Demand Behavior in Manufacturing Supply Chain- Model, and Computational Study

Mary J. Meixell
The Bell Laboratories, Lucent Technologies, Princeton, NJ
S. David Wu
Department of Industrial and Manufacturing Systems Engineering
Lehigh University, Bethlehem, PA

Abstract

Production planning policies have a direct impact on demand propagation and the overall performance of manufacturing supply chains. The relationship between production planning and supply chain behavior is important because the planning decisions play a major role in determining how demand propagates through the supply tiers. In this paper, we study supply chain demand behavior via computational experiments. The purpose of our study is to examine the experimental results against the propositions from (Wu and Meixell, 1998), and to gain a better understanding of the demand behavior and the operational costs. For the purpose of studying supply chain behavior, we first develop a solution method for the multi-level, multi-period lot-sizing model on which the supply chain demand model is based. This is followed by a discussion on the hypotheses pertaining to supply chain behavior, and a compendium of the results from the computational experiments. Each of the experiments is evaluated relative to the propositions posed in (Wu and Meixell, 1998). We conclude from the study several managerial insights and practical recommendations that may improve supply chain performance.

1 A Solution Method for the Supply Chain Production Model

We first describe a particular version of the multi-level, multi-period lot-sizing model to study demand propagation in manufacturing supply chains. This model is a special instance of the general model discussed in (Wu and Meixell, 1998). The objective we use here follows a conventional lot-sizeing cost objective, minimizing setup and inventory costs.

$$\text{Minimize} \sum_{i=1}^{N} \sum_{t=1}^{T} (h_i y_{it} + c_i d_{it}) \qquad (1)$$

subject to:

$$y_{i,t-1} + f_i x_{i,t-L_i} - y_{it} - \sum_{j=1}^{N} a_{ij} x_{jt} = r_{it} \qquad ,\forall i,t \qquad (2)$$

$$\sum_{i \in I_k} (b_{ik} x_{it} + s_{ik} d_{it}) \le CAP_{kt} \qquad ,\forall k,t \qquad (3)$$

$$x_{it} - q_{it} d_{it} \le 0 \qquad ,\forall i,t \qquad (4)$$

$$d_{it} = 0{,}1, \qquad x_{it} \ge 0\ , \qquad y_{it} \ge 0 \qquad ,\forall i,t \qquad (5),(6),(7)$$

<u>Indices:</u>

i = 1,…, N	index of products and items
j = 1,…, N	index of products
t = 1,…., T	index of planning periods
k = 1,…,K	index of facilities
I_k =	index set of items that are produced at facility k

<u>Parameters:</u>

h_i	= inventory of holding cost ($ per unit of item *i*)
c_i	= set-up cost ($ per set-up of item *i*)
L_i	= minimum lead time for item *i*
f_i	= yield of item *i* (fraction)
a_{ij}	= number of units of item *i* required for the production of one unit of item *j*
r_{it}	= demand for item *i* in period t
b_{ik}	= capacity utilization rate of item *i* at facility *k* (capacity units per unit)
s_{ik}	= set-up utilization of facility capacity k by item *i* (capacity units)
CAP_{kt}	= capacity of facility *k* at time *t* (units of capacity)
q_{it}	= upper bound on the production of item *i* that can be initialized in period *t*

<u>Variables:</u>

y_{it}	= inventory of item *i* at the end of period *t*
δ_{it}	= 1 if item made in period *t*, 0 otherwise
x_{it}	= production of item *i* initialized in period *t*

1.1 Computation Complexity of Multilevel Lot-sizing Algorithms

The multi-level lot-sizing problem is computationally very difficult to solve. Many authors have shown that even finding feasible solution in some cases is NP-Complete. In the case of lot-sizing problems with setup times, Maes, et al, (1991) stated that it is "a very difficult task, far more difficult than has been suggested in the literature." They show that finding a feasible solution to this problem is NP-complete. At best, setup costs are sometimes used as a Lagrangean substitute for setup times. Maes and his co-authors use an LP-based heuristic with a rounding routine to solve a production lot size problem for the assembly structure case.

Salomon (1991) addressed the complexity of algorithms for a variety of lotsizing sub-problems, and ultimately proceeds to use heuristic search techniques to solve his lotsizing problem. Importantly, Salomon notes that a problem with multiple facilities and multiple items with general production costs but no setup costs or setup times is polynomially solvable. This problem is equivalent to a transportation problem, which is solvable in polynomial time. Salomon also shows that in the case of a single facility with multiple items, the problem is NP-hard if *either* setup times or setup costs are present. He notes that the un-capacitated item-level problem has a polynomial time dynamic programming algorithm that could be useful in solving this style of problems. This well-known algorithm is attributed to Wagner & Whitin, and in its original version, has $O(n^2)$ complexity. Federgruen and Tzur (1991) and Wagenmans, et.al (1992) have each published the results of their research that improves on this complexity with algorithms that solve the general dynamic lot size problem in $O(n \log n)$ time, with a special case that solves in $O(n)$ time.

Florian, et al (1980) address the complexity of algorithms for special cases of the single-item, capacitated production planning problem without setup times. When the capacities are equal for this single-item problem, and when a concave cost function is specified (e.g. a setup cost exists), a polynomial algorithm of $O(n^4)$ does exist (Florian & Klein, 1971). If the capacities are not equal, however, the single item problem is NP-hard. In fact, Garey and Johnson (1979) list the related feasibility problem as one of the set of known NP-complete problems.

From this brief survey it is easy to conclude that the existence of a polynomial time algorithm for our manufacturing supply chain model is extremely unlikely. Alternatively, a good approximation algorithm is needed.

1.2 A Solution Method for the Supply Chain Production Model

The complexity results are at first discouraging in terms of finding a good solution method for the supply chain production model. But as we will show, they are useful in selecting an approach to solving the production planning model that serves as a driver for the demand model we develop here. A number of authors (Tempelmeier & Derstroff (1996), Shapiro (1993), Diaby, et.al. (1992), Billington, et. al. (1986), Thizy & Wassenhove (1985)), have shown that a Lagrangean-based heuristic is a promising approach for solving the lot-sizing problem. Kimms (1997) discusses Lagrangean relaxation as a useful approach for multi-level lot sizing when standard MIP solvers cannot find a solution in reasonable time. Lagrangean relaxation both provides a lower bound value for evaluation purposes, and a starting point solution for a feasibility-restoring heuristic. We too rely on a Lagrangean-based heuristic with good average-case efficiency for the computational studies of this supply chain demand model.

A number of options exist for simplifying this problem via Lagrangean relaxation. One approach is to relax only the capacity constraint, yielding *M* single end item multi-level un-capacitated lot sizing problems. Unfortunately, no efficient algorithm exists for this sub-problem, and although our research objective is to study the behavior of supply chain demand, this approach is not good enough. A second strategy is to relax only the balance constraint, resulting in a single level capacitated problem. This problem will likely be multi-item, and possibly multi-facility. Again, no efficient algorithm exists for this sub-problem.

Attempts have been made to solve the multi-item, multi-level lot-sizing problem by relaxing both the capacity and balance constraint. (Billington, et al, (1983) Tempelmeier and Derstroff, (1996) and Thizy & Wassenhove (1985)) This generates a set of single-item, un-capacitated subproblems with a known efficient algorithm. Unfortunately, the duality gap that result from this relaxation can be quite large, and so the routine that is needed to restore feasibility to the solution will be quite involved and may wander quite a distance from the relaxed solution.

Here, we choose to relax only the capacity constraint and extract the dependent demand term from the balance constraint. This unlinks the items in a single chain and allows for a single item un-capacitated subproblem. The dependent demand is computed and summed within the Lagrangean loop. This detail will be discussed further in the next section on the solution method.

A situation that results from this handling of the dependent demand term is the impact of the lower bound. We no longer have a "true" lower bound when a term is extracted from a constraint of the original problem. We test the quality of this "surrogate" lower bound and report on the results in Section 4 on the performance of this heuristic.

The supply chain problem, then, is an optimization problem that for practically sized problems is too big to be solved with exact methods. Even small supply chains can consist of hundreds of items, across several facilities and tiers. Accordingly, we develop and use a heuristic that decomposes the master problem into smaller sub-problems. The procedure is a Lagrangean-based heuristic that solves a relaxed version of the original problem, compares it to a feasible upper bound solution, and iterates by re-computing multipliers and re-solving the problem until a good solution has been found. Although this sub-problem is readily solvable using a number of techniques, we use the CPEX MIP solver for the 25 period sub-problems. This heuristic procedure, outlined in Figure 1, is described in greater detail in this section.

1.2.1 Generating a Feasible, Initial Upper Bound Solution

An upper bound solution is computed at the start of this algorithm with a procedure that mimics a decentralized MRP system. We solve for each item in the product structure from top to bottom, by first computing the total demand, solving the dynamic lot size problem, and then computing the remaining capacity at each facility. This procedure continues until the capacity in the current period is exhausted. We than make use of capacity from earlier periods, “building ahead” by fulfilling demands using earlier period capacities.

