SAFETY STOCK SAVINGS FROM FLEXIBLE ORDERING: A SIMULATION STUDY

David J. Parsons

Simulation Dynamics, Maryville, Tennessee

parsons@simulationdynamics.com

ABSTRACT

Premise: contracts for products with long lead times and unpredictable demand can benefit the buyer by providing for revision of order quantities after the order is initially placed.

'Flex Ordering' allows the buyer to place an order with substantial lead time, say 22 weeks, but then revise the quantity of the order within windows and percentages specified in the flex ordering contract. This benefits the buyer by reducing the safety stock that would be required to protect against higher than expected demand. It benefits the supplier by providing a degree of predictability of production requirements.

A simulation study was conducted to determine the safety stock savings that result from flex ordering in a benchmark case; and to determine the sensitivity of this result to several context and flex ordering design factors.

1. POTENTIAL SAVINGS IS DICTATED BY CONTEXT

1. Scenario

The context of this study is a U. S. distribution center (DC) supplied by an overseas plant. The DC places orders at the beginning of each week, with delivery at the beginning of the week 22 weeks hence. Customer orders are received by the DC after the DC's order to the plant is placed. Customer orders are filled immediately from inventory. Customer orders average 700 units per week, with benchmark coefficient of variation of 0.4. Demand is normally distributed with a minimum of zero.

The question posed by this study is: what reduction of safety stock inventory will result from use of flex ordering.

Benchmark case: a benchmark case is posed with the following parameters:

Mean demand 700 units per week
Demand variability: 0.4 coefficient of variation.
Order lead time: 22 weeks

1.2. Measurement of order filling performance

Each week, if there is not sufficient stock in inventory to fill the full customer demand for the week, the week is counted as an unfilled order.

Since the product in question has a life cycle of two to three years, periods of 100 weeks are studied.

The benchmark performance goal is to achieve 99% orders filled within 100 week periods 95% of the time. 99% is the Goal (G); 95% is the target Achievement Rate (AR).

1.3. Safety Stock Required without flex ordering

Calculation of savings will be based on reduction of safety stock required without flex ordering.

The plot below indicates the safety stock required to achieve the customer order filling goal with order lead times from one to twenty two weeks. With the benchmark lead time of 22 weeks, safety stock of 3,687 units is required to achieve benchmark goals.

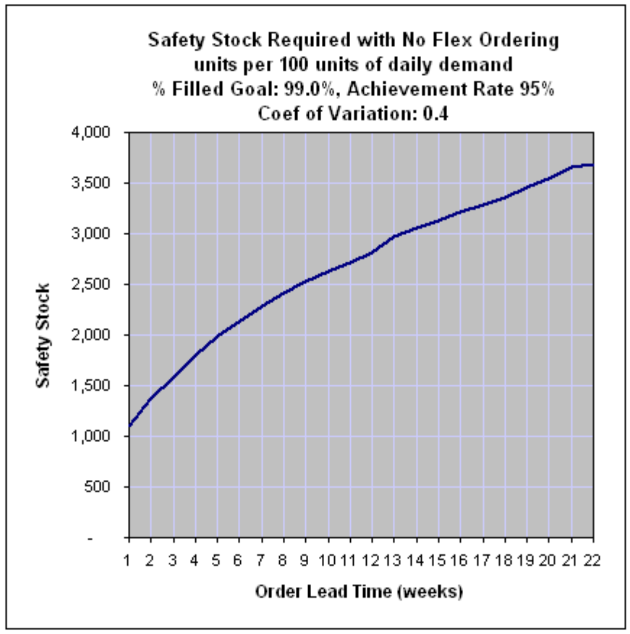

Figure 1: Safety stock as a function of lead time

A model of 10,000 weeks with random variations in weekly demand was repeated 11 times for each flex scenario. This produced a matrix of results that counted the number of out of stock weeks that would have occurred at a range of safety stock levels at 100 unit increments during 1,100 periods of 100 weeks. The safety stock required to reach a specified goal at a specified achievement rate was interpolated from this matrix of results.

1.4. Potential Savings

Potential savings from flex ordering is determined by considering the extreme case where unlimited flexibility is allowed. For example, as shown in Figure 2, the safety stock required for a nine week lead time is about 2500 units. This would be the same as a scenario where unlimited flexibility is offered in the 9th week.

The savings resulting from a unlimited flexibility in the 9th week can be determined by subtracting the safety stock at 9 weeks from the safety stock at 22 weeks. The shaded area indicates savings that would result from unlimited flexibility starting in other weeks.

