

ORLANDO

# Managing the Impact of Supplier Lead Time and Quantity Variation

**John Estep, CFPIM**  
**President**  
**E/Step Software Inc.**

**Nathan Boyd**  
**President**  
**Forparco**

APICS is not responsible for statements or opinions expressed by individuals in its publications or at its meetings. The views expressed are solely those of the individual and are not necessarily endorsed by APICS. APICS programs are noncommercial forums. Speakers have agreed to refrain from the use of brand names and specific product endorsement whenever possible. Speakers understand that the association's podium cannot be used as a place for direct promotion of the presenter's product, services, or for monetary self interest.



## **Introduction**

To meet your customers' service expectations you must manage your suppliers' (both internal and external) lead time and quantity variation. The first step is to work with your suppliers to reduce variation, but often variation can not be eliminated. As a result, demand safety stocks are often expected to cover both demand *and* supply variation, but this approach fails since demand safety stocks include the effects of forecast errors *but not supplier errors*. Frequently, supplier variation is the ignored culprit in supply chain stockouts. We need techniques to measure—and thus manage—supplier variation.

## **The Application**

An apparel manufacturer originally hired us to help them solve some service and inventory problems on the demand side. The client's initial inventory policy was a 12 week fixed time supply: 9 weeks to cover the forecast errors, and 3 weeks to cover the supplier variation. Note this means they carried 12 weeks on every SKU, whether they needed it or not.

As usual, the fixed time supply inventory policy was the culprit; a statistical safety stock—set by SKU—provided the solution. On the demand side this cut the inventory by a third, from 9 weeks (on average) to just 6 weeks. Note here “6 weeks” is not a fixed time supply, but just the aggregate of the statistically calculated safety stocks. So the client began to implement the new finished goods safety stock targets by slowly “draining the pond” to get down to the 9 week aggregate inventory (i.e., the new 6 week safety stock to cover the forecast errors plus the 3 week fixed time supply safety stock to cover supplier variation).

Eventually unexpected stockouts began to occur. The client assumed that the new safety stock wasn't delivering the service target, because the stockouts would have been covered by the original 12 week fixed time supply target. The problem was actually due to supplier variation, however. How was supplier variation confirmed as the source of the problem? First, we evaluated the individual demand transactions against the median and found all were within 3 standard deviations.

Next we compared the aggregate demand for each forecast period and found it was less than or equal to the expected inventory (working stock plus safety stock). We did, however, find the actual stock on hand was less than the expected inventory. This proved the 3 week fixed time supply buffer against supplier variation was not sufficient. The excess finished goods safety stock that we had earlier eliminated had actually been covering up this client's supplier variation problem. Our client's problem revealed the need to manage two issues: lead time variation and quantity variation.

Lead time variation, which is the difference between a supplier's lead time policy (the nominal lead time) and the actual lead time, is important because of lead time's role in the calculation of safety stock. Likewise, the quantity variation, which is the quantity ordered versus the quantity received, is important because it affects the size of the working stock, and working stock contributes to our ability to meet our service target.

The client resolved the problem temporarily by restoring the additional 3 weeks of fixed time supply safety stock, but since he did not want to give up the savings achieved, engaged us to develop a statistical solution to calculating safety stock to cover supplier variation.

We will consider, separately, variations in both lead time and quantity.

### **Supplier Lead Time Variation**

Measuring supplier lead time variation (SLTV) provides the abilities to:

1. Identify the largest opportunities to reduce SLTV.
2. Determine the true cost of a product among competing suppliers.
3. Evaluate the tradeoff between the inventory required to cover SLTV and service and/or expediting levels.

How often are your suppliers supposed to deliver product in 10 weeks, but actually take 12 weeks (or 9 weeks, or 15 weeks) before the product arrives? If you aren't accounting for this variation in your supply chain then you should.

So what is the lead time of a supplier with this degree of variation? If you use the longest lead time, your service will be excellent, but your turns will be poor. If you use the shortest lead time, your turns will be excellent, but your service will be poor. Is there a way to obtain excellent turns and service simultaneously? The answer is to measure the variation between the nominal and actual lead times and then convert it from units of lead time to pieces. With this information you will be able to manage your supplier's lead time variation.

### **Data Requirements**

To measure lead time variation you need the following data:

1. Part number
2. Unique supplier transaction code (such as a purchase order number)
3. Supplier identification code
4. Date the supplier transaction was generated
5. Date the supplier transaction was scheduled to arrive
6. Date the supplier transaction actually arrived
7. (optional) Aggregate grouping codes, such as a product line, style, etc.

The first four are necessary to keep all records unique, to eliminate duplication from partial receipts or multiple records with the same part number with the same date. The dates allow us to compute

the number of days late for each transaction, so we can compute the nominal and actual lead times. The optional grouping codes allow aggregation by common characteristics.

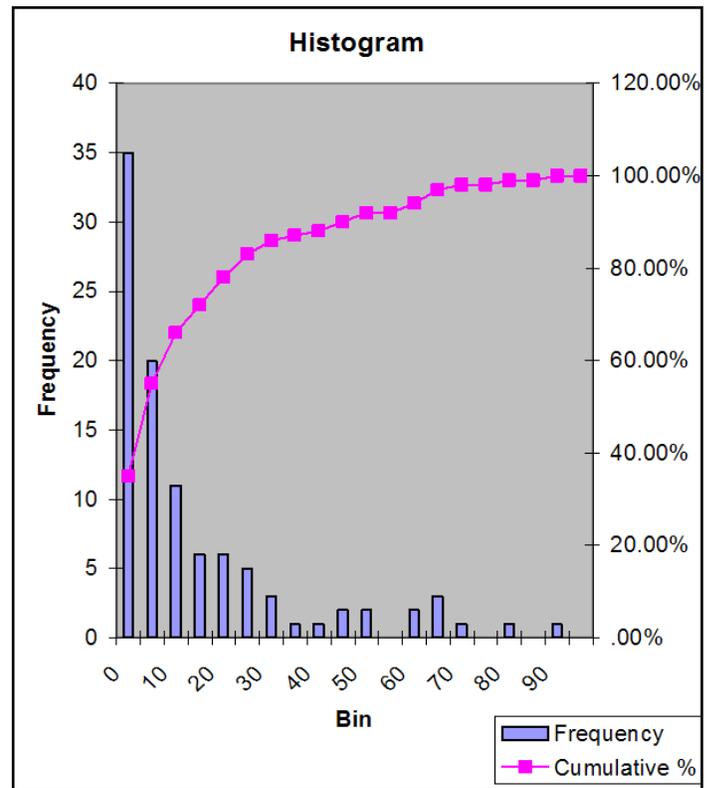
To determine the number of days late, subtract the nominal lead time from the actual lead time. Early orders are assigned a zero value for days late since—although they increase the inventory—they do not increase stockouts, and this analysis was confined to late orders. (See sample orders at right.)

