Technical Memorandum Number 829

Setting Safety Stocks and Production Quantities for Product Families

by

A. Dale Flowers
Weatherhead School of Management, Case Western Reserve University,
dale.flowers@case.edu

Serge Karalli
College of Business Administration, California State University
KaralliS@cba4.cba.csus.edu

September 2010

Department of Operations
Weatherhead School of Management
Case Western Reserve University
330 Peter B. Lewis Building
Cleveland, Ohio 44106

Setting Safety Stocks and Production Quantities for Product Families

Corresponding Author

A. Dale Flowers
Department of Operations
Weatherhead School of Management
Case Western Reserve University
10900 Euclid Avenue
Cleveland, OH 44106
dale.flowers@case.edu
Phone: 216-368-3825
Fax: 216-368-6250

Serge M. Karalli
Department of Management
College of Business Administration
California State University, Sacramento
6000 J Street
Sacramento, CA 95819
KaralliS@cba4.cba.csus.edu

Abstract

In many manufacturing environments, parts or products may be naturally grouped into families based on their processing needs. A family setup must be performed before any item in the family can be made. After the family setup is accomplished, a product setup must be done before a specific part or product in the family can be produced. In this situation, production to stock is the norm. The manufacturing equipment used to produce the family is used to produce many other families of products, and demand for the products is uncertain. If setup times are significant, it would be prohibitively expensive to execute the family and product setups every time a customer orders the product. Thus, manufacturers want to run all the products in the family once the family setup is performed, and want to establish run quantities so that a family setup will not be required again for some interval of time. This means that safety stocks must be kept for each part in order to ensure adequate customer service is provided. The question is: How should run quantities and safety stocks be established so as to minimize the total costs of meeting the customer service requirements? We briefly describe a new approach to computing these values and provide insights as to the benefits the new approach will provide to practitioners.

Keywords: Scheduling, inventory, product families, manufacturing, production

Setting Safety Stocks and Production Quantities for Product Families

1. Introduction and Motivation

There has been a lot of published research on the problem of economic lot scheduling, including situations having families of products. Sox et al. (1999) provide the most recent summary of this research. Similarly, a lot of research has been done on establishing safety stock levels, including consideration of both the length of the lead time and the variability of the lead time. Chopra et al. (2004) is a recent example of such work. Surprisingly, a problem situation that occurs regularly in industry that combines economic lot scheduling for part families with safety stock considerations has not been addressed in the literature. The purpose of this research was to help remedy that situation.

There are many examples of this situation in industry. We know of a company in the pipe, valve, and fittings industry that manufactures brass and stainless steel pipe nipples. A major setup is required for the threading machines to change the diameter of the pipe being threaded. A minor part setup is required to change from one length of nipple to another within the same diameter and type of pipe (brass or stainless). The company produces to stock in box quantities, and then simply picks and ships orders from stock to meet their customer requirements.

Another common situation where this problem occurs is in plastic injection molding. Installing a mold that produces a part with a certain shape is typically a major (family) setup, and then changing the color of the plastic to produce a variety of parts constitute the minor (product) setups. Production is typically done in sequence from lightest to darkest to minimize the product setup times.

The manufacture of quartz tubing, ceiling tiles, and synthetic rubber are additional examples where this situation exists. Many chemical processes tend to exhibit these characteristics. If it were possible to reduce setup times to negligible amounts, this problem would not be of critical importance. However, the process technology often precludes such reductions from being achieved. So production planners are left with tools to try to solve the problem that don't really fit the problem structure. The economic lot scheduling procedures that are available do not consider safety stocks in their optimization procedures, with the result that the production quantities will generally be too large. This also causes safety stocks to be too large as well. The consideration of safety stock costs in a solution procedure will naturally shorten the time between production runs and the resulting safety stock required to provide the desired customer service. In this paper, we will illustrate the impact of considering this safety stock for some situations of interest to practitioners.