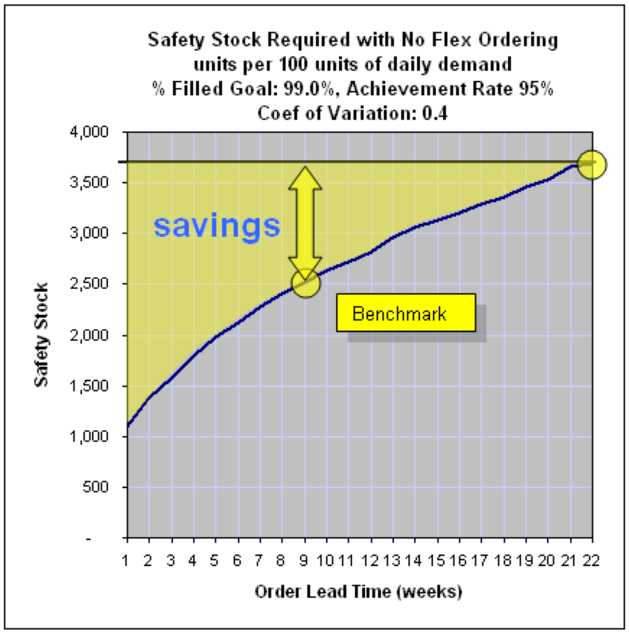

Figure 2: Potential savings

Figure 3 shows potential safety stock savings as a separate curve.

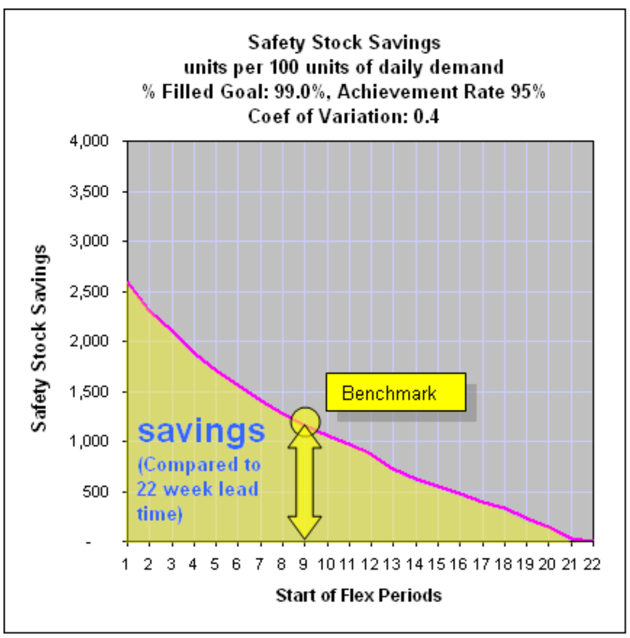

Figure 3: Potential Savings

The study of alternative flex ordering designs can now be framed in terms of the percentage of potential savings that each strategy realizes.

2. ACHIEVEMENT OF POTENTIAL SAVINGS IS DICTATED BY FLEX ORDERING DESIGN

1. Flex Ordering Contract Design

Flex Periods: The time from when an order is placed to when it is delivered is divided into a number of flex periods, as shown in the table in Figure 4.

The flex ordering contract specifies for each period the percentage by which an order can be increased or decreased during that period.

The starting point for increasing or decreasing the order is its quantity when it enters each flex period. The percentages are based on the initial amount of the order.

In the benchmark flex ordering contract, flex percentages are the same in each period. This need not be the case.

In this example, an initial order of 700 units could be increase three times by 10% to a maximum of 910 units.

Each week, two order amounts are determined. First, the forecasted demand for the week 22 weeks in the future is projected. Second, the difference between the forecasted and actual demand for the week just completed is calculated. This forecast error might be a positive or negative number. With no flex, the forecasted demand for week 22 and the forecast error from the prior week would be summed to make up the order for delivery in 22 weeks. With flex ordering, the forecast error from the prior week, positive or negative, would be used to adjust the available flex period nearest to delivery.

Benchmark Flex Ordering Contract

Period	Weeks until delivery: First	Weeks until delivery: Last	Percent increase	Percent Decrease
1	8	0	0	0
2	12	9	10%	10%
3	16	13	10%	10%
4	21	17	10%	10%
5	Order placed with 22 week lead time			

Figure 4

1. Sensitivity to Start of Flexibility

Potential savings is dictated in part by the first week in which flexibility is available (as shown in Figure 3). In all of the examples discussed below, results are shown for start values of week 1 to week 20.

2.2. Sensitivity to Flex Percent

The plot shown in Figure 5 shows the percent of potential savings that result from our benchmark flex ordering design, starting in week 9, with four week periods of 10% flex. About 77% of the potential savings is realized. In addition, the percentage of potential savings is shown for the same flex design but with 5% and 20% flex

The plot in Figure 6 shows the actual safety stock savings from these three flex designs.