Step 1. Compute an initial upper bound solution
For each item in i=1..N
Compute demand for item i
Solve the item-level dynamic lot-size problem using remaining capacity
Re-compute remaining capacity at facility k
End For
Step 2. Find the lower bound solution form the capacity-relaxed supply chain problem (Subgradient Search)
2.1 For each item in i=1..N
Compute demand for item i
Solve the item i sub-problem
End For
2.2 Compute the lower bound
2.3 Re-compute gradients, step sizes, multipliers
2.4 Check the stopping criteria, terminate if satisfied, keep iterating if not
Step 3. Restore feasibility (Upper Bound Calculation)
3.1 Identify capacity in-feasibilities for all facilities for all time periods
3.1.1 Find item produced at facility k in period t
3.1.2 Search backward in time for period where item can be built ahead
3.1.3 Adjust production plan for item i at facility k in period t
3.2 Identify balance in-feasibilities resulting from capacity-based plan adjustments for all items in all time periods
3.2.1 Re-run item-level sub-problem

Figure 1: The Lagrangean Heuristic for the Supply Chain Problem

As shown in Figure 1, the first step in this procedure is to compute demand for the first item - the external and the internally generated demands. The external demand is an input to the procedure. The derived demand can be computed using the dependent demand term, $\sum_{j=1}^{N} a_{ij} x_{jt}$. It is important in this procedure that the items are sequenced so that parent items are always computed before component items.

The resulting dynamic lot size problem is then solved for the single item. The capacity constraint treats each item as though it is the only item produced at that facility. Once a production plan is derived for the item, the "remaining capacity" for the facility is computed and used for the remaining item problems. This step ensures that capacity feasibility is maintained. $RCAP_{kt}$ is thus defined as:

$$CAP_{kt} - \sum_{j=1}^{M} \left(b_{jk} * x_{jt} + s_{jk} * d_{jt}\right)$$

where: $RCAP_{kt}$ — capacity of facility k at time t after items *1-M* have been produced

$j = 1,\ldots, M$ — index set of parent products and items for i

The upper bound sub-problem, then, is:

Minimize $\sum_{t}^{T} (h * y_t + c * d_t)$ (12)

subject to:

$$y_{t-1} + f * x_{t-L} - y_t = r_t \qquad ,\forall t \qquad (13)$$

$$(b_k x_t + s_k d_t) \le RCAP_{kt} \qquad ,\forall k,t \qquad (14)$$

Also needed are constraint sets (4) through (7).

1.2.2 The Lagrangean Relaxed Problem

The next step of the heuristic is to solve the relaxed version of the original problem, compare it to a feasible upper bound solution, and iterate until a good enough solution is found. At each stage of this iteration, the algorithm re-computes multipliers and re-solves the relaxed version of the problem.

The capacity constraint but not the balanced constraint, is relaxed and pulled into the objective function in this formulation. This is done for two reasons. The balance constraint in a supply chain problem is a critical constraint. In practice, a customer facility in a supply chain cannot begin production of any product unless all the components are available in-house. The balance constraint is, then, in effect, a material availability constraint. By leaving it intact in the sub-problem, we can assure that the material availability constraint is given higher priority. A second reason for relaxing only one set of constraints is to maintain high quality heuristic solutions.

This leaves the issue of the dependent demand term in the balance constraint, which bundles the constraints by linking the item sub-problems. We handle this complicating term by extracting and then embedding it in the loop that executes the sub-problems within each Lagrangean relaxation iteration. Treating the balance constraint in this way carries a penalty, however. The lower bound generated by the heuristic is a lower bound to the specific problem only, not to the master problem that we're ultimately trying to solve. The stopping criteria ensure proper handling of the multipliers. But the quality of the final heuristic solution in some problem implementations may be affected.

In this way, the supply chain production planning problem is broken down into a set of manageable item-level sub-problems. This sub-problem is:

$$\text{Minimize } \sum_{t}^{T}(h*y_t + c*d_t) - \sum_{k=1}^{K} l_{kt}(CAP_{kt} - b_k x_t - s_k d_t) \qquad (15)$$

subject to:

$$y_{t-1} + f*x_{t-L} - y_t = r_t \qquad ,\forall t \qquad (16)$$

< constraints (4) through (7)>.

The Lagrangean multipliers are computed using sub-gradient optimization. The gradient G_{kt} is the difference between the capacity and the time required for setup and production at each facility in each time period. Specifically, it is computed as follows:

$$G_{kt} = \sum_{i \in I_k} \left(CAP_{kt} - s_{ik} d_{it} - b_{ik} x_{it} \right)$$

The step size for multiplier update is computed as follows:

$$\Gamma = p\left(z_{UB} - z_{LB}\right) \Big/ \sum_{t} G_{kt}^{2}$$

The multipliers are updated at each iteration as the sum of the previous multiplier value, l'_{kt}, and the product of the step size and the gradient. Since the capacity constraint is an inequality, the multiplier value must always be non-negative. Specifically,

$$l_{kt} = \max\left(0,\ l'_{kt} + \Gamma \cdot G_{kt}\right)$$

The procedure iterates until no further improvement can be found, an optimal solution is found, or a maximum number of iterations is reached.

1.2.3 Restoring Feasibility to the Lower Bound Solution

The third step of the heuristic is to restore feasibility to the solution generated by the Lagrangean relaxation. This is accomplished by first identifying and resolving capacity violations, followed by a step that identifies and resolves balance violations.

The algorithm identifies facilities that have capacity violations by scanning all periods for each facility in the product structure. As each violation is detected, backward scheduling is used to eliminate the violation. The routine searches through the list of items produced at that facility and finds the first one that has production scheduled in that violated time period. Then, the excess production for that item is shifted to the next earliest time period with available capacity.

The use of backward scheduling in this routine is consistent with how violations are handled in practice. Since backordering or adding capacity is not allowed, producing an item early is the only allowable action for an over-scheduled facility. The excess demand is “built ahead” to ensure that customer demand is met on time.

Shifting production for any item in the supply chain certainly changes the demand pattern for that facility's suppliers, and possibly drives a need to change the supplier's production plan as well. In this way, a schedule change propagates down the chain. Therefore, once a production plan has been modified for capacity shortfalls, the algorithm searches for balance violations. These balance violations are corrected by re-computing the schedule for that specific violated item. This routine iterates through all facilities, items, and time periods to ensure that the resulting solution is fully feasible.

1.2.4 Performance of the Heuristic

Table 1 illustrates the performance of the heuristic procedure relative to the optimal solution. Sixteen smaller problem instances were randomly selected for this comparison, each with 10 items, 4 demand periods scheduled across 8 time periods in 3 facilities. The table illustrates the following statistics for these problem instances:

- Initial upper bound
- Value of the heuristic solution using the surrogate lower bound.
- Value of the optimal solution for the problem,
- Deviation of the heuristic solution using a surrogate lower bound vs. optimal
- Value of the heuristic solution using the true lower bound
- Deviation of the heuristic solution using the surrogate lower bound vs. true lower bound

Three comparisons can be drawn from Table 1. First, the heuristic improves the initial feasible solution in almost all cases. Recall that the initial UB resembles a local optimization routine, where each item in turn across the entire supply chain is optimized. The items are solved from the top of the product structure to the bottom, in a manner similar to a typical MRP driven process. The improvement versus the UB is important in this study in part because the computational effort of the heuristic is justified, but also because the initial feasible solution gives an estimate of the performance of an uncoordinated supply chain. This comparison suggests that coordinating across even a small supply chain can improve performance.

The heuristic solution using a surrogate lower bound deviates from the optimal solution for this set of problem instances by 16% on average. For this set of small supply chain problems, the optimal solution can be found by solving the full supply chain problem as a single MIP. The range of deviation from optimal is displayed in column 5.

The table also shows that using a surrogate lower bound instead of a true lower bound in the Lagrangean relaxation carries a very small penalty. Recall that we chose an item-level subproblem that did not fully de-couple the problem constraints. We artificially removed the dependent demand term in the balance constraint and computed its value in-between the sub-problem iterations. Most of the problems (75%) we analyzed here had the same solution regardless of whether the surrogate or true lower bound was used. The average objective value penalty was 1%.

Problem Instance	**Initial UB Solution**	**Heuristic Solution (surrogate LB)**	**Optimal Solution**	**Deviation From Optimal**	**Heuristic Solution (true LB)**	**Deviation From True LB**
G1	2,479	2,262	2,019	.12	2,262	.00
G2	5,906	5,646	5,278	.07	5,646	.00
G3	6,152	5,970	5,190	.15	5,970	.00
G4	5,900	5,900	5,678	.04	5,900	.00
G5	6,893	6,436	6,214	.04	6,436	.00
G6	6,152	5,970	5,107	.17	5,763	.04
G7	1,803	1,803	1,712	.05	1,755	.03
G8	1,975	1,925	1,721	.12	1,811	.06
A1	2,000	2,000	1,217	.64	2,000	.00
A2	15,648	13,101	10,681	.23	13,534	.03
A3	10,323	5,243	5,196	.01	5,243	.00

Problem Instance	Initial UB Solution	Heuristic Solution (surrogate LB)	Optimal Solution	Deviation From Optimal	Heuristic Solution (true LB)	Deviation From True LB
A4	6,052	5,097	4,992	.02	5,097	.00
A5	14,724	11,608	9,026	.29	11,608	.00
A6	15,485	11,117	10,547	.05	11,117	.00
A7	4,921	4,853	3,273	.48	4,853	.00
A8	6,974	5,935	3,707	.14	5,935	.00
			Average=	.16	Average=	.01

2 Computational Studies

In (Wu and Meixell, 1998), some theoretical results are discussed pertaining to various demand behaviors in the supply chain. Specifically, we analyze the effects of order batching, multiple schedule releases, product design and capacity leveling. To verify and refine the theoretic results in a more complex and realistic setting, we constructed a set of computational experiments using the supply chain model described earlier. We implement the Lagrangean-based heuristic procedure using AMPL/CPLEX and use the heuristic solution to drive the supply chain decisions through the tiers. This provides us with a “best of conditions” simulation of production decisions in action. The experiments are ran in a 200 MHz Pentium PC with 64 MB RAM. Details of this study are presented in the following sections.