Date	Sched. Lead Time	Actual Lead Time	Days Late
9/11/02	46	50	4
9/11/02	46	61	15
9/11/02	46	50	4
11/14/02	46	81	35
11/20/02	46	104	58
11/20/02	46	35	0
11/20/02	46	89	43
11/25/02	46	63	17
11/25/02	46	74	28
12/4/02	46	65	19
12/4/02	46	62	16

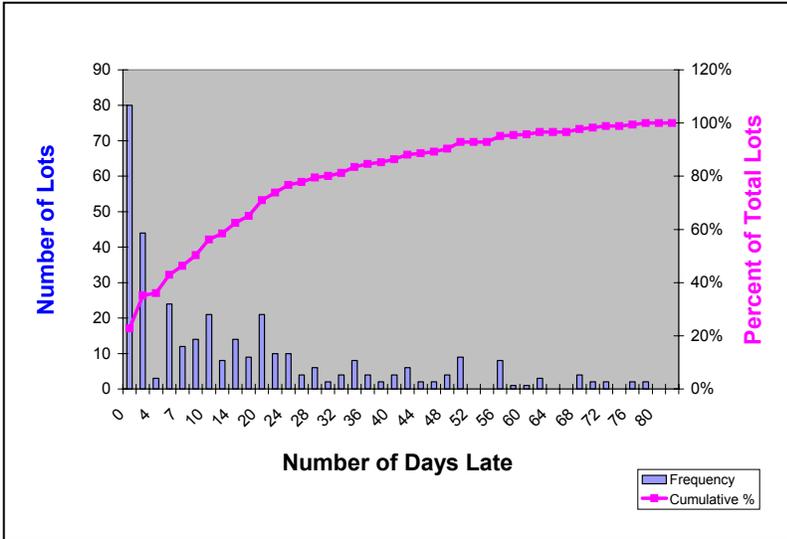
**Determine Distribution Type**

To determine the distribution of the data, we used the Histogram tool in Microsoft Excel (Tools/Data Analysis/Histogram). We pointed it to the list of days late and also specified the bin size (range of days late to include in each group). The tool counts the

Days Late	Bin Size 5 Series	Bin Size	Frequency	Cumulative Percent
86	0	0	35	35.00%
77	5	5	20	55.00%
66	10	10	11	66.00%
31	15	15	6	72.00%
15	20	20	6	78.00%
9	25	25	5	83.00%
36	30	30	3	86.00%
23	35	35	1	87.00%
62	40	40	1	88.00%
62	45	45	2	90.00%
15	50	50	2	92.00%
0	55	55	0	92.00%
0	60	60	2	94.00%
15	65	65	3	97.00%
3	70	70	1	98.00%
1	75	75	0	98.00%
1	80	80	1	99.00%
1	85	85	0	99.00%
9	90	90	1	100.00%
5	95	95	0	100.00%
0	100	100	0	100.00%
9	105	105	1	100.00%
21	110	110	0	100.00%
21	115	115	0	100.00%
5	120	120	0	100.00%
<b>Total</b>			<b>100</b>	



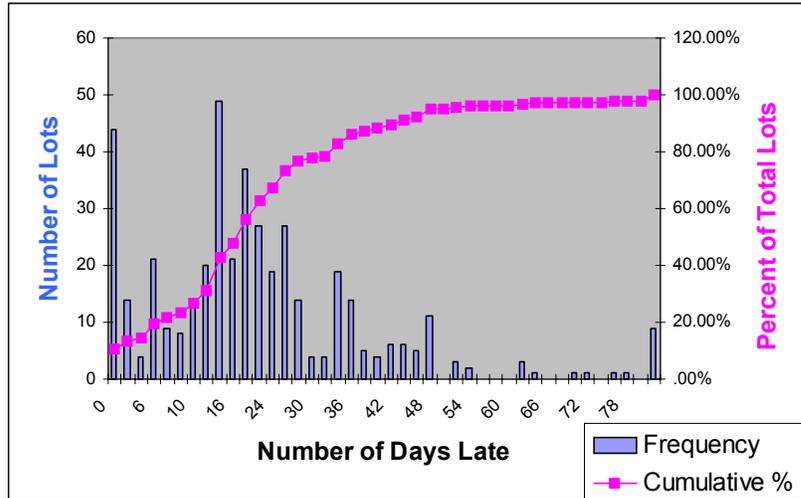
frequency with which replenishment orders in each group of days late occurs, arranges them in order, and computes the cumulative percentages. It also displays a histogram of the frequency distribution as well as a line graph of the cumulative probability. We then visually confirmed the type of distribution observed, comparing it to known distributions.



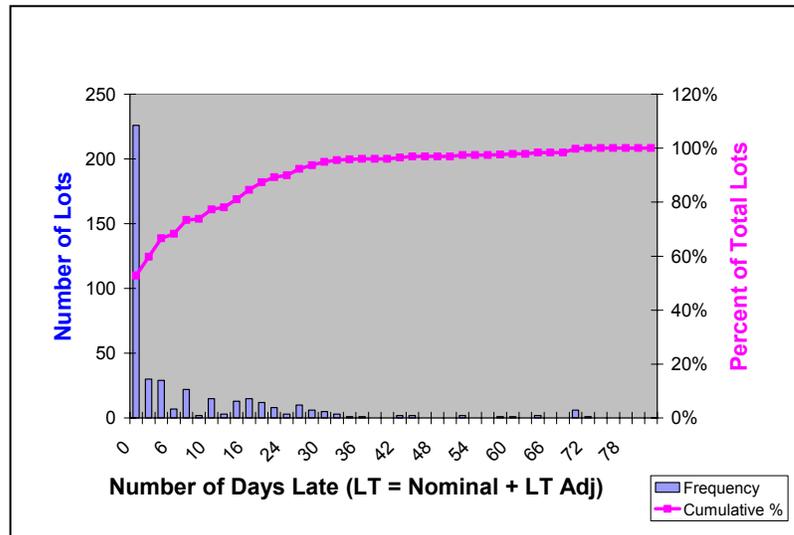
Initially we determined that 30% of the SKUs had days late data distributions of the Normal (Gaussian) type; the other 70% matched the Exponential distribution.

Here is an example of an exponential distribution. The horizontal axis is the number of days late. The vertical axis is the number of replenishment lots. The shape of the curve tells us that it is exponential. Note that there is a high probability of low errors and a low probability of high errors. Most orders are delivered around the due date, while some orders are delivered late.

In contrast, here is an example of days late which are normally distributed. The median for this SKU is 19 days late. A normal distribution here essentially means that the nominal lead time is wrong. We can correct this problem by increasing the nominal lead time by the median days late. We then save the median value for future lead time calculations, recalculating the new days late values based on the old nominal plus the median value.



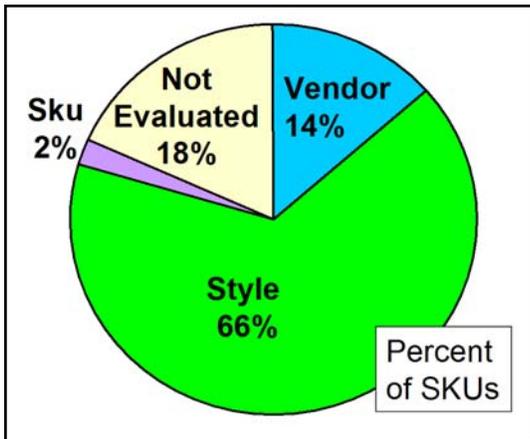
After adjusting the nominal lead time to be the old nominal plus the median, and then recomputing the number of days late based on this new “nominal” lead time, the result is an exponential distribution. Thus all normal distributions are converted to exponential distributions



**Statistically Unusable Data**

For some SKUs we were unable to determine the distribution of the days late data. If there were no late lots then there was no lead time error—in which event no additional action was required, as the nominal lead time was accurate. Sometimes the data was unusable simply because there were too few replenishment lots for the particular SKU—typically because there was little customer demand. In other cases it was because the vendor was doing a good job, and very few replenishment lots were arriving late. In these circumstances the distribution was either uniform or had no recognizable distribution.