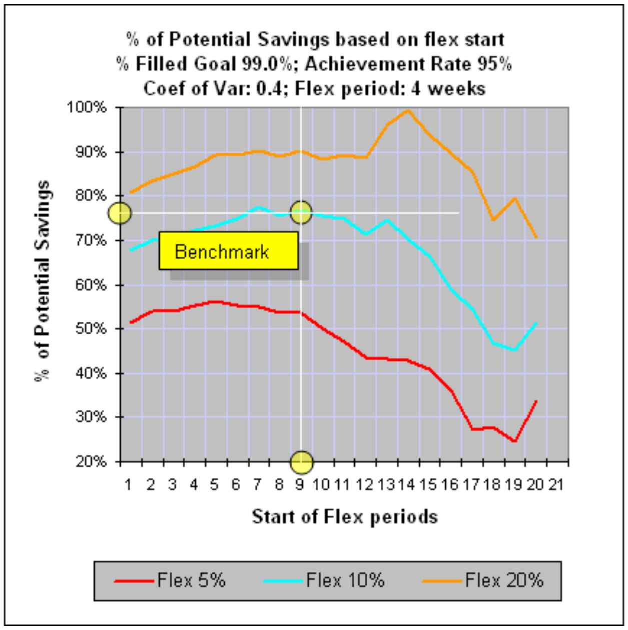

Figure 5: Sensitivity to Flex Percent

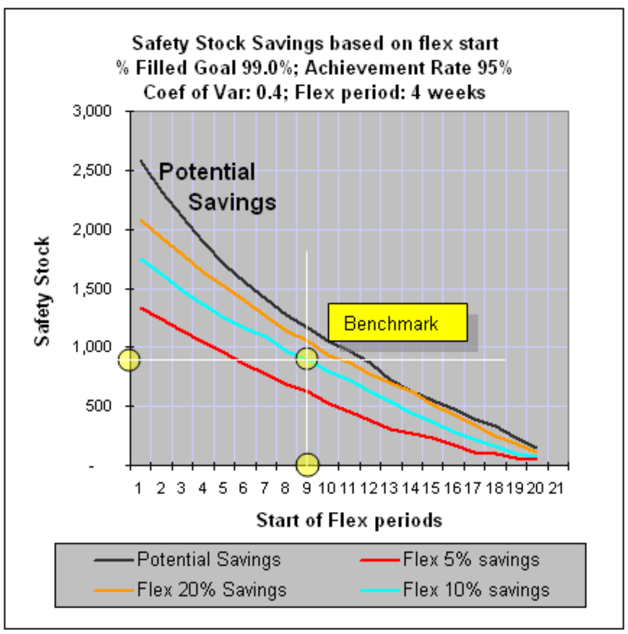

Figure 6: Sensitivity to Flex Percent

4. Sensitivity to Length of Flex Periods

Figure 7 compares our benchmark case to flex designs with different lengths of flex period. In each of these cases, the flex increase and decrease amounts are 10%. Case FP1 has flex periods of one week. This means that flex amounts could be increased or decreased each week. For example, with FP1, an order of 700 units could be increased thirteen times by week 9 to a quantity of 1610 units. FP13 is a design with only one flex period: each order can only be increased or decreased by a total of 10%.

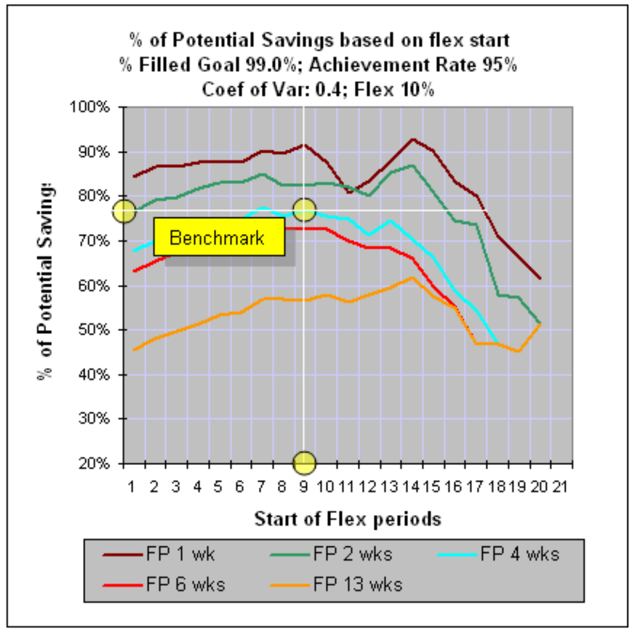

Figure 7: Sensitivity to Length of Flex Periods

Figure 8 gives the actual savings from each of these flex period scenarios. Clearly, the shorter the flex period, the greater the percent of potential savings realized.