2.1 The Design of the Supply Chain Experiments

We use analysis of variance and multiple regression analysis to investigate the relationship between manufacturing planning and supply chain performance. For each of three experiments, we identify the factors that we believe will explain the differences in the response variables, and then estimate the size of these differences. Both main effects and their interactions are studied. Blocking is used wherever possible to increase the precision of the experiments, and randomization is used in cases with no known source of variation to block. In all cases, the structure of the model errors are evaluated and verified by inspecting the residual plots and the normal probability plots, with statistical consistency checks where necessary. The three main experiments are summarized as follows:

1. The first experiment is a *screening experiment* where we use demand amplification and total supply chain inventory as response variables. Four main factors examined are product structure, capacity utilization , coefficient of variation (c.v.) for end item demand and setup cost.

2. In the second set of experiment we study the *influence of end item order patterns* on supply chain performance. For this experiment, we used a multi-factor design, building on the results of the screening experiment, with four treatment levels to examine the magnitude and direction of the factor effects.

3. In the third experiment, we investigate the influence of *schedule release policies* on supply chain performance. This experiment also builds on the result of the screening experiment, adding a factor to examine two schedule release policies: using the most recent schedule release as internal demand, vs. using the mean of all releases up-to-date.

2.2 Test Problems

Data that represents a wide variety of supply chain environments are necessary to study behavior of manufacturing supply chains. For this reason, we used the data that Tempelmeier and Derstroff developed to test a multi-item, multi-stage production model (Tempelmeier and Derstroff, 1996). This data contains a large number of randomly generated problem instances that vary systematically in:

- product and operation structure,
- setup time,
- time between order profile,
- capacity utilization,
- CV of end item demand, and,
- setup cost.

We use their definition of stages to define the tiers in supply chains. We selected at random from the 1200 available instances the set of problem instances necessary to populate our experimental design. Figure x illustrates the two product structures contained in this database.

The problems each consist of 40 items with 16 period demand, produced at 6 different facilities, distributed across 5 tiers in a manufacturing supply chain. In all cases, lead-time is set equal 1 and the initial inventory is 0. The factors we focused on that were present in the database include product structure, capacity utilization, CV demand, and setup cost. To this set, we added two additional characteristics of interest in this study - order pattern type and multiple schedule releases type. In this way, we supplemented the original database to suit our experimental needs.

We used a demand series for each end item as determined by Tempelmeier & Derstroff. These series were generated by fixing the mean of demand for each end item and a specified CV={.1,.9}. A truncated normal distribution was then used to generate one demand series for each end item in each of the two product structures. All component demand in the system is derived demand, so other demands are computed as the solution is executed.

The utilization data we selected contained two different capacity utilization profiles, with each facility's target utilization *p*={50%, 90%}. Available capacity for each problem is computed to be the mean demand divided by this target capacity utilization assuming a lot-for-lot system. In a system with setup times, it is impossible to preset the utilization a priori to knowing the solution to each problem. Tempelmeier & Derstroff computationally derived an appropriate adjusted utilization level for each subclass of problem instances that approximates the desired target utilizations. This is the capacity utilization level used in these computational studies.

The feasibility of the demand series relative to capacity was also verified by Tempelemeier & Derstroff to ensure that the cumulative sum of demand over time for each time period does not exceed the cumulative sum of capacity. When this occurred, demands were exchanged in the series to eliminate the violation, as it creates an infeasible problem instance. These exchanges cause an increasing trend in the demand series when setup times exist with heavy capacity utilization.

3.3 Experiment 1: Determine the Effects of Main Factors

The first set of experiments examines the effects of product structure, planned capacity utilization, end-item demand variation and setup costs. The experiment is set out to address the following questions:

1) Which of these factors have a significant influence on demand amplification?
2) What combination of factor levels lead to better supply chain performance?
3) Can supply chain performance be managed with a design factor alone, and if so, how?

This experiment is a screening experiment where we use demand amplification and total supply chain inventory as response variables. Total demand amplification TDA is computed as the maximum difference in demand variation between any two tiers from the second through the fifth. Tier-to-tier demand variation between tier k and l is computed as follows:

$$TDA = Max_{k,l}\{CV(X^{l}) - CV(D^{k})\}$$

where X^{l} and D^{k} are the demands for all items over all time periods in tiers l and k, respectively. The first tier demand is excluded from the calculation because the variation in end-item demand will influence the result while being exogenous to the system. A factorial design with three factors and two levels is used. For each of the eight treatment combinations, there are eight replicates in each design block, selected randomly across the un-controlled parameters in this experiment.

The factors and their levels are product structure (assembly/general), capacity utilization (.95/.55), and variation in end item demand (.9/.1). The assembly and general product structures are the same as those defined in Tempelmeier and Derstroff (1996), and are described in Figure 2. The capacity utilization is the planned level – for each experiment, we average demand at each facility, and compute the necessary capacity to drive an average utilization level as defined for each experiment. The end item demand variation is the coefficient of variation of the 16-period demand stream.

3.3.1 Results for Experiment 1: The Effect of Product structure, Utilization and Setup Cost on Demand amplification

The screening experiments look at the impact of the three initially selected planning factors - product structure, capacity utilization, and CV(end item demand) - on both response variables - demand amplification and inventory cost. The impact of setup cost is then evaluated using regression analysis for each of these measures. This experiment shows that of the three selected production planning factors, only product structure is significant at a confidence level of 95%. The mean of the response variable, demand amplification, is 0.55 for general product structures and 0.33 for assembly product structures. The analysis of variance also indicates that none of the factor interactions is significant. Table 2 displays the details of these ANOVA results. Residual plots and the normal probability plot are displayed in Appendix x. The residuals appear structure-less and normally distributed.

A graph of the CV of demand across tiers (Figure 3) further illustrates the significance of product structure, and leads to an extension of this analysis. Note how the CV of demand appears to be larger for general product structures than for assembly product structures.

Table 2. ANOVA Results for the Production Planning Screening Experiment – Demand Amplification Variable

Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
MAIN EFFECTS					
A: Utilization	0.2853	1	0.2853	2.07	0.1559
B: CV_Demand	0.0001	1	0.0001	0.00	0.9812
C: Product Structure	0.7843	1	0.7843	5.68	0.0205
INTERACTIONS					
AB	0.0110	1	0.0110	0.08	0.7788
AC	0.1819	1	0.1819	1.32	0.2558
BC	0.0083	1	0.0083	0.06	0.8067
RESIDUAL	7.866	57	0.1400		
TOTAL (CORRECTED)	9.1370	63			

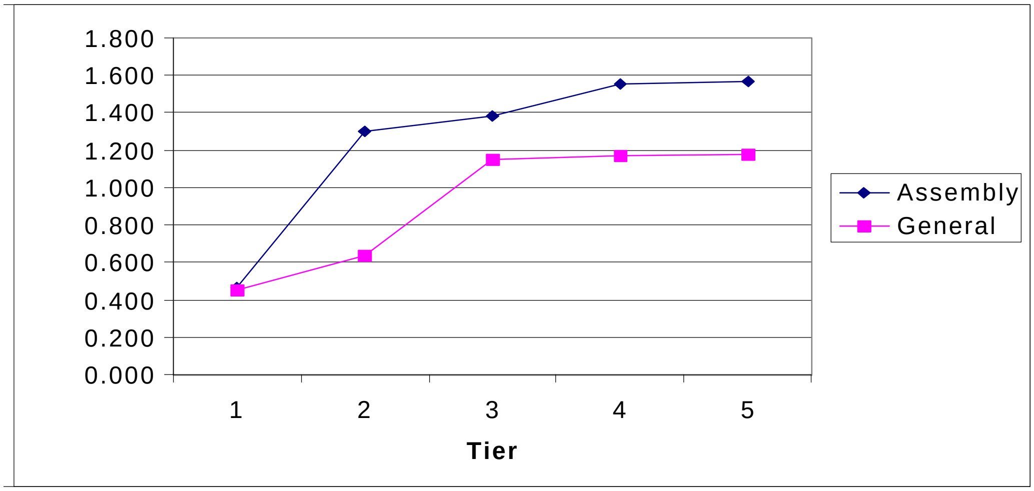

Figure 3. Comparing Total Demand Amplification between Assembly and General Production Structures

At a first glance, it appears that there is a higher but more consistent level of variation in assembly supply chains. Closer review of the included cases indicates, however, that the problem instances chosen for assembly differ from those chosen for general in setup cost. There are 3 low setup cost and 11 high setup cost instances for assembly, and 9 low setup cost and 3 high setup cost instances for general product structures. It is conceivable that more high setup cost cases populating the assembly product structure set would influence the variation in demand. Although this factor was not included in the original screening experiment, we now include an extension to study the impact of setup cost on demand amplification.