In the next sections, we will carefully define the problem and the mathematical notation we use, as well as the form of the solution we generate. We will illustrate the impact of our approach on the time between production runs versus an approach that does not consider the safety stock costs in the optimization procedure. We will then examine the impact of family setup costs, demand variation, and the selected service level for our approach versus the alternate procedure.

2. Assumptions and Form of the Solution

The assumptions and characterization of the problem include the following:

- Multiple product families
- Multiple products within each family
- Major setup required for each family
- Minor setup required for each product within the family
- Single, high speed production equipment on which all families are manufactured, one unit at a time, each with a known, constant production rate
- Production to stock, with safety stock required to ensure an adequate level of customer service
- Relevant costs include:
    - Major setup costs
    - Minor setup costs
    - Cycle inventory holding costs
    - Safety stock inventory holding costs

For each family, there is a

sequence-independent setup cost

sequence-independent setup time

For each item, the demand is

time-stationary,

normally distributed, with

known mean and standard deviation,

uncorrelated with other items,

not substitutable with other items

For each item,

there is a sequence-independent setup cost,

there is a sequence-independent setup time,

there is a specified service level,

safety stock is maintained in order to meet the specified service level, and

there is no backlogging

In addition to these assumptions, we have chosen a basic period approach to develop a solution to the problem. Moon et al. (2002) provide a summary of the three approaches commonly taken to the form of the solution for economic lot scheduling problems:

- Common cycle approach
- Basic period approach
- Time-varying lot sizes approach

Moon et al. (2002) list these descriptors for the basic period approach:

- Allows different cycle times for different products
- Requires each product cycle time to be an integer multiple of the basic period
- All lots of each item are the same size
- Gives better solutions that the common cycle approach
- Is NP-hard
- Its main drawback is ensuring feasibility

- Gives inferior solutions to the time-varying lot sizes approach

We chose the basic period approach over the common cycle approach because of its demonstrated advantages over the latter. We chose it versus the time-varying lot size approach for two primary reasons. The first is our explicit inclusion of safety stock in our formulation. Safety stock varies with the length of time between runs, the variance of demand for the item, and the service level specified. To consider time-varying lot sizes would require the readjustment of the target safety stock level every time the lot size changed. This would complicate both the analysis and the implementation of the resulting solution. The second reason is the popularity of Just-In-Time production systems, and the attempt to smooth production as much as possible in order to achieve coordination in the supply chain. Monden (1998) indicates that the Toyota Production System is able to accommodate changes in demand of only +/- 10 per cent. The repeating schedule that results from our approach will facilitate supplier/customer relationships in Just-In-Time manufacturing environments.

The form of our solution to the problem consists of a basic period, as well as a family multiplier for each family, and an item multiplier for each product. This allows us to establish a safety stock inventory level for each item that is needed to protect against stockouts between production runs. To establish safety stock levels in our problem, we use a service level criterion as opposed to a fill rate. The service level criterion is the complement to the probability of a stockout in an inventory replenishment cycle. We have chosen this criterion due to the costly consequences of stocking out of a particular product before its family is scheduled for production. This would appear to be consistent with the practitioner's point of view in such service environments.

3. Mathematical Formulation of the Problem

We seek to minimize the total average relevant cost per unit time in a cyclic schedule (a schedule that eventually repeats itself). The relevant costs include family setup costs, item setup costs, and inventory holding costs for both cycle and safety stocks. We allow a unique service level to be specified for each item. A high service level may be required for items that generate the bulk of the firm's revenues. Low-revenue generating items may be given lower service level requirements.

For notational convenience, we take N to be both the number of families and items within each family, without loss of generality (w.l.o.g.). When the number N exceeds the actual number of items in any family, or exceeds the number of families, we simply create dummy items and/or families and assign the value of zero to their parameters.

The following parameters are inputs to the problem:

S_i Setup time for family i.

s_{ij} Setup time for the item j in family i.