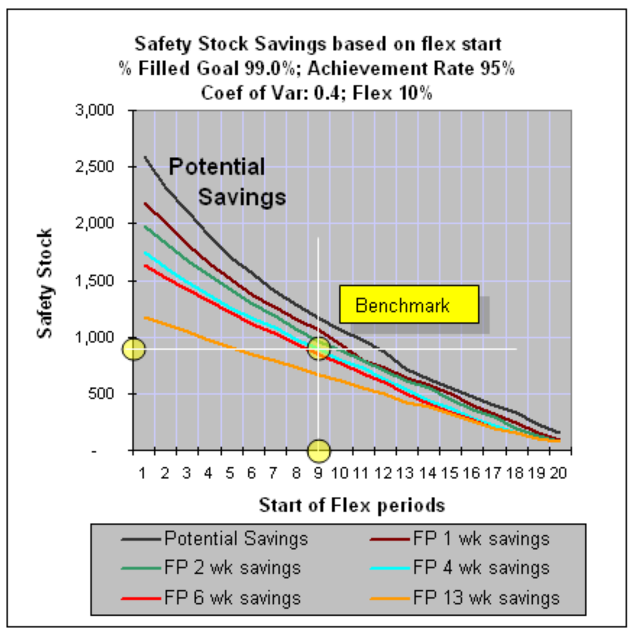

Figure 8: Sensitivity to Start of Flex Periods

5. Mixed Flex Rates

Figures 9 and 10 show the sensitivity of savings to changes in the flex rate of the first and third periods of the benchmark flex contract design. The parameters of these contracts are given in the tables below. The results shown here indicate that safety stock savings are very sensitive to changes in flex rates in the first flex period, while not sensitive at all to changes in the third period. Contract #3, with increased flexibility in the third period has essentially the same performance as the benchmark case. Contract #2 with a lower first period flexibility has substantially reduced performance. contract #4, which differs from contract #2 by having an increased third period, has the same performance as contract #2.

Flex Ordering Contract #2				
Period	Weeks until delivery		Percent increase	Percent Decrease
	First	Last		
1	8	0	0	0
2	12	9	5%	5%
3	16	13	10%	10%
4	21	17	10%	10%
5	Order placed with 22 week lead time			

Flex Ordering Contract #3				
Period	Weeks until delivery		Percent increase	Percent Decrease
	First	Last		
1	8	0	0	0
2	12	9	10%	10%
3	16	13	10%	10%
4	21	17	20%	20%
5	Order placed with 22 week lead time			

Flex Ordering Contract #4				
Period	Weeks until delivery		Percent increase	Percent Decrease
	First	Last		
1	8	0	0	0
2	12	9	5%	5%
3	16	13	10%	10%
4	21	17	20%	20%
5	Order placed with 22 week lead time			