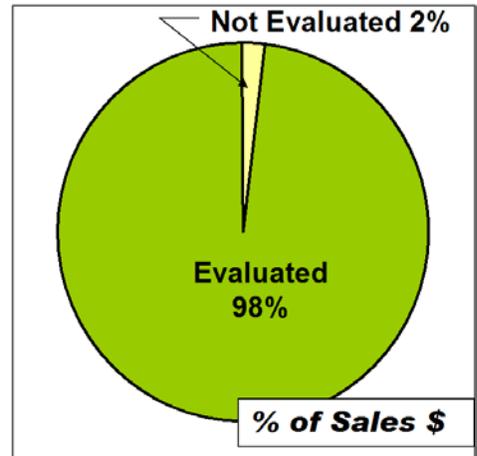
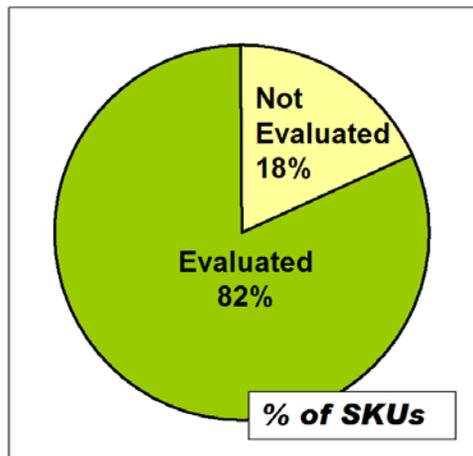
Rather than ignore this data, we aggregated the data of these products to a higher group level (e.g., from the SKU/Size level, to the Style level, and if necessary, to the vendor level). We then repeated the above process to identify the groupings whose lead time variation data was statistically usable.



We repeated the aggregation and evaluation process until either all data was a member of a statistically usable distribution or all aggregate groupings had been exhausted. If any data could not be grouped to form a statistically usable distribution, we dropped it from consideration in the SLTV analysis. In our study this was usually because of very few orders, which in turn was because of very low volumes. As a result, omitting these items had negligible impact on the total result. As the chart at left shows, only 2% of the SKUs had usable distributions at the SKU level; however 66% of the SKUs had usable distributions after aggregation to the Style level. Fourteen percent required aggregation to the

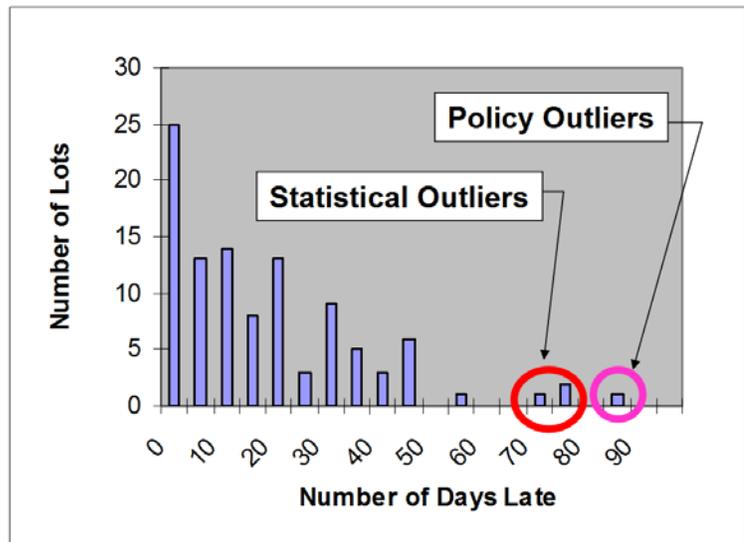
vendor level to obtain a usable distribution, and for 18% of the SKUs we could not determine any usable distribution.

Even though 18% of the SKUs had unusable distributions, the particular SKUs amounted only to 2% of the sales dollars, so we were able to ignore them without any significant effect on the results of this study.



## Outliers

We identified and eliminated outliers in the SLTV data for both policy and statistical reasons. As a matter of policy, the client did not want to include in the analysis any replenishments that were more than 90 days late; so we eliminated these. In addition we identified statistically any replenishments falling outside a 99 percent confidence limit and eliminated this data from consideration as well. In these cases the client would have to cover SLTV with expediting or pass the variation through to its customers as service misses. An example of a SKU with both statistical and policy outliers is shown above.



Given that all the (eventual) distributions were exponential, the precise logic used for statistical outliers is to exclude all replenishments (or groups) where the number of days late exceeds:

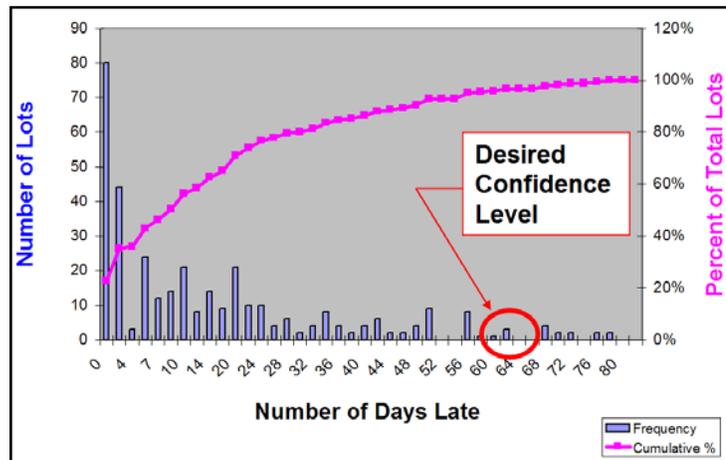
$$\text{AverageDaysLate} \times \text{ABS}(\ln(1 - 0.99))$$

where  $\ln(x)$  = natural log ( $x$ ), and  $\text{ABS}(x)$  = absolute value ( $x$ ).

To use any other statistical limit, substitute the desired confidence limit for 0.99 in the above expression. Thus for a 95% outlier limit, use  $(1-0.95)$ , etc.

## Expected Days Late

Now that the data has been prepared, it is time to calculate the Expected Days Late for each SKU or aggregate group. We used a 99% confidence limit for this calculation. That means that—excluding outliers—there is a 99% probability that the actual days late will not exceed the Expected Days Late. The example at right shows a SKU for which the desired 99% confidence limit requires adding 64 days to the nominal lead time in order to receive 99% of the replenishment lots “on time”, i.e., within nominal lead time plus the expected days late.



While the exponential distribution uses the average as its defining parameter, we used an *exponentially smoothed* average instead, to reflect explicitly any trend in the days late; i.e., in case the number of days late is improving or getting worse. (If it is staying the same then it makes no difference whether we use an average or an exponentially smoothed average.) We used a smoothing rate of 10%. This is quite useful because as the client works with the vendors to reduce the lead time variation, the additional safety stock carried because of SLTV will automatically reduce.

The expression for the Expected Days Late is:

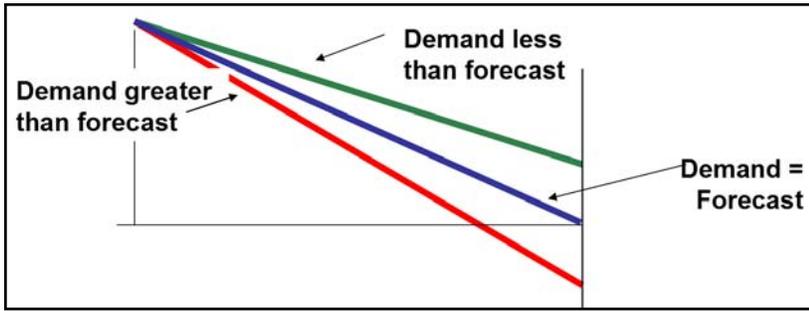
$$\text{ExpectedDaysLate} = \text{ExponentiallySmoothedAverageDaysLate} \times \text{ABS}(\ln(1 - 0.99))$$

where  $\ln(x)$  = natural log (x), and  $\text{ABS}(x)$  = absolute value (x).

To use an alternative confidence limit, substitute it for 0.99 in the above expression.