A_i Setup cost for family i.

a_{ij} Setup cost for item j in family i.

d_{ij} Demand mean for item j in family i.

σ_{ij}^2 Demand variance for item j in family i.

p_{ij} Production rate for item j in family i.

ρ_{ij} $= d_{ij} / p_{ij}$, production capacity needed by item j of family i.

ρ $= \Sigma_{i\in N}\Sigma_{j\in N} (d_{ij} / p_{ij})$, total production capacity required

h_{ij} Inventory carry cost for item j in family i.

Z_{ij} Standard deviation from N(0,1) corresponding to the service level required for item j in family i.

The decision variables of the problem are:

T Length of Basic Period

K_i Multiplier of family i

k_{ij} Multiplier for item j of family i

The following additional notation is needed:

$C(\cdot)$ Cost Function Evaluated at $(\cdot)$

$G(\cdot)$ Feasibility constraint, see formulation below

N Set of subscripts of i or j from 1 to N.

K N-Vector of K_i's

k (N x N)-matrix of k_{ij}'s

$\mathbf{k}_{i\bullet}$ The vector of all the item multipliers in family i.

P $=\{2^p: p \in Z_+\}$, set of integer-powers-of-two

Ideally, the firm would like to adopt an inventory policy that:

- minimizes the total relevant operating costs,
- is feasible to implement given the facility's existing capabilities—production rates and capacity, and
- complies with its service criteria.

We represent the solution to the problem by (T, **K**, **k**), where T, **K**, and **k** are defined above. The objective is to find the cycle time, T, and the multipliers K_i and k_{ij} for each family i and its member item j [i.e., to find (T, **K**, **k**)] so as to

$$\text{Minimize} \quad C(T,\mathbf{K},\mathbf{k}) = \sum_{i\in\mathbf{N}} \frac{A_i}{TK_i} + \sum_{i\in\mathbf{N}} \sum_{j\in\mathbf{N}} \left(\frac{a_{ij}}{TK_i k_{ij}} + b_{ij} TK_i k_{ij} + g_{ij}\sqrt{TK_i k_{ij}} \right) \quad (1)$$

$$\text{Subject to} \quad G(T,\mathbf{K},\mathbf{k}) = \sum_{i\in\mathbf{N}} \frac{S_i}{K_i} + \sum_{i\in\mathbf{N}} \sum_{j\in\mathbf{N}} \frac{s_{ij}}{K_i k_{ij}} - (1-\rho)\cdot T \le 0 \quad (2)$$

$$T > 0 \quad (3)$$

$$K_i \in \mathbf{P} \qquad \forall\ i \in \mathbf{N} \quad (4)$$

$$k_{ij} \in \mathbf{P} \qquad \forall\ (i,j) \in \mathbf{N}\times\mathbf{N} \quad (5)$$

$$\text{Where} \quad b_{ij} = \frac{1}{2} h_{ij} d_{ij} \left(1-\rho_{ij}\right) \quad (6)$$

$$g_{ij} = h_{ij} Z_{ij} \sigma_{ij} \quad (7)$$

The objective function, (1), is an average total cost per unit time of all family and item setups, and of holding working and safety stocks. Constraint (2) is a necessary feasibility condition. Constraint (3) can be relaxed to a weak inequality without loss of accuracy. Constraints (4) and (5) restrict the family and item multipliers, respectively, to integer-powers-of-two. Roundy (1989) showed that using integer-power-of-two multipliers, obtained by his polynomial-time Roundoff Algorithm (a heuristic), results in a total cost that is at most 6% higher than the lowest cost of the continuous relaxation of the economic lot scheduling problem.