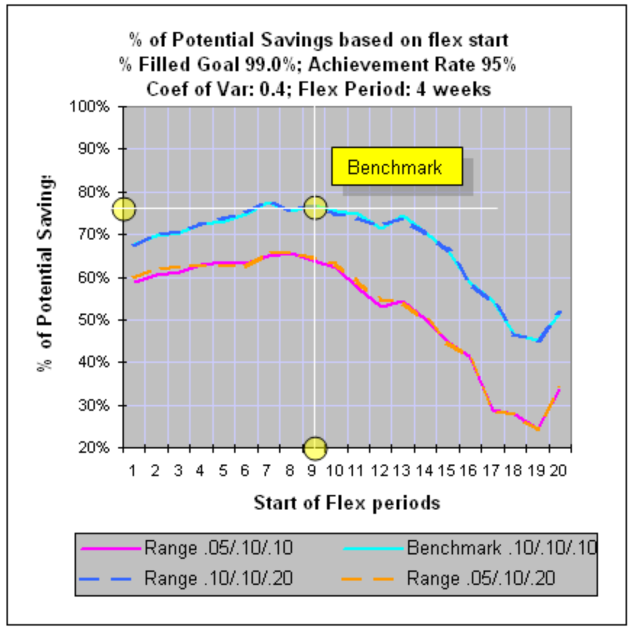

Figure 9: Results from mixed flex rates

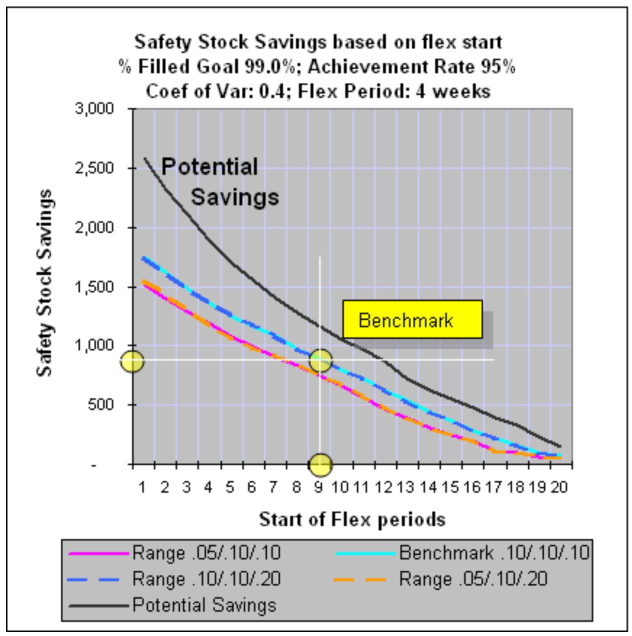

Figure 10: Results from mixed flex rates

2.6. Flex Increase vs. Flex Decrease

Figure 11 shows the savings from flex contracts where only the increase or only the decrease are allowed. With flex starting in the 9th week, flex increase contributes about 45% of the savings potential, while flex decrease contributes 37%. It is not obvious that the two components should 'add up' to the performance of the combined flex contract, yet that appears to be the case with our benchmark example. Figure 12 gives the actual safety stock savings from this experiment.

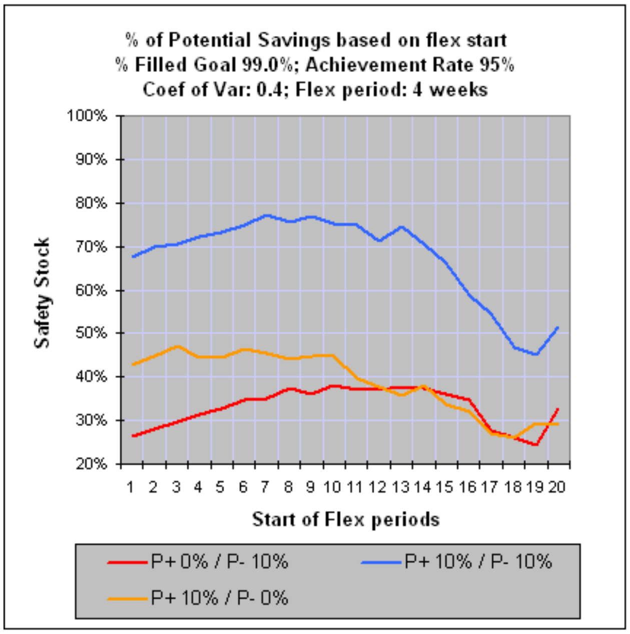

Figure 11: Percent of potential from Flex Components

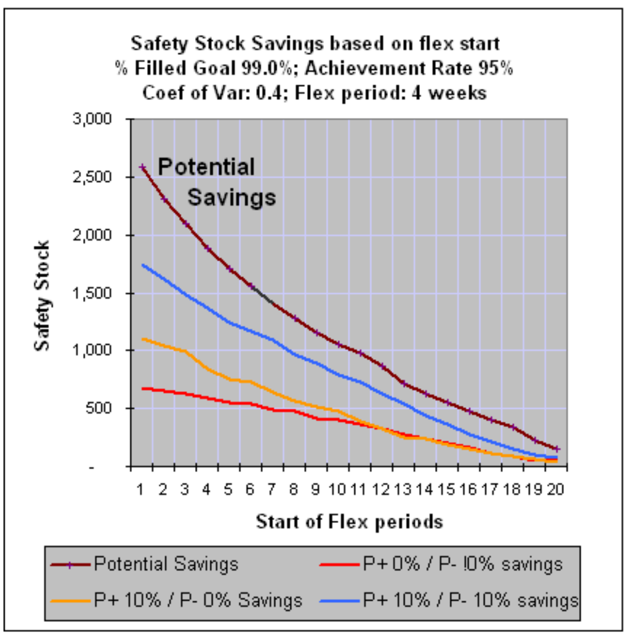

Figure 12: Safety stock savings from Flex Components

3. IMPACT OF HIGHER DEMAND VARIABILITY
Safety stock requirements increase as demand variability increases as shown in Figure 13. Figure 14 shows the corresponding potential savings from these three levels of demand variability.