Since the experimental data in this analysis is unbalanced for setup costs, a multiple regression analysis is done to study the influence of product structure, utilization, and setup cost on demand amplification. CV demand of the end item was evaluated and found to be not significant at the 90% confidence level, and so was removed from the model. The regression results for the 3-factor model are listed in Table 3, followed by the analysis of variance for the unbalanced model in Table 4. Note that the P-value for the model is 0.01, suggesting a high confidence level (99%) that the variables selected provide a valid description of demand amplification. The adjusted R^2 for this model is 58.4%. The Drubin-Watson statistic, used to test the residuals for autocorrelation, is 1.64, sufficiently higher than the critical statistic 1.4.

Table 3. Multiple Regression Analysis for the Demand Amplification Screening Experiment

Parameter	Estimate	Standard Error	T Statistic	P-Value
CONSTANT	-0.1496	0.1083	-1.3808	0.1826
ind_sucost	0.5881	0.1048	5.6141	0.0000
ind_utilization	-0.1980	0.0761	-2.6013	0.0171
ind_product structure	0.2886	0.1018	2.8345	0.0102

Table 4. ANOVA for the Demand Amplification Screening Experiment

Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
Model	1.1619	3	0.3873	11.76	0.0001
Residual	0.6584	20	0.0329		
TOTAL (Corrected)	1.8203	23			

Figures 4 through 7 provide further insight relative to the behavior of demand in supply chains. The regression analysis suggests that setup cost is a significant factor in demand amplification, and we see how that is manifested in Figure 4. High setup costs appear to drive more variation as demand propagates through a supply chain for both general and assembly product structures.

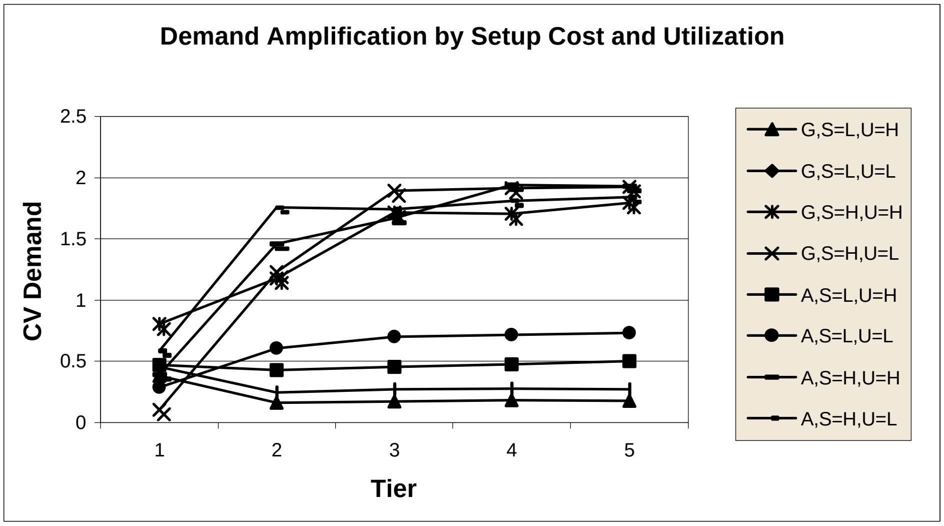

Figure 4. Demand Amplification by Setup Cost and Utilization

At first glance, it appears that there is a higher but more consistent level of variation in assembly supply chains. Closer review of the included cases indicates, however, that the problem instances chosen for assembly differ from those chosen for general in setup cost. The analysis also indicates that utilization is a significant factor in demand propagation. At low utilization, more demand amplification occurs in supply chains than with high utilization cases. Figure 6 illustrates this point. Note that for both assembly and general product structures, the high utilization line falls below the low utilization line. The analysis also suggests that product structure influences demand amplification, as we can see from Figure 5. For low setup costs, a slight demand de-amplification is observed for general product structure- there is less variance at lower tiers than higher tiers in the supply chain. However, this de-amplification behavior does not occur in assembly product structures.

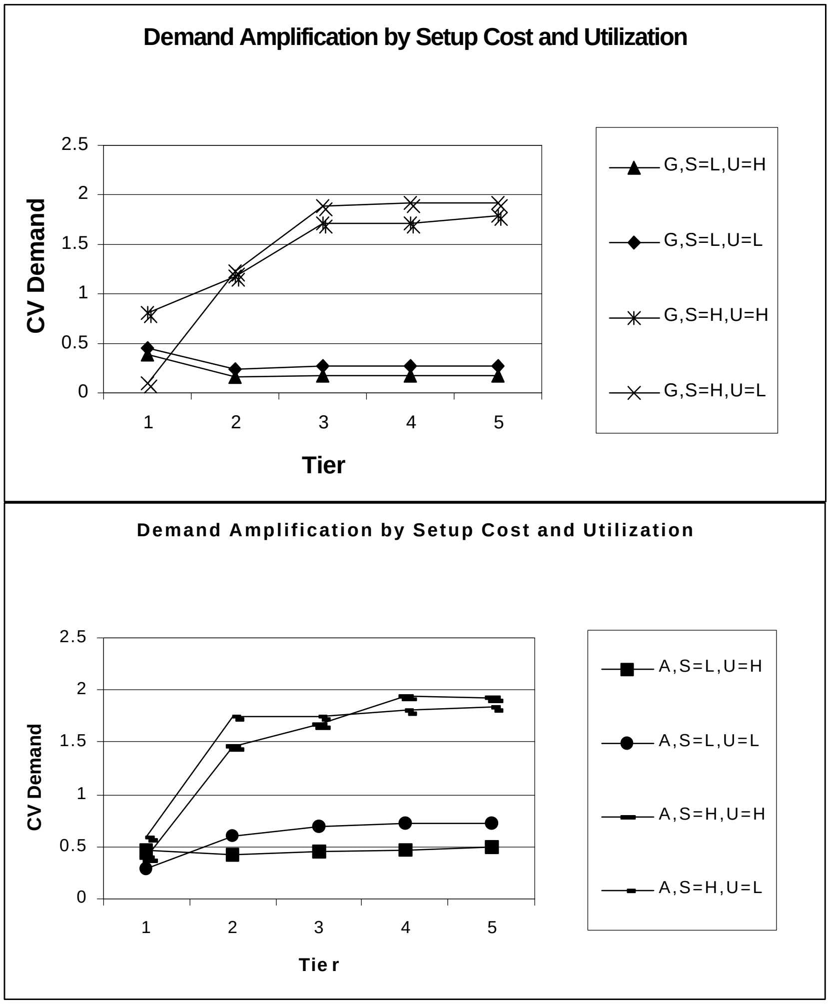

Figure 5. Demand Amplification by Setup Costs and Utilization

Also note the tendency for variation in demand to reach a limit, or "cap", especially for the high setup cost cases, and level off. For high setup costs for both product structures types, this demand variation cap is in the range of 1.7-1.9. This cap on variation naturally occurs in a manufacturing supply chains because of manufacturing capacity. Lot sizes are limited by the amount of manufacturing capacity available in any one production period, since we do not allow setups to carryover between production periods.

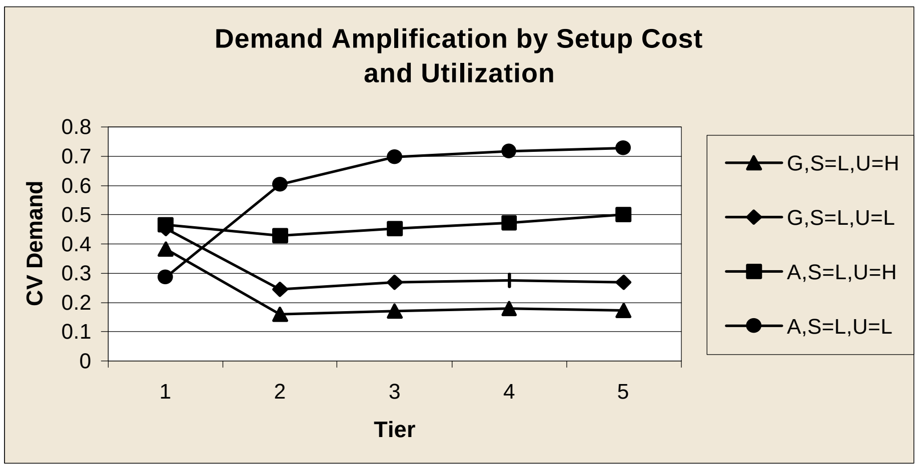

Figure 6. Demand Amplification by Setup Cost and Utilizaiton

3.3.2 Results for Experiment 1: The Effect of Setup Cost and Utilization on Inventory performance.

A second response variable is also investigated relative to manufacturing supply chain performance, and shows that of the three selected production planning factors both product structure (at 95% confidence level) and utilization (at 90% confidence level) are significant. The mean of the response variable, total inventory in the chain equalized by total demand, is 0.44 for general product structures and 0.65 for assembly product structures. This inventory measure is 0.46 for low utilization case and 0.63 for the high utilization case. The analysis of variance indicates that none of the factor interactions is significant. A plot of the residuals and the normal probability plot are displayed in Appendix x. The residuals appear structure-less and normally distributed. Table 4 displays the details of these results.