The details of the solution procedure for the above problem are beyond the scope of this paper, but are available in Karalli and Flowers (2006). Obtaining the solution to our problem does not guarantee that a feasible schedule can subsequently be determined, since we have no guarantee that the families and products can be assigned to production in specific periods so as to be within the capacity constraints of the process. Dobson (1987) demonstrated how the time-varying lot sizing approach facilitated the development of feasible schedules. We have independently addressed this issue for the basic period approach taken in this research (Flowers and Karalli 2008). We developed an algorithm to load families and products to the basic periods in a production cycle and tested 200 example problems. The problems were created utilizing scenarios that seemed reasonable in practical environments. In four per cent of these 200 problems, the algorithm was unable to find feasible schedules. Of course, these results might not be as good if capacity is severely constrained. On the other hand, practitioners often have some flexibility with capacity due to resources such as overtime, so our approach certainly appears to be practical.

4. The Impact of Safety Stock on Production Quantities

Since the previously published procedures for establishing production quantities do not include safety stocks in the cost optimization, it seems reasonable to assume that practitioners would either (1) use procedures that ignore the safety stock costs or (2) make judgmental corrections to try to account for them. We know of no software that considers safety stocks in the context of setting production quantities for product family situations, especially in standard ERP software products. We did not have the ability to test assumption (2) with company data. So we used assumption (1) above, namely, that safety stocks would be ignored when production quantities were established. If production quantities were established that were intended to meet expected demand for some time period, then the practitioner would have to compute the additional safety stock required to meet the desired customer service level after the fact. The safety stock would depend upon the time between production runs resulting from the production quantity. We used this approach to provide a benchmark for comparison to our approach. We generated 30 random problems using a uniform distribution for each input with the parameters shown in Table 1. These parameters were carefully chosen to be representative of the relationships between these variables that we have observed in various real world settings. Each problem had five families each consisting of five products. These problems are available from the authors.

[Insert Table 1 here]

Table 2 shows a summary of the results for these 30 problems for both the benchmark approach and our improved procedure that considers safety stock costs in the optimization. The improved approach generated an average cycle length that was 37% smaller than the benchmark solution approach. For this set of problems, our approach produced solutions with total costs 7.36% lower that the benchmark method. This is a substantial difference for practitioners to exploit in the competitive environments in which they operate. Of course, the values in the table do not cover every situation that can occur in practice. However, it seems clear that this improved approach is worthy of

[Insert Table 2 here]

consideration by practitioners. The last column in the Table reveals that our procedure produced solutions with total costs averaging only 0.6% above a continuous relaxation lower bound for the problem.

5. The Impact of Family Setup Times

To illustrate the impact of family setup times for the benchmark versus our improved procedure, we continued to work with the same 30 problems but we replaced the randomly generated family setup costs from Table 1. Specifically, we resolved all 30 problems using the four constant values of the family setup cost for all families as $100, $1000, $3000, and $5000. These values are discrete points within the range listed in Table 1. By substituting these constant values we could assess how our procedure performed versus the benchmark along this continuum. We created four tables similar to Table 2 for each of the four discrete setup values, and then used the summary information from them to produce Figure 1. It reveals greater percentage savings for the improved procedure for lower values of the setup costs. This is logical since the safety stock costs should dominate in such situations, but the benchmark procedure does not consider them until after the fact. As we move to higher setup costs, the setup costs start to dominate and so the savings level off, but they are still substantial.

[Insert Figure 1 here]

6. The Impact of Demand Variability

To assess the impact of demand variability we performed an experiment similar to that for the setup costs, except this time we used constant values for the coefficient of variation of demand for all products in a family. Since our experience has been that product families occur more often for industrial products than consumer products, and the demand for industrial products is highly variable, we used discrete values of 0.5, 0.7, 0.85, and 0.95 to replace the randomly generated values from Table 1. Again, we solved the resulting 30 problems for each of these four discrete values and generated four tables similar to Table 2, from which Figure 2 was produced. It reveals greater benefit for our procedure versus the benchmark with increasing demand variability. This is logical since the safety stock costs would become more and more dominant with increasing variability, and the benchmark approach is not considering them in the optimization.