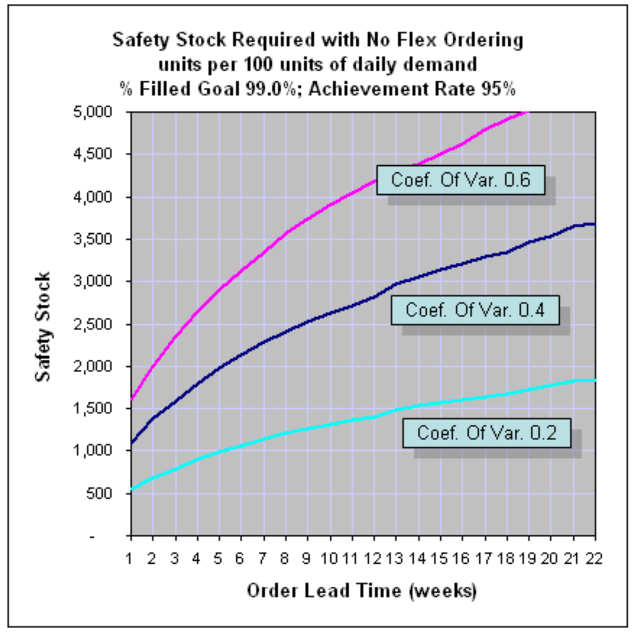

Figure 13: Safety stock requirements

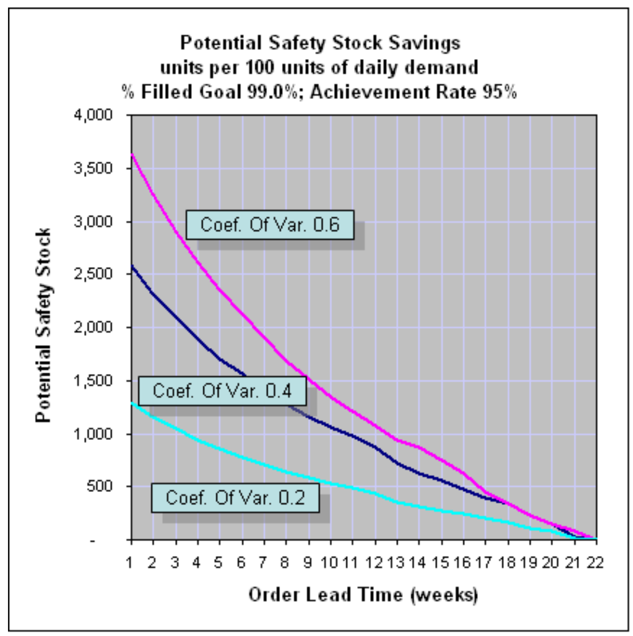

Figure 14: Potential Safety Stock Savings

3.1 Demand Variability Effect on Percent of Flex
When demand variability is increased to a coefficient of variation of 0.6, the percent of the potential savings from flex ordering strategies goes down. Figure 15 shows the percent of potential savings from 5%, 10%, and 20% flex at demand variability of 0.4 and 0.6. The reduction is higher for lower percentage flex rules.

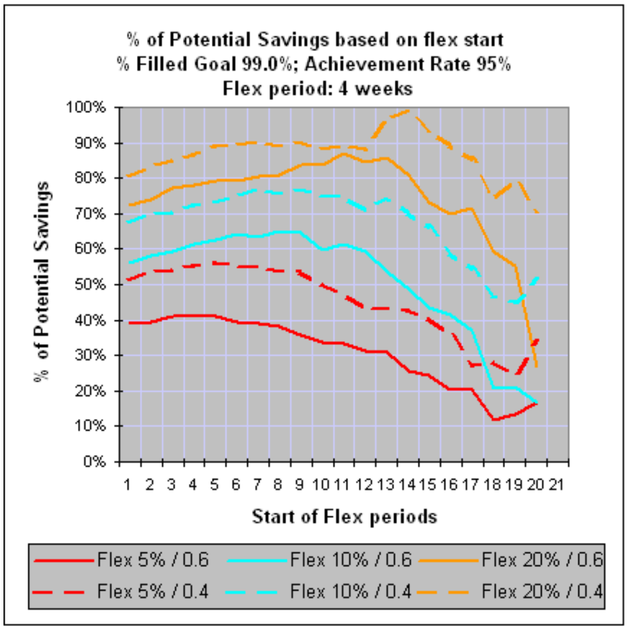

Figure 15 Demand Variability / Flex Percent

Figure 16 shows that the actual savings from 5% flex does not increase in the 0.6 scenario, even though there is higher potential savings. With the 20% flex design, there is significant increase in safety stock savings as demand variability increases.