Table 4. ANOVA Results for the Production Planning Screening Experiment – Inventory Response Variable

Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
MAIN EFFECTS					
A: Utilization	0.4691	1	0.4691	3.45	0.0685
B: CV_Demand	0.3446	1	0.3446	2.53	0.1170
C: Product Structure	0.7616	1	0.7616	5.60	0.0214
INTERACTIONS					
AB	0.0226	1	0.0226	0.17	0.6851
AC	0.1218	1	0.1218	0.90	0.3480
BC	0.0015	1	0.0015	0.01	0.9167
RESIDUAL	7.7531	57	0.1360		
TOTAL (CORRECTED)	9.4744	63			

Again, we specify and fit parameters for the model using regressions analysis to fit a linear model with unbalanced data, to determine the influence of setup cost on the inventory response variable. An analysis of the P-values indicates that the product structure and end item demand variation indicator variables were not statistically significant at the 90% confidence level, and so both were removed from the model. The remaining variables - setup cost and utilization - do have a significant influence at the 99% confidence level on the inventory in the type of supply chains studied in this study, as indicated by the results of the regression analysis. The adjusted R^2 for this model is 81.2%. The Durbin-Watson statistic, used to test the residuals for autocorrelation, is 1.93, sufficiently higher than the critical statistic 1.4.

Table 5. Multiple Regression Analysis for the Inventory Screening Experiment

Parameter	Estimate	Standard Error	T Statistic	P-Value
CONSTANT	0.1063	0.0599	1.7758	0.0890
ind_sucost	0.7468	0.0716	10.4241	0.0000
ind_utilization	-0.2007	0.0714	-2.8093	0.0100

Table 6. ANOVA for the Inventory Screening Experiment

Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
Model	3.5658	3	1.7829	55.07	0.0000
Residual	0.7446	23	0.0324		
TOTAL (Corrected)	4.3104	26			

Consequently, we can reject the null hypothesis that the inventory means are equal for setup cost and utilization with 99% confidence, but we cannot reject the null hypothesis that the means are equal for the product structure and CV of end item demand factors.

3.3.3 Results from Experiment 1: Discussion

We will now examine the questions posed earlier regarding Experiment 1.

1. *Which of the production planning factors has the greatest influence on demand amplification?*

– Product structure appears to have a significant influence on supply chain performance. In fact, we see that in general product structures, demand variance may actually be *reduced* both for single release (due to negative correlation and to the effect of a summation of random upstream lot sizes) and for multiple release demand series.

– Setup cost also appears to significantly influence supply chain performance. Because high setup costs tend to increase batch size, they also tend to increase demand variance along a chain.

– Utilization appears to influence supply chain performance too, because binding capacity in any period at any facility reduces the lot size and so reduces downstream demand variation.

2. *What combination of factor levels lead to better supply chain performance?*

– General product structures with low setup cost and high utilization drives a lower amount of demand amplification. These are the favorable conditions for minimum demand amplification.

– Low setup cost and low utilization drives the least amount of inventory in supply chains, regardless of product structure. Some inventory must be incurred, then, if utilization is high, even when setup costs are low. Inventory and demand amplification do not have identical response to these design factors.

– Another attractive combination of factor levels is a level order pattern with low setup costs and low utilization. In this case, regardless of the product structure, a zero variation demand pattern at the end item level will translate without alteration throughout the supply chain. Note that this alternative requires an operational control – leveling the order pattern – in addition to the design factors.

3. *Can supply chain performance be managed with a design factor alone, and if so, how?*

– Designing a supply chain with a general product structure, low setup cost, and high utilization throughout a supply chain can reduce or eliminate demand amplification. This finding is supported by both the propositions from (Wu and Meixell, 1998) and the experimental evidence presented here in this paper.

In the following, we summarize the experimental results in light of the theoretic results from (Wu and Meixell, 1998).

1. ***The effects of order batching.*** *Propositions 1*states that when there is no incentive to consolidate batches in the supply chain there will be no item-component demand amplification. In the experiments, the incentive for batching is provided by high setup costs as related to the inventory holding cost. Thus, the experimental results strongly support the above point (see Figures 4 and 5); in the low setup cost cases there is almost no demand amplification across the tiers. On the other hand, *Proposition 2* and *Corollary 3.1* state that when demands are consolidated into production batches over time, the mean and variance for a lower tier component will

both increase. As evident from the experimental results (Tables 4-6), when the setup cost is high, more inventory is carried in the chain because of the larger variance in internal demands. When tier-to-tier demand amplification is used, as is the case in the experiment, *Propositions 7.1* and *7.2* predict that demand amplification increases from one supply tier to another as a function of the lot size. While *Proposition 6* predicts that the setup capability of the system and the system capacity will limit the batch size, therefore the amount of amplification. The experimental results (see Figures 4 and 5) strongly support the theoretic prediction. Not only does the tier-to-tier demand amplification increase drastically after first batching occurs from Tier 1 to Tier 2, the amplification taper off at the third tier since limited capacity prevent larger lot sizes to be formed.

2. ***The effects of product design and component sharing.*** *Proposition 5* and *Corollary 4.1* states that when the components manufactured in the supply chain are shared by a large number of upstream products, the fluctuation in end-item demands tends to create a more significant level of variance. However, when putting into the perspective of total volume, the amplification effect for products with general structure is less then that of the assembly products. The experimental results support this point, showing that the general product structure tend to experience a much less significant demand amplification then that of the assembly products. When order batching is not a factor (in the low setup cost cases), the general product structure demonstarte a negative trend in demand amplication. This is again supported by *Corollary 4.1.,* which predicts a negative demand amplification for general product structure when the upper tier demands are independent.

3. ***The smoothing effects of limited capacity.*** *Propostion 6* states that when the effects of order batching is excluded from consideration, limited capacity has a smoothing effect on demand amplification. When manufactured in a capacity-bounded facility, the orders eminated from the facility will invariably have a lower variance than the orders entering it. As pointed out above in 1, this phenomenon explains a tapered

demand amplification below a certain level of the supply tier. Since the test problems used in the experiments have indeed a capacity level close to its true expected demand, the observed amplificaiton profile is well explained by the Proposition. Also, we see from the results in Section 3.2 that high capacity utilization drives more inventories in a supply chain, regardless of product structure. When the capacity is limited, some items must be built ahead that would otherwise be built in a later production period. Since demand is realized in a later production period, it follows that higher capacity utilization *drives more inventory* in a supply chain.

3.4 Experiment 2: The Influence of End Item Order Pattern on Supply Chain Performance

In this set of experiments we study the influence of end-item demand patterns to demand amplifications and total inventory levels in the supply chain. We test four different end item demand patterns (see Figure 7) to determine which order pattern, if any, may potentially improve overall supply chain performance. These patterns could be induced by a variety of mechanisms, such as a demand management mechanism that prices or provides other motivation to customers to place orders according to this pattern. Also, the end item manufacturers could carry finished inventory to cause this pattern to occur.

Order pattern can be also influenced at the end item level via scheduling and inventory mechanisms, such as order leveling, a common practice in the industry: if a particularly large number of orders with a critical component are received in the current time period, delay some of these orders and schedule them in the next time period (and possibly incur late or expedited delivery), or in the previous time period, if known with sufficient lead time (incurring inventory costs). On the other hand, order balancing can be accomplished by scheduling a specific version for a critical component in every other, or possibly every third period. Again, inventory or late delivery charges may apply here.

The first of the four patterns tested represents random demand, as orders might flow direct from customers. This random demand case shows a moderately small amount of variation across the sixteen period planning horizon. To create the second pattern, the level demand case, we set each period demand equal to the average demand. The balanced demand case is the third pattern, created by batching orders into four, four-period blocks, across the planning horizon, staggered to provide synchronized order flow over all order types. The fourth case, positively correlated demand, batches each of these four-period blocks into the same time slot – a worst case alternative. Lee, et al (1997) analyzes these same order batch rules.

It is the goal of this experiment to address questions pertaining to:

1. Is production leveling, a cornerstone in JIT production systems, a useful tool for reducing demand amplification in a multi-tier environment?
2. Does balancing demand for end items improve supply chain performance?
3. How does such a case compare to the extreme scenario when all orders arrive together (the correlated demand case)?
4. How do each of these alternatives compare to the base scenario, the random case?

In this set of experiments, we study the influence of end item order patterns on supply chain performance. We measure the response in terms of demand amplification and total supply chain inventory, equalized for differences in total demand between the two product structure problem sets. We use a multi-factor design, building on the results of the screening experiment, with four treatment levels to examine the magnitude and direction of the factor effects. The four levels of order pattern (Figure 7)- random, level, balanced, and positively correlated - each contain twelve replicates, identical except for the values that comprise the end item demand array.

3.4.1 Results from Experiment 2: The Effect of End-Item Order Patterns on Demand Amplification

This experiment explores the benefit that might be gained by controlling the end item demand patterns at the top tier in a supply chain. We tested the significant factors from the screening experiments for demand amplification – product structure, setup cost, and utilization – along with the order pattern factor. We found numerous significant interactions between the main effect factors. Table 7 presents the analysis of variance for the demand amplification response variable. A plot of the residuals and the normal probability plot are displayed in Appendix x. Again, the residuals appear structure-less and normally distributed.