[Insert Figure 2 here]

7. The Impact of Service Level

We performed a similar set of experiments with the same 30 original problems except that we used four discrete values for the service level including 90%, 93.5%, 97%, and 99.99%. Again, these values sampled the spectrum of the range of values from Table 1 for this variable. Figure 3 reveals that the advantage of the improved procedure increases with increases in the service level, as would be expected. In such situations, the safety stock costs become more dominant and the benchmark procedure is ignoring them; thus the advantage. Note the familiar exponential increase the Figure suggests since we are assuming that demand is normally distributed.

[Insert Figure 3 here]

8. Summary and Conclusions

We have provided a computational demonstration that for product family manufacturing situations with independent demand, the consideration of safety stock costs as well as family and part setup costs, and cycle stock costs can provide important benefits to the practitioner. The results reveal that if a manufacturing environment resembles the experimental problems that we generated for the analysis reported in this research, then savings of 5 – 10% may be achievable. In addition, shorter production intervals and lower safety stocks are likely to occur. This produces a reduction of risk of obsolescence in the safety stock investment.

Our goals for future research on this problem include one or more implementation projects, extending the results to consider a fill rate criterion, and extensions to allow for trend and seasonal effects on product demands.

References

Chopra, S., Reinhardt, G., & Dada, M. (2004). The Effect of Lead Time Uncertainty on Safety Stocks. *Decision Sciences*, 35, 1-24.

Dobson, G. (1987). The Economic Lot-Scheduling Problem: Achieving Feasibility Using Time-Varying Lot Sizes. *Operations Research*, 35, 764-771.

Flowers, A. D., & Karalli, S. M. (2008). The Family Periodic Loading Problem. Working Paper, Department of Operations, Weatherhead School of Management, Cleveland, Ohio 44106.

Karalli, S. M., & Flowers, A. D. (2006). The Multiple-Family ELSP with Safety Stocks. *Operations Research*, 54, 523-531.

Monden, Y. (1998). *Toyota Production System: An Integrated Approach to Just-In-Time,* 3rd Edition. Engineering & Manaagement Press, p. 28.

Roundy, R. O. (1989). Rounding off to the Powers of Two in Continuous Relaxations of Capacitated Lot Sizing Problems. *Management Science*, 35, 1433-1432.

Sox, C. R., Bowman, R. A., & Muckstadt, J. A. (1999). A Review of the Stochastic Economic Lot Scheduling Problem. *International Journal of Production Economics*, 62, 181-200.

Table 1: Input Data Distributions for Benchmarking

Parameter	**Distribution**
Family setup time (weeks)	U(0.015, 0.025)
Family setup cost ($)	U(100, 5000)
Item setup time (weeks)	U(0.0012, 0.018)
Item setup cost ($)	U(50, 150)
Item holding cost ($/unit/week)	U(0.01, 1.25)
Item demand mean (units/week)	U(10, 1000)
Item coefficient of variation	U(0.5, 0.95)
Item production rate (units/week)	U(10,000, 100,000)
Item service level (probability all demand is met)	U(0.90, 0.9999)