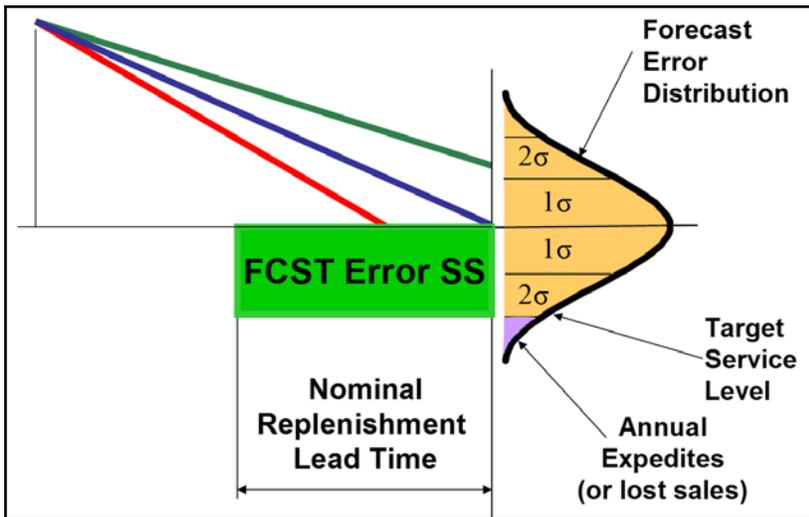
### SLTV Safety Stock

Recall that the purpose of safety stock is to provide a cushion in the event that demand exceeds the



forecast during the replenishment lead time. Let's build a picture of this, starting first with a single replenishment cycle and the 3 possibilities comparing the actual demand with the forecast.

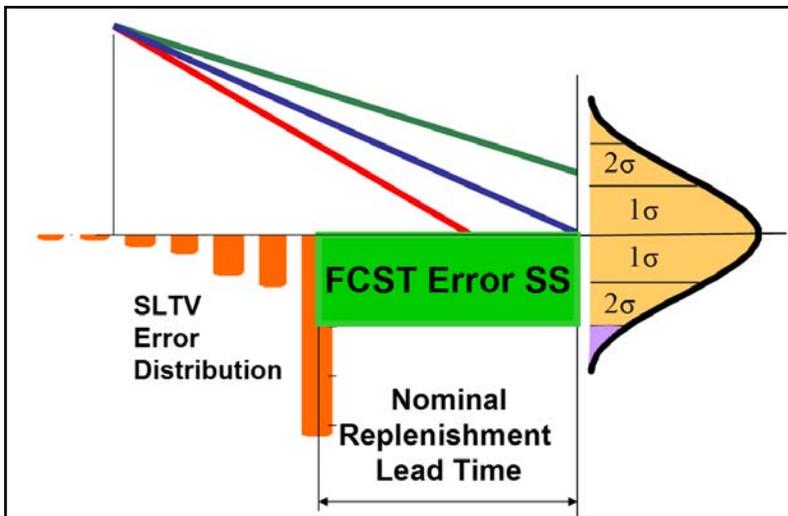
Now the safety stock required will be calculated from the (in this example) nominal lead time, taking into account the forecast error and error distribution for this SKU, in such a way as to achieve the desired customer service target.



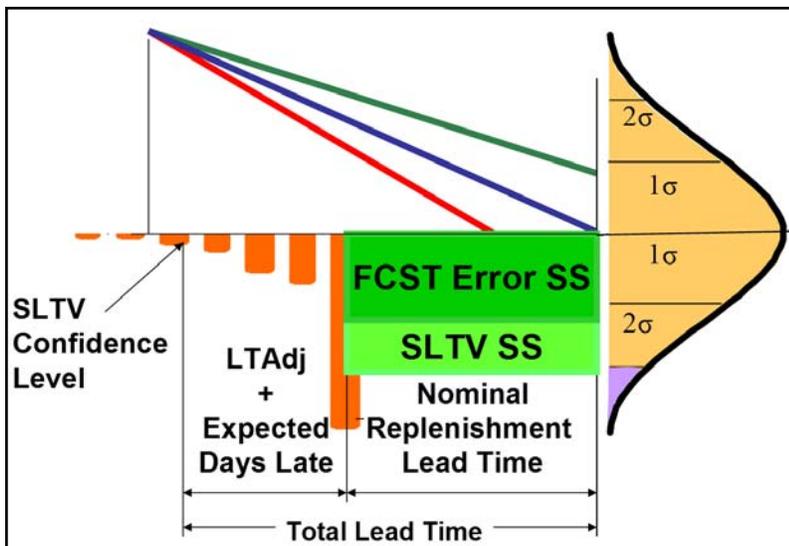
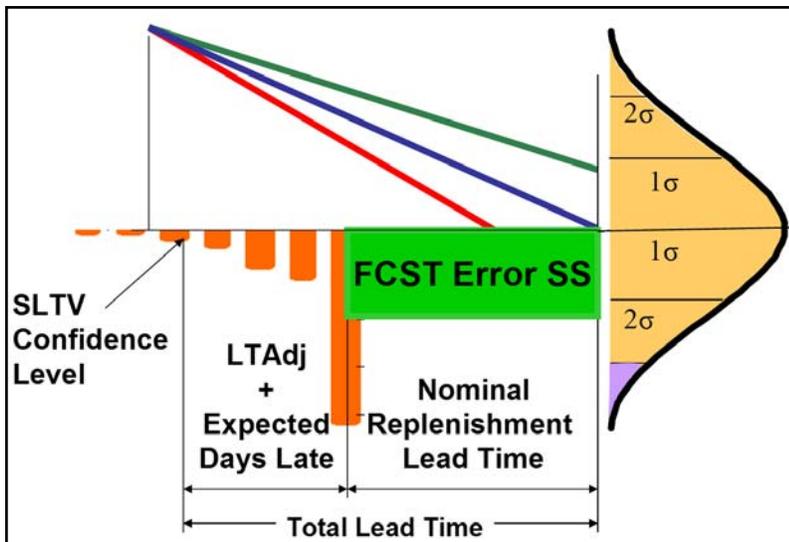
The tan shaded area under the error distribution curve represents the probability that is covered, either by the demand being less than the forecast, or by the cushion provided by the safety stock.

The purple shaded area under the curve is the remaining portion of

the probability not so covered, and is thus lost sales or perhaps expedites, if expediting is a possibility. The height (amount) of the safety stock is computed from the nominal lead time, the desired service target or expediting level, the forecast error, and the specific error distribution.



If we rotate the SLTV exponential distribution curve and flip it upside down next to the front end of the nominal lead time, we can see how it can be used to extend the nominal lead time to achieve full coverage for the nominal lead time plus any nominal adjustment (from a normally distributed SLTV) plus the expected days late.



On the left we now see the lead time adjustment plus the expected days late, based on the desired SLTV confidence level (here 99%). This new Total Lead Time is considerably longer than the original nominal lead time, which has the effect of spreading out the error distribution. That means that our original safety stock computed to cover the forecast error over the nominal lead time no longer suffices to deliver our target service.

We can restore the expected service by adding an amount of safety stock computed using the total lead time adjustments (nominal adjustment plus expected days late). We call this the *SLTV Safety Stock* or *Supplier Safety Stock*.

We typically compute the total safety stock using the Total Lead Time, but for comparison purposes, we want to convert the SLTV from days of additional lead time to pieces. We do this by subtracting the original safety stock (based on the original nominal lead time) from the new safety stock (based on the Total Lead Time). This gives us the

number of units of additional inventory required because of our vendor's inability to meet the nominal lead time. We can then multiply this *supplier safety stock* by the cost of the product to yield the number of dollars of additional inventory investment required. Thus:

$$\text{SupplierSafetyStock} = \text{TotalSafetyStock} - \text{OriginalSafetyStock}$$

where:

$$\text{OriginalSafetyStock} = \text{SafetyFactor} \times \sqrt{(\text{MSE} \times \text{NominalLeadTime})}$$

and:

$$\text{TotalSafetyStock} = \text{SafetyFactor} \times \sqrt{(\text{MSE} \times (\text{NominalLeadTime} + \text{LeadTimeAdj} + \text{ExpDaysLate}))}$$

where:

MSE = Mean Squared Error (the forecast error)

and:

all lead times and adjustments are expressed in the same units of time as the forecast calendar.