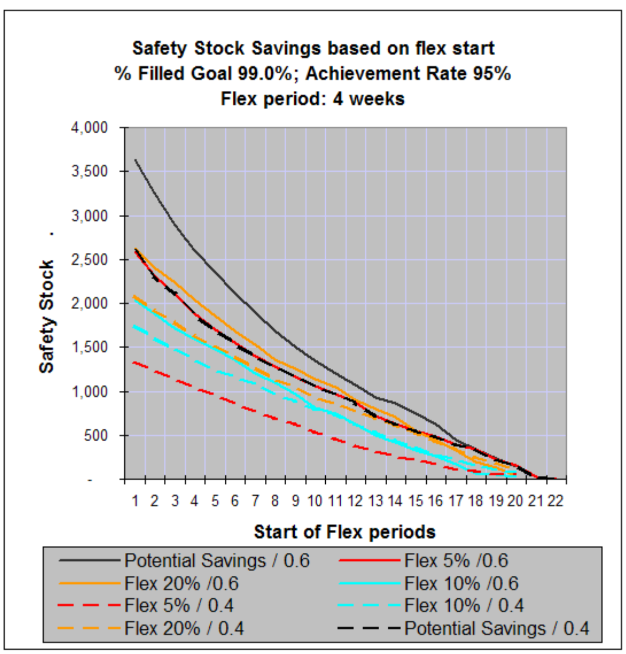

Figure 16: Demand Variability / Flex Percent

3.2. Demand Variability effect on Flex Period Length

Figures 17 and 18 show flex designs with flex periods of 1, 2, 4, 6, and 13 weeks, with design variability of 0.6. These results can be compared to Figures 7 and 8.

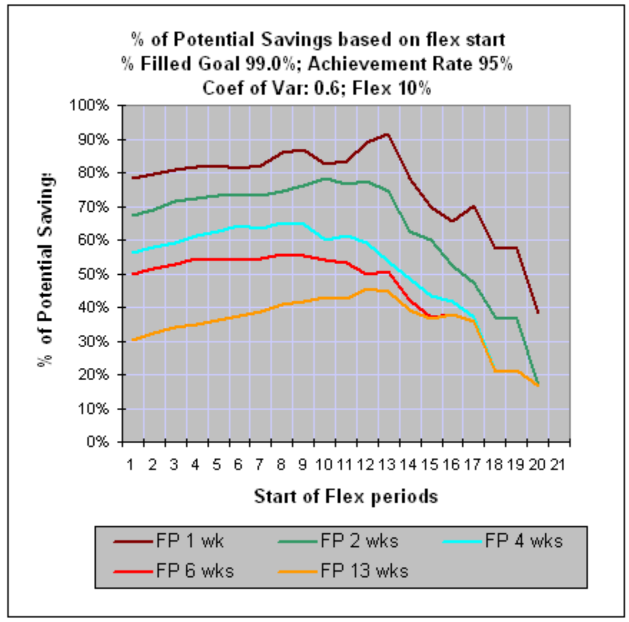

Figure 17: Demand Variability / Flex Period Length

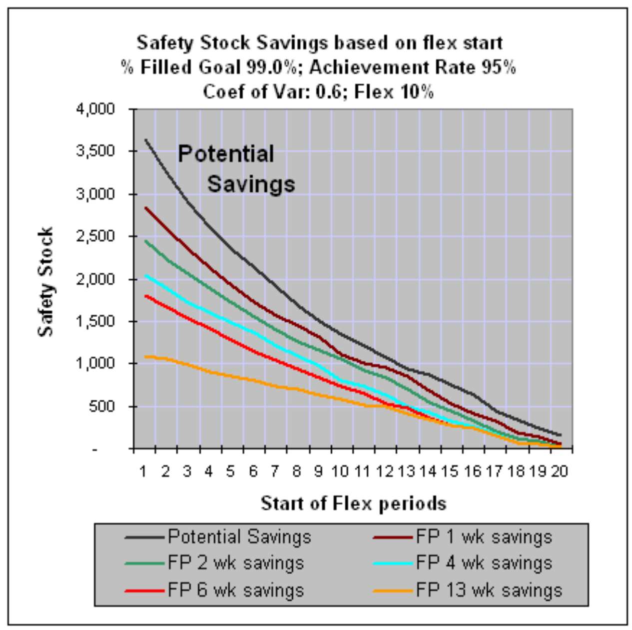

Figure 18: Demand Variability / Flex Period Length

The table in Figure 19 compares percentage and absolute savings with different flex period lengths at 0.4 and 0.6 demand variability. In these scenarios, flexibility starts in the 9th week. Note that the savings with periods of six weeks and longer are actually less with the higher coefficient of variation.