Problem	Improved Cycle Time	Benchmark Cycle Time	Improved Total Cost (TC)	TC Improvement Over Benchmark Cost	TC Difference from Lower Bound
1	2.44	3.54	$30,979.69	7.54%	0.76%
2	0.97	1.77	$29,408.47	10.66%	0.78%
3	1.23	2.66	$24,829.34	7.06%	0.78%
4	1.13	2.34	$24,381.88	17.96%	0.64%
5	1.34	1.94	$21,939.60	4.73%	0.51%
6	1.40	2.03	$31,282.69	4.84%	0.46%
7	1.48	2.30	$27,834.24	6.97%	0.76%
8	2.12	2.21	$29,531.07	6.13%	0.47%
9	2.09	2.14	$34,352.77	7.90%	0.89%
10	1.19	2.28	$21,476.60	9.08%	0.12%
11	0.91	2.10	$30,916.01	5.08%	1.09%
12	1.39	2.44	$24,276.36	6.95%	0.31%
13	2.42	2.86	$27,075.30	4.95%	0.67%
14	1.20	2.07	$28,537.15	7.52%	0.42%
15	2.06	3.02	$18,888.43	4.14%	0.65%
16	1.38	2.32	$29,174.21	7.61%	0.41%
17	0.78	2.17	$34,224.12	5.79%	0.50%
18	1.18	1.89	$24,008.61	8.27%	0.25%
19	1.45	2.10	$28,491.49	4.68%	0.41%
20	3.25	4.32	$26,581.10	3.65%	0.82%
21	2.00	3.02	$26,556.87	5.75%	0.94%
22	1.55	2.18	$27,679.74	6.98%	0.56%
23	1.29	3.33	$22,381.54	6.02%	0.52%
24	1.40	2.11	$35,307.31	5.48%	1.01%
25	1.69	2.57	$29,106.95	5.67%	0.45%
26	1.27	1.73	$27,368.77	5.68%	0.69%
27	1.93	2.35	$28,319.57	9.96%	0.21%
28	1.22	2.60	$27,207.19	6.78%	0.33%
29	0.83	1.64	$27,544.82	10.08%	0.89%
30	0.97	1.91	$15,844.66	16.74%	0.74%
Mean	1.52	2.40	$27,183.55	7.36%	0.60%
Std Dev	0.555	0.583	$4,351.37	3.24%	0.25%
Max	3.25	4.32	$35,307.31	17.96%	1.09%
Min	0.78	1.64	$15,844.66	3.65%	0.12%

Table 2: Results of the Original Runs

Figure 1: Impact of Family Setup Costs

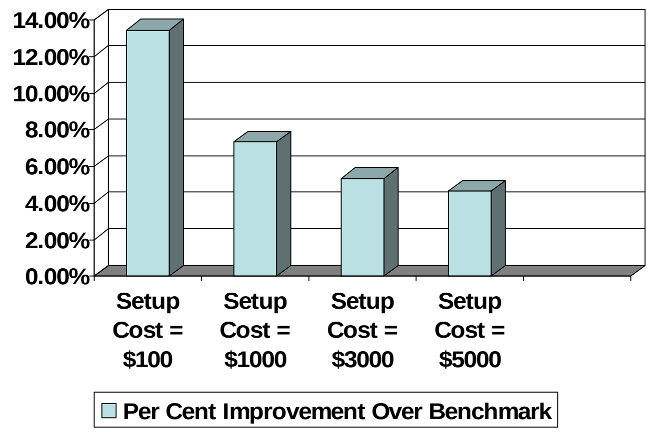

Figure 2: Impact of Demand Variability

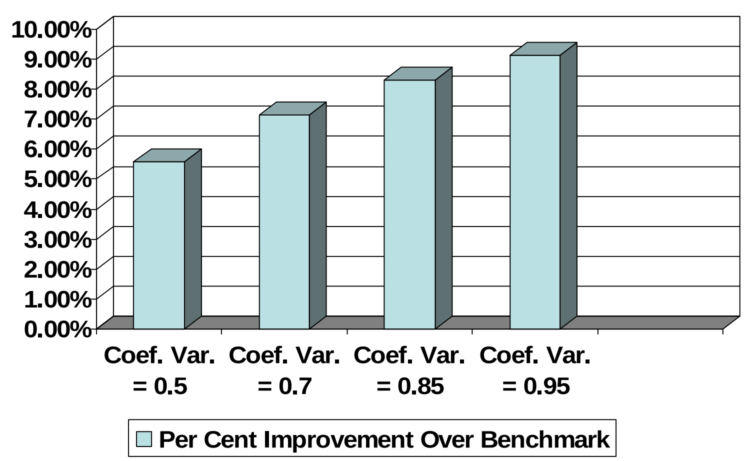

Figure 3: Impact of Service Level