Table 7. ANOVA Results for the Demand Amplification/Order Pattern Experiment

Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
MAIN EFFECTS					
A: Order Pattern	0.2049	3	0.0683	2.99	0.0521
B: Product Structure	1.3581	1	1.3581	59.40	0.000
C: Setup Cost	2.2356	2	1.1178	48.89	0.000
D:Utilization	0.4569	1	0.4569	19.99	0.0002
INTERACTIONS					
AB	0.1726	3	0.0575	2.52	0.0833
AC	0.1735	6	0.0289	1.26	0.3117
AD	0.0618	3	0.0206	0.90	0.4559
BC	0.8119	2	0.4059	17.76	0.0000
BD	0.0554	1	0.0554	2.42	0.1333
CD	0.4244	2	0.2122	9.28	0.0011
RESIDUAL	0.5258	23	0.0229		
TOTAL (CORRECTED)	6.4809	47			

Table 7 shows that order pattern is not significant due to the fact that product structure, setup cost and utilization have a significant effect on lot sizing. After examining the interaction plot for order pattern and product structure it appears that the interaction occur with the negatively correlated order pattern. This behavior supports *Corollary 4.1,* which states that a negatively correlated demand will reduce demand amplification for general product structure. So, the negative correlation at the end item tier initially causes de-amplification, and a reduced level of demand variation. The negative correlation pattern is immediately absorbed in tier 2 of the supply chain, followed by the usual batching and capacity smoothing effects. In other words, the benefit seems to merely delay the point where demand amplification starts.

As a result, we do not see an overall improvement in supply chain behavior driven by any of supposedly beneficial end item order patterns. Common belief is that balancing or leveling orders with respect to critical components at any level in a supply chain will reduce amplification across a single release. We see here, however, that demand amplification does not differ among the end item order pattern types.

At a first glance, the weak influence of order pattern on demand amplification maybe counter-intuitive. After some further examination it is clear that the lot sizing decisions at each supply tier absorb the ultimate effect of end-item demand variation. Where setup costs are sufficiently high to warrant batching orders from future periods, any value due to beneficial order patterns at the end item level will not translate down the chain. Likewise, the capacity smoothing effects tend to mitigate the influence of end-item demand pattern. This failure to translate order patterns is beneficial when top tier patterns are undesirable. The positively correlated order pattern is quickly distributed across production weeks at the top tier, and so bottom tiers see little or no effect. Demand patterns are constantly absorbed in middle tiers of the supply chain. What eventually translates through to the last tier is the long term trend in total volume. This suggests that lower tiers suppliers do not experience and do not need to react to short term peaks and valleys.

In conclusion, we can not reject the null hypothesis that the mean for demand amplification across different order patterns are the same. This is likely due to the role of the lot sizing policy, which absorb the initial order pattern within one or two supply levels.

3.4.2 Results from Experiment 2: The Effect of End-Item Order Patterns and Inventory

Our experimental results again show a significant factor interaction on the inventory response between order pattern and setup cost at the 95% confidence level. Utilization shows to be a significant factor again in this inventory response measure. A plot of the residuals and the normal probability plot are displayed in Appendix x. The residuals appear structure-less and normally distributed. Table 8 is the ANOVA table summarizing these results.

Table 8. ANOVA Results for the Inventory/Order Pattern Experiment

Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
MAIN EFFECTS					
A: Order Pattern	2.7000	3	0.8999	2.54	0.0754
B: Setup Cost	97.2337	2	48.6168	137.06	0.0000
D: Utilization	9.0679	1	9.0679	25.56	0.0000
INTERACTIONS					
AB	7.2656	6	1.2109	3.41	0.0109
AC	1.4236	3	0.4745	1.34	0.2807
BC	0.2574	2	0.1287	0.36	0.6987

RESIDUAL	10.6413	30	0.3547		
TOTAL (CORRECTED)	128.589	47			

This experiment illustrates that the lower the level of planned capacity utilization the lower the inventory cost. Lower utilization means that orders are built ahead infrequently, and so the total inventory cost is lower.

Similar to the demand amplification response, the interaction plot reveals that the interaction occurs for level and random order patterns for the inventory response variable. The gap in inventory is wider between high and low setup costs for both the level and the random order patterns. For the level pattern, for low setup costs, the inventory is lower than expected. For the random order pattern, for high setup costs, inventory is higher than expected. This interaction muffles the results of the analysis of variance.

A closer inspection of the experimental results illustrates some interesting behaviors of supply chains. Figure 8 shows three statistics for two types of systems with level end item order pattern. Relative inventory is the average inventory at each tier, adjusted for demand. Equalized lot size is the average number of weeks of demand built in a single lot, and is computed for each tier. The third measure is the CV of demand at each tier, and it reflects the amplification in the system. Note that the level order pattern with both low utilization and low setup costs illustrated in Figure 8 mimics the behavior of a well-structured JIT system. End item orders are level across all production periods, and since sufficient capacity is available at all facilities, the level order patterns translates perfectly down the chain. There is no variation in demand, lot-for-lot production is economical, and consequently, there is no inventory in this system.

Contrast this with the high utilization, high setup cost environment for the level order pattern depicted in Figure 8. Even though the same "perfect" order flow is introduced to the chain, demand amplifies, lot sizes are large and growing, and inventory is relatively high.

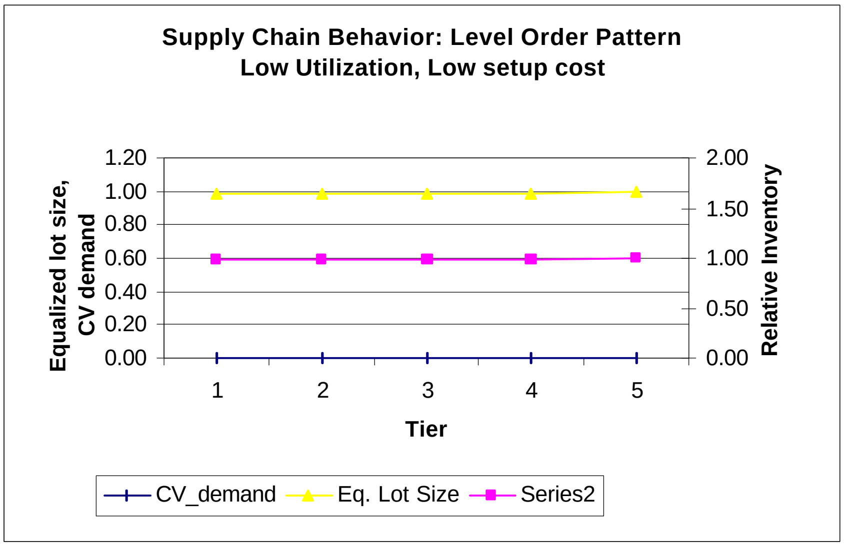

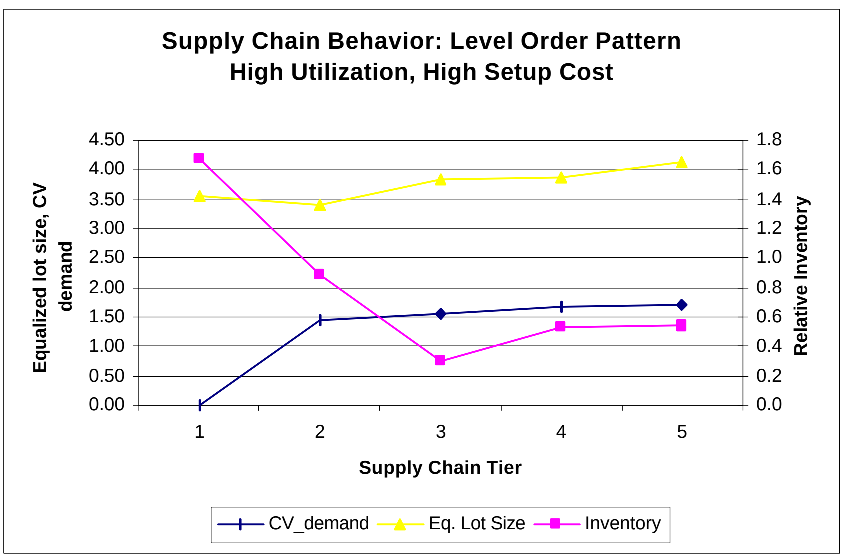

Figure 8. Level Order Pattern with Low and High Setup Costs

3.4.3 Results from Experiment 2: Discussion

We will now examine the questions posed earlier regarding Experiment 2.

1. *Is production leveling, a cornerstone in JIT production systems, a useful tool for reducing demand amplification in a multi-tier environment?*

– Level production is a useful tool to reduce amplification only if both low setup costs and low utilization are in place in a supply chain. Otherwise, a level order pattern will not translate through a supply chain due to lot sizing.

2. *Does balancing demand for end items improve supply chain performance? How does this alternative compare to the base scenario, the random case?*

– It appears that balancing demand, that is, creating a negatively correlated demand pattern, has a de-amplification effect on demand on supply chains. De-amplification is not induced by a random order pattern. The balanced order pattern, however, is quickly absorbed in the supply chain, followed by the usual setup cost, capacity smoothing, and product structure lot sizing effects.