Note the footnote about the lead times. This means—if a SKU is forecasted using a monthly calendar—that the nominal lead time, any adjustment to the nominal lead time (because of an initial normal distribution), and even the “days late” must all be converted to months before computing the safety stock.

**Effect on Reorder Point**

As a practical matter, we should also consider the effect on the reorder point (ROP). Recall that the ROP is the sum of the forecast over the lead time (lead time usage or LTU) plus the safety stock. Or with our slightly modified terminology,

$$ROP = \text{NominalLeadTimeUsage} + \text{FcstErrorSafetyStock}$$

Since our lead time has changed, this affects both terms in the above equation:

$$ROP = (\text{Nominal LT} + \text{LT Adj} + \text{Exp Days Late}) \text{ Usage} + \text{FcstErrorSafetyStock} + \text{SLTV Safety Stk}$$

Now, what can you do with this data? Your first step should be to use this information to try to reduce SLTV. Here are two methods to consider.

**Prioritize Supplier Safety Stock Costs**

Rather than attacking SLTV at all suppliers, a Pareto analysis of the data, aggregated to the supplier level, gives you a good mechanism for prioritizing your suppliers. You can do this on a dollars-of-additional-inventory-investment-required basis, or (multiplying by your carrying cost) an annual-cost-to-carry-SLTV-inventory basis.

Either tool shows you clearly where to start to reduce the problem, as well as the benefit to be achieved by doing so. In the example shown at right, a single supplier accounts for over 50% of the total SLTV safety stock in dollars. If you have time to work with only a single vendor, that’s the one to start with. You may discover that acting on information at this level could eliminate the need to implement the additional inventory required to cover SLTV. And if that’s not the case, only 20% of the vendors account for 90% of the additional inventory dollars.

Line #	Vendor ID	\$ SLTV Safety Stock	Cumulative Value	% Total Value	% Vendors
1	Vendor 422	135,583	135,583	52.2	2.4
2	Vendor 111	25,728	161,311	62.1	4.8
3	Vendor 539	23,977	185,288	71.3	7.1
4	Vendor 101	18,257	203,545	78.4	9.5
5	Vendor 284	8,805	212,350	81.8	11.9
6	Vendor 366	6,299	218,649	84.2	14.3
7	Vendor 419	5,068	223,717	86.1	16.7
8	Vendor 513	4,965	228,682	88.1	19.0
9	Vendor 169	4,512	233,194	89.8	21.4
10	Vendor 307	4,397	237,591	91.5	23.8
11	Vendor 174	3,833	241,424	93.0	26.2

**Use Total Cost to Compare Competing Suppliers**

To evaluate the true cost of obtaining the same product from competing suppliers, you must consider all of your costs. These include, in addition to the quoted cost, the carrying cost of the normal safety stock which protects against forecast errors over the nominal lead time; the carrying cost of the supplier safety stock; and the carrying cost of the working stock. These costs should be divided by the annual forecast to obtain the total unit cost of a product per supplier which is expressed by the following formula:

$$\text{True cost} = \text{Quoted cost} + \text{Supplier cost}$$

where:

$$\text{Supplier cost} = [\text{CC} \times \text{QC} \times (\text{SS} + \text{SLTVSS} + \text{WS})] \div \text{12 Month Forecast in Units}$$

and:

CC = Carrying cost (e.g., 0.18 = 18%)

QC = Quoted cost (price/unit from this supplier)

SS = Safety stock (in units) to cover forecast error over nominal lead time

SLTVSS = Supplier lead time variation safety stock (in units)

WS = Working stock (in units; use ½ the lot size)

Note that the lot size (and thus the working stock) may be different for different suppliers (for example, if each has a different minimum buy quantity).

As the example at right shows, the lowest price quoted may prove

Part Number	Supplier	Quoted Unit Cost	Average Days Late	Annual Carrying Cost per Unit Sold	True Unit Cost	% Savings
ABC	XY10	100.00	35	15.22	115.22	0%
ABC	MN22	105.00	2	3.09	108.09	6.2%

more expensive over time. Buying from the supplier whose price is ostensibly 5% less expensive, is actually over 6% more expensive when all the costs are considered.

### **Implement Supplier Lead Time Safety Stock**

If you are unable to reduce your SLTV to an acceptable level, and you are not able to cover this variation by expediting or cannot pass it along to your customer, you should load the units required to cover supplier safety stock into your MRP system as a safety buffer to reach your target service level. Perhaps over time you can either work with your suppliers to reduce this additional inventory, or develop new suppliers without these problems.

One remaining question is what level of confidence should you use? If your goal is to prioritize SLTV safety stock dollars for the purposes of comparing and managing your vendors, then a 99 percent confidence level works well.

If your goal is to create a safety stock that will cover SLTV, then you will need to evaluate the trade off between inventory, service, and expediting. While we recommend you use this tool to manage your suppliers, if you cannot change a supplier's practices, then you want to use a confidence level that is identical to your service target for the SKUs in question. The chart at right, for example, shows some alternative confidence levels and their associated costs for one SKU.

Confidence Level	SLTV \$	SS	Days' Supply
90	\$355,641		6.2
95	\$453,915		7.9
98	\$578,945		10.1
99	\$670,246		11.7

For this client, we were able to eliminate the 3 week temporary fixed time supply, as well as half of the original 3 week fixed time supply to cover SLTV.

### **Supplier Quantity Variation (SQV)**

For some suppliers, the difficulty is quantity, instead of lead time (or in addition to it). As before, the goal then becomes to measure the variation. For external suppliers this is the difference between the quantity ordered and the quantity received; for internal suppliers, it takes the form of scrap, yield, etc.

With lead time variation, the assumption is that lead time is the same no matter what quantity is ordered. (While usually correct, this assumption may not always be true.) With quantity variation, however, the amount of the error (i.e., the difference between the quantity ordered and the quantity received) is usually related to the order size.

For example, it is not correct to say that the quantity error for lot 1 of 10 pieces (100 ordered, 90 received) for an item is the same as the error for lot 2 (1000 ordered, 990 received) for the same item. Instead, we need to convert the quantity variation to yield and track that. In this example, the yield for lot 1 is 90% (90 divided by 100), whereas the yield for lot 2 is 99% (990 divided by 1,000).

Still, tracking quantity variation via the yield allows us to do two things:

1. Identify the largest opportunities to reduce supplier quantity variation (SQV).
2. Determine a statistical yield for each SKU that provides the desired service.

What is the likely difference between the quantity you order from a supplier and the quantity you receive? If SQV is not managed properly it can create stock outs. A common technique for controlling SQV is to incorporate a standard yield into your MRP system that is equivalent to the range of SQV you allow as a matter of policy (e.g., an order is considered complete if the quantity received is plus or minus five percent of the quantity ordered).

There are two problems with this approach. The first is that it assumes that SQV is uniform for all SKUs. Where SQV varies from the standard yield, you could be suffering needless stock outs from an understated yield, or have too much inventory from an inflated yield. By measuring quantity variation, a statistical yield can be determined for each SKU that will deliver the service you desire for each SKU.

The second problem is that an “allowable” quantity variation is normally nothing more than a guess at what is reasonable. By measuring SQV you can determine the cost of variation and, as a result, set your allowable variation range as a function of cost. Finally, by measuring SQV you can know the cost of this variation and can identify the largest opportunities to attack and reduce it.

