window length of two periods. Although the production variances of the two models do differ, the production variance of S1, as a function of the window length, is very similar in shape to that for S2.

We conclude from the simulation study that for normally-distributed i.i.d demand, model S1 can provide significant production smoothing benefits as observed in the case study. Furthermore, the characteristics of the smoothing behavior are predictable from the factors provided in Table 1. Finally we note that the approximate model S2 has some predictive value with regard to the behavior of S1. However the value of S2 seems to be more for explaining and providing insight into the observed behavior of S1, than for predicting this behavior.

The simulation results are for normal i.i.d. demand streams. Cruickshanks and Drescher [2] provide limited evidence that the above results hold for non-normal demand distributions based on the simulation of a "lumpy" demand stream consisting of a mixture of zero demands and normally-distributed demands. They have also considered a serially-correlated demand stream for which they show that the relative benefits from the smoothing model S1 decrease as the demand is more positively correlated.

5. Discussion

We have proposed a production smoothing approach, given by model S1, for use in a job shop environment. In the previous two sections we provided analyses of the smoothing model S1 that indicate and characterize the nature of the benefits available from the production smoothing approach. Our main results are the specification of both the smoothing measures (i.e. production variance and production change variance) and the expected inventory levels, from either Table 1 or the approximate model, as a function of the length of the planning window. In general, we have shown that a little flexibility, as provided by a planning window, generates dramatic smoothing benefits. Yet the question remains as to the proper choice of the planning window. To resolve this question requires an understanding of both the mechanisms available for creating a planning window of a certain length as well as the costs involved in using this planning window.

In the case study, for one product family we estimated the inventory holding costs and we obtained from the production manager a "best guess" at the cost for changing production levels; for details see Cruickshanks and Drescher [2]. We found that for these costs a two-month window minimized the sum of inventory holding and production smoothing costs. Furthermore this result remained valid when we both increased and decreased the cost for changing production levels by 50% of its original value. Hence, the resulting issue was whether or not the benefits from smoothing over a two month window justified the costs required to implement a two-month window. When we presented this issue to the production manager, he indicated that he could (and would) reduce the production lead time for the family by one month to achieve the two-month planning window, and that this reduction in lead time would be essentially costless due to the built-in slack in the current production cycle. Hence, the choice of the planning window was obvious.

In general one cannot expect such luck. The determination of the appropriate planning window requires the proper consideration of inventory holding costs, production smoothing costs, and implementation costs for the planning window. The earlier analyses are useful for determining the inventory and production smoothing consequences of a given planning window. The costs associated with the planning window itself, however, are not so clear. The implementation of a planning window requires that the planned production time be reduced or the promised delivery time be increased or both. We mention two methods for reducing the production time. First, if the production time includes the procurement time for raw materials and parts, then we may stock long-lead time raw materials and parts to avoid some of the procurement component of the production time. The cost for this option is the inventory-related costs for stocking the critical components. Second, we can reduce production time by increasing production capacity, particularly at production bottlenecks. The cost of this option includes the capacity acquisition costs and the cost of supporting underutilized capacity. One benefit of this option is reduced work-in-process inventory. We can also generate a planning window by increasing the promised delivery time. The obvious consequences of this option are lost sales or less profit for a particular product or both.

The determination of the best mechanism to generate a planning window and consequently the cost of the planning window, clearly depends on the production environment. The models and their analyses in this paper give the benefits to expect from a planning window. The comparison of these benefits with the costs will dictate the choice of the planning window.

ACKNOWLEDGEMENT

The authors wish to thank Professor Ken Baker for his helpful comments on an earlier draft of this paper.

References

1. Baker, K. R. and J. W. M. Bertrand, "An Investigation of Due-Date Assignment Rules with Constrained Tightness", Journal of Operations Management, Vol. 1, No. 3 (February 1981), pp. 109-120.
2. Cruickshanks, A. B. and R. D. Drescher, "A Study of Production Smoothing in a Job Shop Environment", Joint S.M. thesis, A. P. Sloan School of Management, M.I.T., January 1982.
3. Hax, A. C., "Aggregate Production Planning", in Handbook of Operations Research, J. Moders and S. Elmaghraby (editors), Van Nostrand Reinhold, New York, 1978.
4. Magee, J. F. and D. M. Boodman, Production Planning and Inventory Control, second edition, McGraw-Hill, New York, 1967.
5. Silver, E. A., "A Tutorial on Production Smoothing and Work Force Balancing", Operations Research, Vol. 15, No. 6 (November-December 1967), pp. 985-1010.

BASEMENT
Date Due

	Date Due	

Lib-26-67

HD28.M414 no.1296- 82
Cruickshanks, J.A. study of production s
743903 D*BKS 00134596



3 9080 002 020 136