3. *How does the balanced demand case compare to the extreme scenario when all orders arrive together (positively correlated demand)? How does this alternative compare to the base scenario, the random case?*

– The interactions detected from the experiments suggest that positively correlated demand does appear to influence inventory, but not demand amplification, for the low-setup-cost cases. When setup costs are relatively low and orders arrive in a positively correlated manner, more inventory is carried in the chain than would otherwise be required with random order pattern.

Although none of the propositions posed in (Wu and Meixell, 1998) address end-item demand patterns directly, some related issues deserve discussion here:

1. ***The Effect of Order Batching.*** We see in this experiment additional effects of order batching. It is the high setup costs that drive the growth in lot sizes, and as a result the effects of end-item order pattern are absorbed in upper supply tiers. Also interesting in the high utilization/high setup cost system is the inventory behavior. The inventory never bottoms out as with the low utilization, low setup cost case, but it does drop to less than 1/3rd of the peak before increasing again. The reduction in inventory occurs because batching at a higher tier limits the range of lot-sizing opportunities at the lower tiers. Tier 1 batches orders in this example equivalent to about 3.5 periods of demand. Tier 2, and the others down the chain, are compelled to build these same lot sizes, as the cost structure provides no opportunity for

breaking down large orders into smaller lot sizes. The tiers down the chain carry less inventory because their order sizes are essentially lot-for-lot.

2. ***The Effect of Capacity Limitation.*** Note that in the high-setup high-utilization case in Figure 8, the demand variation reaches a limit, imposed by the capacity of the system, and levels off, here at about tier 2. At the same time, the inventory first drop off due to the lot-for-lot production suggested above, until the point where capacity is no longer sufficient to build these lot sizes in the same period. In which case it starts to build ahead and thence increase the inventory level.

In conclusion, we find that we could not reject the null hypothesis that the mean for inventory across the order pattern levels are the same. There are again a number significant supply chain effects in play here, as stated in the propositions. We did not find that order pattern influences either deamnd amplification or inventory as a main effect, but that interactions do influence the results.

3.5 Experiment 3:The Influence of Multiple Schedule Releases on Supply Chain Behavior

In the third experiment, we investigate the influence of multiple schedule releases on supply chain performance. How sales and production planners at each tier in the supply chain view and process demand also influences demand amplification. In practice, planners typically process each order release from their customers as a complete, and deterministic, description of demand. Alternatively, planners can treat an order release as a single instance of the probability distribution that describes demand. In this experiment, we investigate the importance of treating demand as random in reducing demand amplification and inventory costs by addressing these questions:

1. *Is it better to treat demand as a random variable, and schedule supplier production using the expected value of that demand?*

2. *Is there a supplier scheduling policy that would improve supply chain performance?*

In conducting the experiments, two responses were tested - demand amplification and total supply chain inventory. This experiment also builds on the result of the screening experiment, adding a factor to represent the two types multiple schedule release policies: LR, the last release, and MN the mean of the distribution of demand, each with twelve replicates. LR corresponds to the freeze-up-to policy discussed in (Wu and Meixell, 1998), and MN correspond to the policy using expected demand of all available releases.

The data for the LR experimental runs are created assuming that demand is normally distributed (μ,σ). We exclude, however, from the sampling space any values less than the mean value, μ, to avoid a backordered condition. This is, in effect, a “half normal” distribution. Four variates were selected from this distribution for each of the twelve replicates. Conceptually, then, these variates represent four instances of the “last customer release”, and are used here to contrast to the behavior of a supply chain when the expected value of the demand distribution is used. The mean demand case is then adjusted to the true expected value of this truncated distribution.

This third experiment explores the benefit that might be gained from processing the expected value of the distribution of demand instead of the latest demand signal generated by the customer’s MRP system. We look at the influence of this factor on both demand amplification and on inventory.

3.5.1 Results from Experiment 3: The Influence of Schedule Release Policy on Demand Amplification

We first consider the effect of multiple schedule releases along with product structure on demand amplification. Here schedule release policy does not significantly influence demand amplification at the 90% confidence level, and so we cannot reject the null hypothesis that the mean values of the treatment levels are equal. The analysis of variance also indicates that the factor interactions are not significant. A plot of the residuals and the normal probability plot are displayed in Appendix x. The residuals appear structure-less and normally distributed. Table 9 displays the details of these results.

Table 9. ANOVA Results for the Multiple schedule releases Experiment Response: Demand Amplification

Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
MAIN EFFECTS					
A: Release Policy	0.0002	1	0.0002	0.01	0.949
B: Product Structure	0.0614	1	0.0614	1.11	0.304
INTERACTIONS					
AB	0.0237	1	0.0237	0.43	0.520
RESIDUAL	1.106	20	0.0553		
TOTAL (CORRECTED)	1.19092	23			

To consider setup cost along with multiple schedule releases, we use regression analysis since the data from the experiment is unbalanced when setup costs are included. We see that no gain can be demonstrated in terms of demand amplification at the 90% confidence level. Table 10 displays this demand amplification result. The adjusted R^2 for the demand amplification model is 5.3%. The Durbin-Watson statistic, used to test the residuals for autocorrelation, is 1.48, nominally higher than the critical statistic 1.4.

Table 10. Multiple Regression Analysis for the Demand Amplification/Multiple schedule releases Experiment

Parameter	Estimate	Standard Error	T Statistic	P-Value
CONSTANT	0.0209	0.0909	2.5409	0.0199
ind_Release policy	-0.0062	0.0904	-0.0682	0.9463
ind_Product structure	-0.0851	0.0922	-0.9226	0.3678
ind_Setup cost	0.2108	0.1054	1.9992	0.0601
ind_Utilization	-0.0966	0.1085	-0.8903	0.3844

Table 11. ANOVA for the Demand Amplification/Multiple schedule releases Experiment

Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
Model	0.2590	4	0.0674	1.32	0.2985
Residual	0.9319	19	0.0491		
TOTAL (Corrected)	1.1909	23			

Scheduling the mean instead of the last release appears to have no impact on demand amplification. Consequently, we cannot reject the null hypothesis that the means for demand amplification across the two schedule release policies are the same. Note that the design of the experiment is to test for amplification effects when demand is random across multiple releases. The set of “last releases” for each observation models the random nature of demand. The style of demand amplification considered here in our demand amplification measurement considers variation growth across a single release.

3.5.2 Results from Experiment 3: The Influences of Schedule Release Policy on Inventory Performance.

The second part of Experiment 3 considers the effect of multiple schedule releases and product structure on inventory. Here schedule release policy is significant to the inventory response, and so we reject the null hypothesis that the mean values of the treatment levels are equal with 95% confidence. The analysis of variance also indicates that the factor interaction is not significant. A plot of the residuals and the normal probability plot are displayed in Appendix x. The residuals appear structure-less and normally distributed. Table 12 displays the details of these results.

Table 12. ANOVA Results for the Multiple schedule releases Experiment Response: Inventory Cost

Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
MAIN EFFECTS					
A: Product Structure	140.8	1	140.8	1.23	0.280
B: Release Policy	902.0	1	902.0	7.90	0.011
INTERACTIONS					
AB	6.156	1	6.156	0.05	0.819
RESIDUAL	2284.1	20	114.203		
TOTAL (CORRECTED)	3333.07	23			

To consider setup cost along with multiple schedule releases type, we use regression analysis since the data in the experiment are unbalanced when setup costs are included as a factor. The results of the regression analysis do state with 99% certainty that inventory is influenced by the schedule release policies in a manufacturing supply chain. The inventory cost for scheduling the last release for this set of representative supply chain problems is 23.1 across the entire chain, contrasted with 11.8 when the mean is used. This is a 55% reduction in inventory.

Table 13 displays the inventory result. The adjusted R^2 for the demand amplification model is 66.5%. The Durbin-Watson statistic, used to test the residuals for autocorrelation, is 1.80, sufficiently higher than the critical statistic 1.4.

Table 13. Multiple Regression Analysis for the Inventory/Multiple schedule releases Experiment

Parameter	Estimate	Standard Error	T Statistic	P-Value
CONSTANT	12.569	2.844	4.419	0.0003
ind_Product Structure	4.845	2.844	1.703	0.1040
ind_Release Policy	-12.261	2.844	-4.311	0.0003
ind_Setup Cost	14.829	2.844	5.214	0.0000

Table 14. ANOVA for the Inventory Multiple schedule releases Experiment

Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
Model	2362.33	3	787.445	16.22	0.0000
Residual	970.74	20	48.537		
TOTAL (Corrected)	3333.07	23			

Using the mean instead of the last release does appear to drive better inventory performance. We can therefore reject the null hypothesis with 99% confidence that the means for inventory across the two schedule release policies are the same. The results of this study indicate, then, that it is beneficial for a production planner to use the expected value of the demand quantity in lieu of the latest release.

3.5.3 Results from Experiment 3: Discussion

We will now examine the questions posed earlier regarding Experiment 3.

1. *Is it better to treat demand as a random variable, and schedule supplier production using the expected value of that demand?*

– Scheduling the mean instead of the last release does not appear to have a significant impact on demand amplification, however, it does have a significant impact to the inventory performance. This suggests that using the expected demand values reduces the variance in demand down the supply tiers thence reduces the needs for keeping a high level of inventory.