**Data Requirements**

The data required to measure SQV is the same as that required to measure the SLTV, with two exceptions. The dates of scheduled and actual arrival are replaced by the quantities ordered and received. We computed the quantity short or long. We also divided the quantity received by the quantity ordered to determine the individual yield for each order. A yield above 100% means the supplier shipped too much; a yield below 100% means too little. For SQV, we are interested in both shortages and overages. The chart at right shows the yields on some sample replenishment lots.

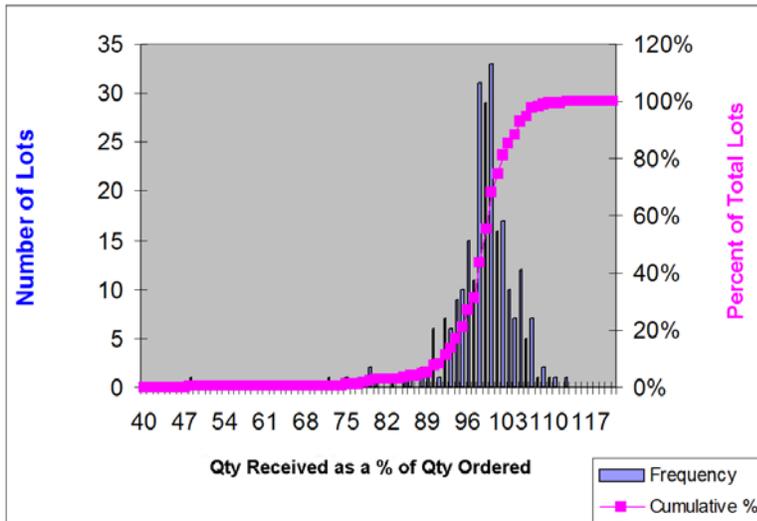
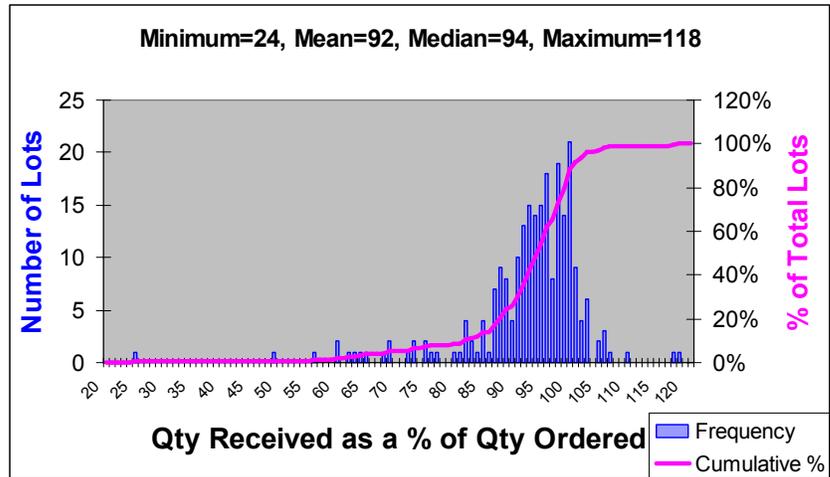
Date	Qty Ordered	Qty Received	Yield %
10/3/2002	586	444	76
12/18/2002	1,155	1,122	97
1/9/2003	1,155	1,127	98
1/9/2003	1,155	1,093	95
1/9/2003	1,155	1,074	93
1/16/2003	1,155	1,060	92
1/23/2003	1,155	1,059	92
1/23/2003	1,155	868	75
2/5/2003	1,155	1,035	90
2/27/2003	1,122	1,059	94

**Determine Distribution Type**

We started by displaying the yields in a histogram in order to determine the applicable distribution. In this study two distributions emerged. One (which occurred rarely) was the normal (Gaussian) distribution. The other (and far more typical) was a distribution that gave some appearance of normality, but was in fact not symmetric. We refer to this as a “skewed distribution”. While one could simply assume normality, the skew is to the shortage side, so this assumption would increase the risk of running out unacceptably often. Our initial approach was to concentrate primarily on the left side of the skewed distribution to reduce the probability of a shortfall to the desired confidence level. (It is possible that additional research could lead us to refine this technique.) In both cases the left side of the curve shows the distribution of supplier shortages, while the right side of the curve shows the distribution of supplier overages.

At right is an example of a skewed distribution. The horizontal axis is the yield percentage. The vertical axis on the left side is the number of lots. The left and right sides are *not* symmetrical, but in all other respects it appears normal. The left side rises more slowly than the right side falls.

The magenta line is the cumulative probability overlaid on the histogram. The scale for the cumulative probability is the vertical axis on the right side.

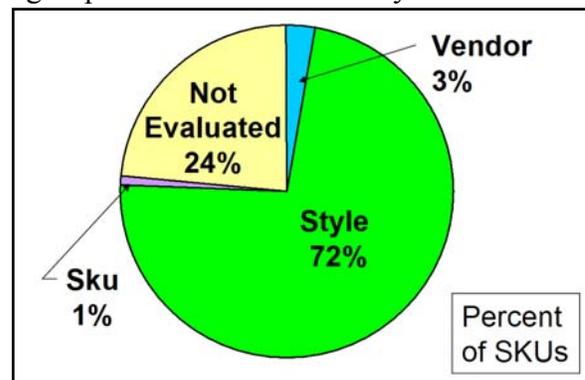


At left is an example of a SKU on the normal distribution. The left and right sides *are* symmetrical—or they will be, once we remove the policy and statistical outliers (*see below*).

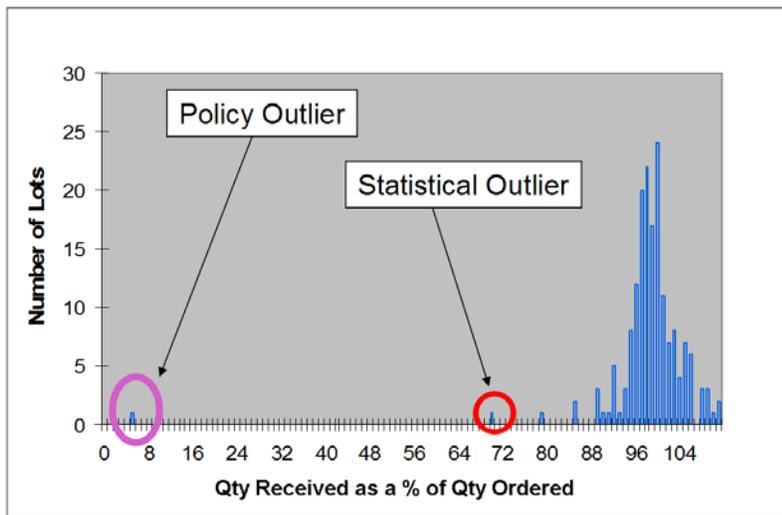
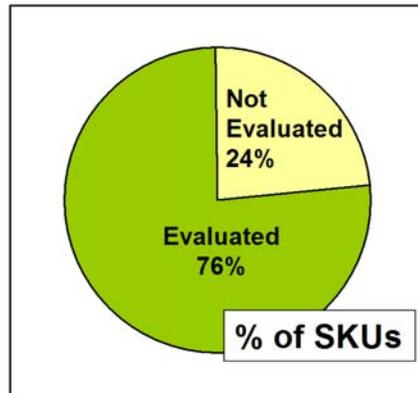
**Statistically Unusable Data**

As with SLTV, the next step is to determine if the SQV data is statistically usable. Like the SLTV process, the data for the products that are not statistically usable was aggregated to a higher grouping and the evaluation process repeated to identify the groupings whose quantity variation distributions were acceptable. If any data could not be grouped to form a statistically usable distribution it was not considered in the SQV analysis. As above, such instances occurred chiefly due to low volumes and the effect of eliminating them from consideration was negligible.