Savings as function of period length and demand variability (Flexibility starting in 9th week)

	Length of Flex Periods (weeks)

	1	2	4	6	13
Coef. of Variation 0.4					
% of potential	92%	82%	77%	72%	57%
Savings	1,064	956	893	841	658
Coef. of Variation 0.6					
% of potential	87%	76%	64%	55%	42%
Savings	1,308	1,150	971	835	630

Figure 19: Savings as function of period length and demand variability

4. SENSITIVITY OF SAFETY STOCK REQUIREMENTS AND SAVINGS TO THE PERFORMANCE GOAL

In this study, the benchmark performance goal is to achieve 99% orders filled within 100 week periods 95% of the time. There has been no tradeoff between the benefits of performance and the cost of safety stock. This section provides safety stock results for the benchmark flex design parameters under varying goal and achievement rate assumptions. In each scenario, flexibility starts in the 9th week, flex periods are 4 weeks long with flexibility at 10% increase and decrease in each period. These results are shown in Figure 20.

As one would expect, the more aggressive the performance goals, the higher the safety stock required. Potential savings under our benchmark flex contract also increase with higher performance goals. In all scenarios with the 95% goal, potential savings are a higher fraction of the 22 week safety stock than in the benchmark case. It is a curiosity that achieving a 95% goal 99% of the time requires significantly more safety stock than our benchmark – achieving a 99% goal 95% of the time.

Actual savings from the benchmark flex design vary roughly in proportion to the 22 week safety stock. This indicates that with the 95% goal, there is high potential savings, but the benchmark flex design is not able to reap them. Actual safety stock, and therefore, safety stock savings are roughly proportional to the 22 week safety stock.

Customer Percent Orders Filled Scenarios		Benchmark Flex Design: Flex start: week 9 Flex Periods: 4 weeks Increase & decrease: 10%			
Goal	Achievement Rate	22 week safety stock	Max Potential Savings	Actual Safety Stock	Actual Savings
99%	99%	4,300	1,340	3,400	956
	95%*	3,687	1,160	2,794	893
	90%	3,367	1,031	2,553	814
98%	99%	4,180	1,413	3,320	862
	95%	3,562	1,172	2,683	878
	90%	3,196	999	2,457	739
95%	99%	3,920	1,540	3,000	920
	95%	3,227	1,127	2,425	802
	90%	2,957	1,057	2,221	736
* Benchmark					

Figure 20: Safety stock variation with goal and achievement rate

Assuming that the cost of the benchmark flex contract is a constant, we can infer from this table the relative cost in units of safety stock of the nine order filling scenarios studied. For example, our benchmark contract requires 573 more units of safety stock than the 95% / 90% scenario (2,794 – 2,221). With no flex ordering, the 99% / 95% scenario would require 730 units more than the 95% / 90% (3,687 – 2,957).

5. CONCLUSIONS

The goal of this study was to gain insight into the sensitivity of the benefits of flexible ordering to several exogenous and endogenous factors. We have determined that the potential for savings is a function of:

- Order lead time
- Customer order filling goals
- Demand variability
- Start of flex periods

The amount of this potential that is realized is a function of:

- Length of flex periods (number of periods)
  - Shorter flex periods result in greater savings
- Percent of flex increase and decrease
  - Increase and decrease both contribute to savings
  - Flexibility in early periods is far more important than in later periods
  - Increase flexibility has greater benefits as demand variability increases

The goal of a negotiated flex ordering arrangement would be to find the lowest total cost to buyer and supplier, with a reasonable division of the savings over conventional fixed contracts. This study views flex ordering from the buyer's perspective. In such a negotiation, if the buyer has the ability to determine the savings from alternative proposed contracts, and has the supplier's proposed cost for such contracts, an informed decision can be made as to which contract has the best combination of cost and performance for the buyer. In addition to highlighting the factors that contribute to savings from flex ordering, this paper suggests that simulating inventory behavior under alternative flex ordering schemes would provide valuable support in such a negotiation.

SIMULATION METHODOLOGY

The overall study or which these results were a part was modeled in Extend V.5.4, a product of Image That! Inc. The model foundation was Supply Chain Builder, a toolkit developed by Simulation Dynamics for a wide range of supply chain and logistics models. A parallel model was developed in Excel to generate repetitions for the analysis presented in this paper.

AUTHOR'S BIOGRAPHY

DAVID J. PARSONS is a principal of Simulation Dynamics. He is responsible for development of SDI's Supply Chain Builder, a supply chain simulation and analysis tool. His experience with simulation began in 1965 with experiments in the use of natural selection algorithms to evolve architectural designs. During the 1980's he designed, built and operated several dairy-processing plants using simulation of key systems as an integral tool for design, value engineering, and trouble- shooting. Mr. Parsons received a B. A. from Harvard College and a Master of Architecture degree from the Harvard Graduate School of Design.