2. *Is there a supplier scheduling policy that would improve supply chain performance?*

– There is experimental evidence that suggests that scheduling the mean of the distribution of demand instead of the latest release will improve inventory performance.

As related to the propositions from (Wu and Meixell, 1998), we have the following observation. *Propositions 3* and *4* suggest that when upper tier customers make multiple schedule releases, it is preferable to follow the expected demand over all up-to-date releases rather than following any particular single release. As stated in *Proposition 4*, following the expected value of the multiple releases rather than following any particular schedule release will result in a much lower variance. On the other hand, as stated in *Proposition 3*, if the releases and the period-to-period demands are iid, multiple releases do not differ from single release in terms of demand amplification. Both propositions are supported strongly from the experiments. While the experiment shows no significant difference in demand amplification between the two policies, the difference in inventory performance is significant. This corresponds to the prediction that a higher variance is to be expected for the LR policy, which results in a higher inventory level.

4. Conclusions

Demand propagation is a fundamental behavior of manufacturing supply chains. Understanding this phenomenon is essential to improving manufacturing supply chain performance. The growth of variation in supply chains – demand amplification - is particularly detrimental to performance because it adds unnecessary cost to the system and drives un-reliability in delivery. In this chapter, we study these behaviors of manufacturing supply chains using an optimization-based modeling framework. Major findings of the experiments can be summarized as follows:

1. Setup cost, product structure, and capacity utilization (roughly in that order) have a significant influence on supply chain performance in terms of demand amplification and inventory level.

2. General product structures with low setup cost and high utilization drives a lower amount of demand amplification. These are the favorable conditions for minimum demand amplification. Low setup cost and low utilization drives the least amount of inventory in supply chains, regardless of product structure. Some inventory must be incurred, then, if utilization is high, even when setup costs are low. Inventory and demand amplification do not have identical response to these design factors.

3. Designing a supply chain with a general product structure, low setup cost, and high utilization throughout the chain can reduce or eliminate demand amplification. This finding is supported by both the propositions from (Wu and Meixell, 1998) and the experimental evidence presented here in This paper.

4. The end-item order patterns appear to be quickly absorbed in the first one or two tiers in the supply chain due to order batching. As a result, manipulating end-item order patterns appears to be quite fruitless unless the supply chain has both low setup costs and low utilization. A combination unlikely to be adopted by manufacturers.

5. Schedule release policy that uses the mean of the distribution of demand instead of the latest release will improve inventory performance due to a reduction in demand variance. However, no significant effect should be expected in terms of demand amplification.

.

Bibliography

Bahl, H.C., L.P. Ritzman, J.N.D. Gupta, (1987) "Determining Lot Sizes and Resource Requirements: A Review" *Operations Research*, Volume 35, number 3, May-June, pages 329-345

Baker, K.R. (1993) "Requirements Planning" in *Logistics of Production and Inventory*, *Handbooks in Operations Research and Management Science*, Volume 4, ed. S.C. Graves, et al, Elsevier Science Publishers, Amsterdam, The Netherlands.

Billington, Peter J., John O. McClain, and L. Joseph Thomas. (1983). "Mathematical Programming Approaches to Capacity-Constrained MRP Systems: Review, Formulation and Problem Reduction." *Management Science*, vol. 29, no. 10, October, pages 1126-1141.

Box, George E.P., William G. Hunter, and J. Stuart Hunter, (1978*) Statistics for Experimenters: An Introduction to Design, Data Analysis, and Model Building,* John Wiley & Sons, New York.

Diaby, M. Bahl, H.C., Karwan, M.H., Zionts, S., (1992) "A Lagrangean relaxation approach for very-large-scale capacitated lot-sizing", *Management Science* volume 38, pages 1329-1340.

Federgruen, A., and M Tzur, (1992) "A Simple Forward Algorithm to Solve General Dynamic Lot sizing Models with n periods in *O(n log n)* or *O(n)* Time." *Management Science*, volume 37, number 8, pages 909-925.

Florian, M., J.K. Lenstra, and A.H.G. Rinnooy Kan, (1980) "Deterministic Production Planning: Algorithms and Complexity" *Management Science*, vol. 26, no 7, July, pp. 669-679

Garey, M.R. and D.S. Johnson, (1979*) Computers and Intractability: A Guide to the Theory of NP-Completeness.* W.H. Freeman and Company, San Francisco.

Harrison, T.P. and H.S. Lewis, (1996) "Lot Sizing in Serial Assembly Systems with Multiple Constrained Resources," *Management Science*, 42, 1, pp. 19-36.

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

Kimms Alf, (1997) Multi-Level Lot Sizing and Scheduling: Methods for Capacitated, Dynamic and Deterministic Models, Physica-Verlag, Heidelberg, Germany.

Lee, H.L., V. Padmanabhan, and S. Whang, (1997) "Information Distortion in a Supply Chain: The Bullwhip Effect," *Management Science*, 43, 4, pages 546-558.

Maes, J., J.O. McClain, L.N. Van Wassenhove, (1991) "Multilevel capacitated lotsizing complexity and LP-based heuristics" *European Journal of Operational Research*, 53, pp.131-148.

Montgomery, Douglas C. (1991) *Design and Analysis of Experiments*, 3rd edition, John Wiley & Sons, New York.

Neter, John, William Wasserman, and Michael H. Kutner (1985) *Applied Linear Statistical Models*, 2nd edition, Richard D. Irwin, Inc. Homewood, Illinois.

Salomon, Marc, (1991) "Deterministic Lotsizing Models for Production Planning." *Lecture Notes in Economics and Mathematical Systems*, Managing Editors: M Beckmann and W. Krelle, Springer-Verlag, New York.

Shapiro, J.F. (1993) "Mathematical Programming Models and Methods for Production Planning and Scheduling" in *Logistics of Production and Inventory, Handbooks in Operations Research and Management Science, Volume 4*, ed. S.C. Graves, et al, Elsevier Science Publishers, Amsterdam, The Netherlands.

Tempelmeier, H. and M. Derstroff, (1996) "A Lagrangean-based Heuristic for Dynamic Multilevel Multiitem Constrained Lotsizing with Setup Times," *Management Science*, 42, 5, 738-757.

Thizy, J.M. and L.N. Van Wassenhove, (1985) "Lagrangean Relaxation for the Multi-Item Capacitated Lot-Sizing Problem: A Heuristic Implementation" IIE Transactions, Volume 17, Number 4, pages 308-313.

Wagelmans, A. S. van Hoesel, and A. Kolen, (1992) "Economic Lot Sizing: An O(n log n) Algorithm that Runs in Linear Time in the Wagner-Whitin Case" *Operations Research*, volume 40, supplement number 1, January-February, pages s145-s156.

Wagner, H.M. (1975) *Principles of Operations Research, with Applications to Managerial Decisions*, 2nd edition, Prentice-Hall, Inc. Englewood Cliffs, New Jersey.

S.D. Wu and M.J. Meixell, "Relating Demand Behavior and Production Policies in the Manufacturing Supply Chain," *IMSE Technical Report 98T-007,* Lehigh University, (http://www.lehigh.edu/~sdw1/meixell.PDF)

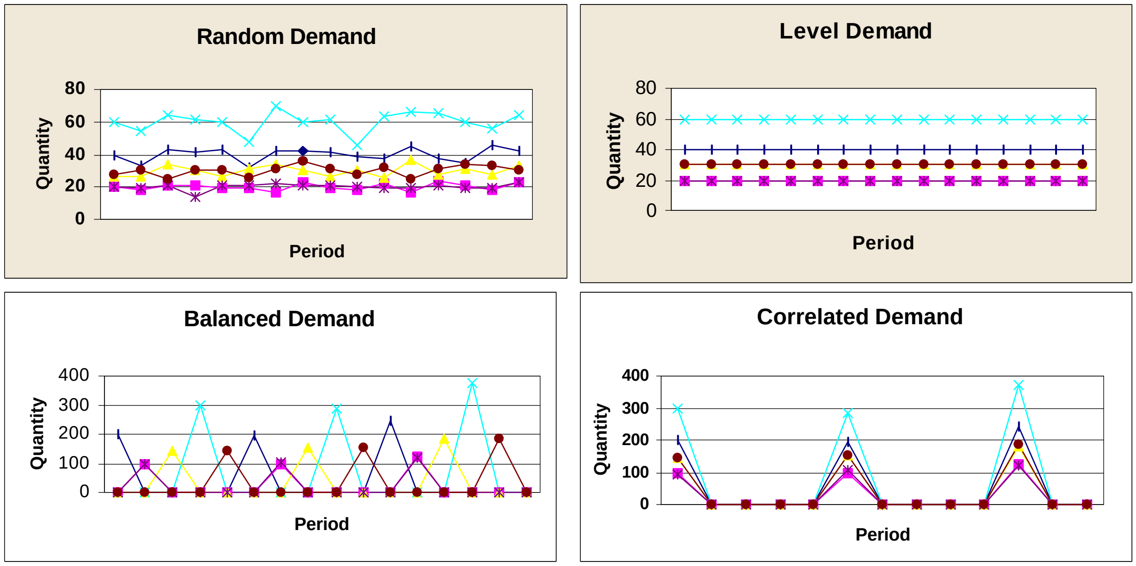

Figure 7. End Item Demand Patterns

Figure 2: Product Structures