As the pie chart at right shows, the level of aggregation required to obtain usable data was similar to that of SLTV. Most of the analysis was conducted at the style level. The 24% of SKUs not evaluated due to unusable data was a bit higher than the 18% for SLTV.



This increase was not as large as it seems, however, given that the 24% of SKUs not evaluated comprised just 6% of the sales dollars. While this 6% was larger than the corresponding 2% in the SLTV study, the figure is still not large enough to invalidate this approach for SQV.



**Outliers**

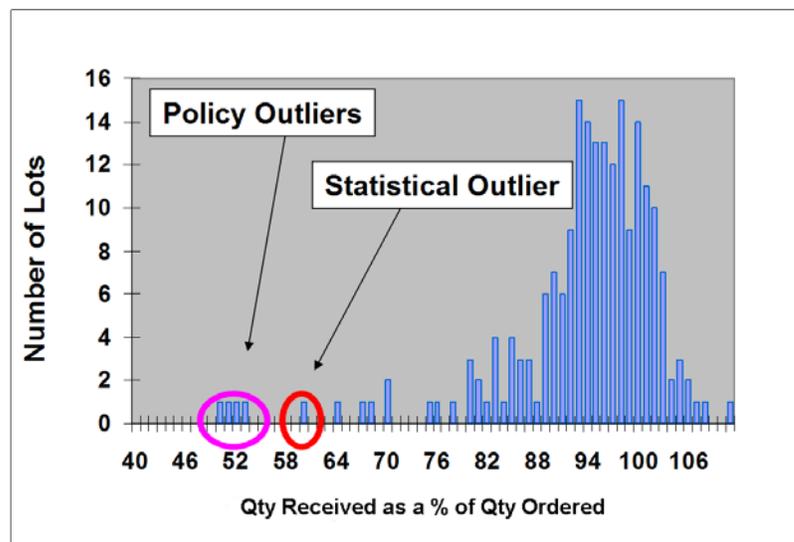
As with the SLTV study, we handled both policy and statistical outliers in the SQV data. The client set the policy limit to ignore any yields less than 60%.

At left is an example of a policy outlier and a statistical outlier for a SKU with a normal distribution.

Statistical outliers with a normal distribution were excluded if they fell outside a 4 standard deviation limit.

Similarly, at right is an example of a SKU on the skewed distribution with multiple policy outliers and one statistical outlier.

As with SLTV we used a 99% threshold for determining statistical outliers for SKUs on the skewed distribution.



**Lowest Yield and Expected Yield**

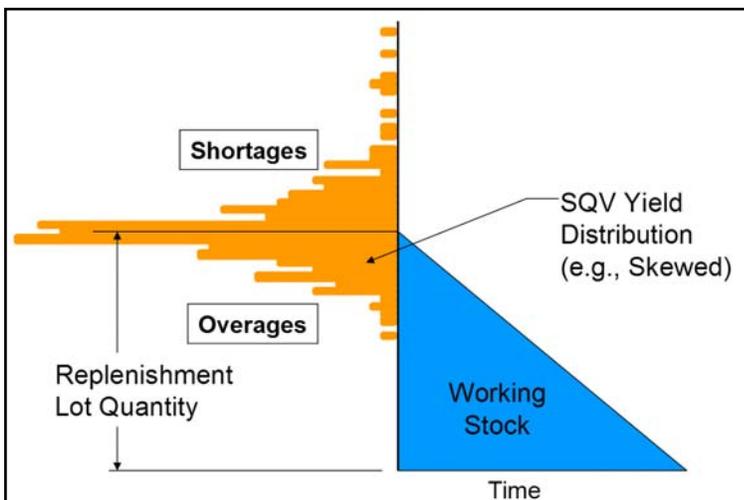
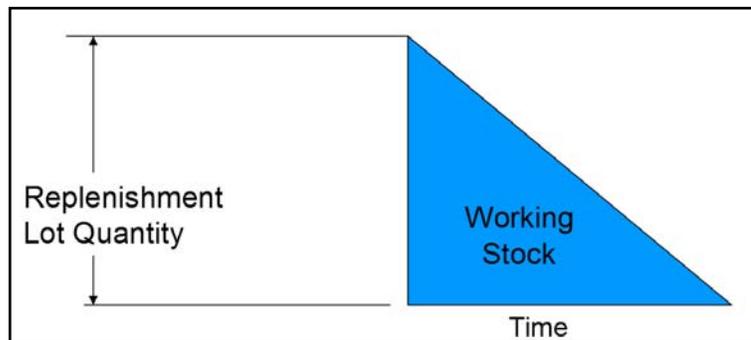
Now that the data is prepared, it is time to calculate the SQV. For normally distributed data, one simply computes the standard deviation of the yield and, plugging in a safety factor based on the desired confidence level, say,  $4\sigma$  for 99%, computes the yield which meets that level. You can then load this yield into your MRP system.

For example, let's say you calculate a lowest yield of 75% at a 99% confidence level. Note that a higher confidence level would give you a lower yield, while a lower confidence level gives a higher yield. If you need to receive 100 units, you divide 100 by 0.75 to calculate the quantity to be ordered of 133. It is important to maintain the nominal 100 RLQ and the additional 33 in separate fields, in order to compute the additional cost required.

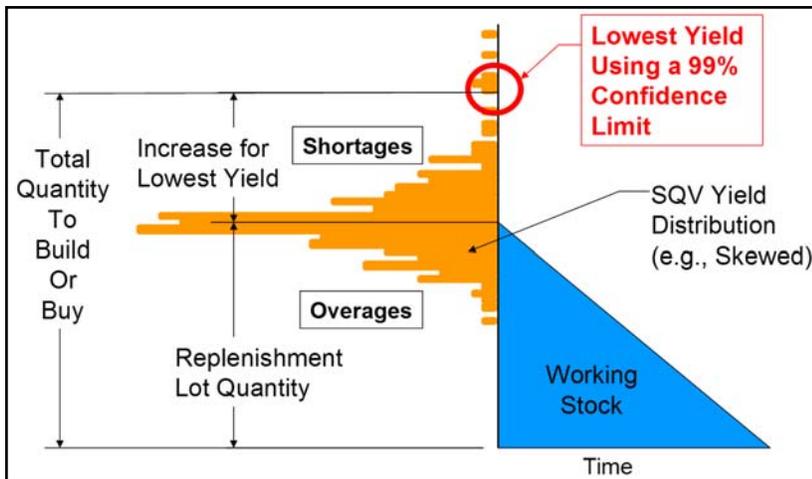
Let's further assume that not only is the data normally distributed, but also that the mean is 100%. In this case, since the expected yield for each lot is 100%, there is no penalty (in increased inventory) for ordering using a yield at the 99% confidence level.

For items with the skewed distribution—the great majority of the data we observed—we essentially do the same thing, except that we work only with the left half of the data (up to the mean) and (reversing it horizontally) treat it as exponentially distributed. As with SLTV, we then compute the yield associated with the desired confidence limit.

Let's look at this graphically. At right is a picture of a single replenishment cycle, with time on the horizontal axis and quantity on the vertical axis. The height is the Replenishment Lot Quantity or RLQ. Over time the demand consumes this supply until we're back down to zero stock on hand (assuming no safety stock).



Now we reverse the skewed distribution for this SKU and rotate it counterclockwise 90 degrees, centering the mean (the expected yield) vertically on the RLQ. As you can see from the picture, if we keep the current RLQ that means we will receive deliveries which are short of the RLQ. So let's move up to a point where that will happen less frequently.



To do that we find the point on the vertical axis corresponding to a 99% confidence limit. This is the lowest yield we can expect with 99% certainty.

Dividing the nominal RLQ by this yield gives us the total quantity to build or buy. Subtracting the RLQ from that tells us the additional amount of stock that must be added to our RLQ in order to receive the expected value with 99% confidence.

To recap the formulas:

$$\text{LowestYield} = \text{Median (yield)} - 4 \times \text{StandardDeviation (yield)}$$

$$\text{ExpectedYield} = \text{Median (yield)}$$

for normally distributed SKUs/groups

or:

$$\text{LowestYield} = \text{Mean (yield)} - \text{ABS (Mean(LeftSideYields) x ln (1.0 - 0.99))}$$

$$\text{ExpectedYield} = \text{Mean (yield)}$$

for SKUs/groups on the skewed distribution

where:

ABS(x) = Absolute Value (x)

ln (x) = Natural Logarithm (x)

LeftSideYields = All of the yields that are less than or equal to the mean of all the yields

At right are some examples of how the Expected and Lowest Yields compare with the client's 5% rule of thumb yield policy. If you use the same 5% yield on all SKUs then you will have great service and poor turns on some SKUs, and just the opposite on the others. But by calculating individual yields you will have the best possible turns and service on all SKUs.

Part Number	Annual Replenishment Frequency	Rule Of Thumb Yield	Expected Yield	Lowest Yield
PN#1	120.0	95.00%	98.22%	77.75%
PN#2	102.3	95.00%	98.63%	79.38%
PN#3	92.7	95.00%	94.24%	61.51%
PN#4	79.2	95.00%	94.15%	65.80%
PN#5	101.8	95.00%	98.61%	79.66%
PN#6	80.5	95.00%	94.02%	58.94%
PN#7	75.1	95.00%	92.91%	50.07%
PN#8	65.8	95.00%	92.92%	57.80%
PN#9	65.3	95.00%	93.14%	59.57%
PN#10	55.7	95.00%	97.44%	74.99%

### SQV Excess

Implementing these individual yields—rather than applying a standard yield percentage across all parts and suppliers—has the effect of increasing the working stock, which in turn raises customer service—restoring it to the service target, one hopes. It also raises inventory somewhat—even more so when a supplier has a history of shipping too much as well as too little on separate occasions. If a supplier ships too much on one delivery, the customer adjusts his next order to be released

downward to compensate. However, depending on the volume and lead time, the customer may already have several released orders in the pipeline and these cannot normally be altered since they are firm. Therefore to quantify this additional inventory we first compute the maximum expected number of simultaneous released orders per year and then multiply that by the expected overage per order. This is the additional inventory caused by the supplier's inability to ship the exact amount ordered.

For data from either distribution (skewed or normal), where the expected yield is not equal to 100%, we must compute the effect on inventory. For example, let's consider the item graphed at the top of page 12. Using the process described above, and assuming a desired 99% confidence limit, the derived yield is 50%. Dividing a lot size of, say, 100 units by 0.50 gives us a quantity to order of 200 units. We know there is a 99% probability that the quantity received will be at least half this amount, or 100 units. However the expected yield is considerably higher at 92%, so over some period of time (for convenience, we'll use one year) we can expect to receive an average of 92% of the 200 units, or 184 units per lot. This means we will carry an extra 84 units times the maximum number of simultaneously released orders over a year's time. We refer to this amount as the "SQV Excess".

**Prioritize by SQV Excess Cost**

Multiplying the SQV Excess units by the standard cost and carrying cost enables us to determine the annual cost of SQV Excess. As with SLTV, you can then aggregate this cost to the vendor level to identify where SQV costs are greatest and attack those first.

In the case of our client, just 20% of their suppliers accounted for over 80% of the SQV Excess cost (*see cumulative chart at left*).

Cum Supplier Percent	Cum Supplier Excess Inv Dollar Percent
2.65%	27.22%
6.62%	50.81%
19.87%	80.52%
43.05%	95.11%
65.56%	99.02%
100.00%	100.00%

The top line of the left chart corresponds to the 4 top (which is to say, worst) vendors (*see chart at right*), who cumulatively account for

Supplier	Cum Supplier Percent	Cum Supplier Excess Inv Dollar Percent
ALP	0.66%	9.29%
ALN	1.32%	16.11%
AGD	1.99%	22.14%
GVV	2.65%	27.22%
AXM	3.31%	32.00%
GDG	3.97%	36.31%
ARP	4.64%	40.50%
GDN	5.30%	44.63%
GFQ	5.96%	47.93%
AL1	6.62%	50.81%

over 27% of the SQV Excess inventory. Using the Pareto chart as a priority list allows us to achieve the maximum possible improvement with the minimum possible effort.

**Use Total Cost to Compare Competing Suppliers**

Once you know the annual cost to carry the SQV Excess inventory, you can use that in the same way as with the SLTV above to compare competing suppliers on a total cost basis. Alternatively, you could plug in supplier-mandated variation allowances to determine your costs and then use that knowledge in price negotiations. For example, on a \$10 item where you might be willing to accept delivery of plus or minus 2 percent; but the vendor asks for plus or minus 10 percent; if you calculate that these limits increase your costs by 75¢ per unit, you can agree to the variation if they reduce the cost to \$9.25.

## Summary

We believe that the greatest benefit of SLTV and SQV analysis is in the knowledge it gives you of your costs of doing business with less-than-perfect suppliers, and the leverage that provides for reducing or eliminating the variation. (For suppliers who do perform without variation you could even reward them by accepting a higher price—since they cost you less than other suppliers.)

The usual alternative to such analyses are rules of thumb: for example, add X weeks to all lead times (for all products and vendors) just in case the suppliers deliver late—where X is derived by guess or by trial and error. Or, assume a yield loss of Y percent for all deliveries from all suppliers to cover quantity variation, where Y is similarly derived.

When variation cannot be eliminated, the benefits of replacing such rules of thumb with statistically calculated values appropriate to each product or group of products include both improved service and reduced inventory—but the service is the greater of these. For example, one study we performed on SLTV actually showed a very slight increase in aggregate inventory. Looking at the numbers for each product, however, revealed that the allocation of the inventory changed dramatically. Some items had their safety stocks reduced by thousands of dollars, while others had them increased by thousands of dollars. So although the total inventory generated by the rules of thumb wasn't too far off optimum, the service was significantly (and unnecessarily) reduced because of the incorrect allocation.

## About the Authors

**John A. Estep**, CFPIM, is president of E/Step Software Inc., a Yakima, Washington based company specializing in education and software for finished goods/service parts forecasting and inventory planning. A frequent speaker at industry conferences, Mr. Estep has written dozens of conference and trade journal articles and was a columnist for APS (Advanced Planning & Scheduling) Magazine, writing their "On Demand" column. With a background in mathematics, operations research, and computer engineering, he worked on his first forecasting system for an apparel manufacturer in 1970, and has since counseled hundreds of companies on their forecasting and inventory planning needs. Mr. Estep is the chief architect for his company's Finished Goods Series™ (FGS) demand forecasting and inventory planning software.

Mr. Estep may be reached by phone at (509) 853-5000 or by email at [John.Estep@EstepSoftware.com](mailto:John.Estep@EstepSoftware.com).

**Nathan A. Boyd** currently manages the forecasting, inventory planning, and cash flow for Southern Enterprises Inc. of Coppell, Texas. He is also the President of Forparco, an Application Services Provider (ASP) for E/Step Software Inc. since 1997. In his ASP and consulting capacities, Mr. Boyd has worked with such companies as GE Transportation Systems, Russell Athletics, Caterpillar, Rheem HVAC, Ausimont USA, Hastings Manufacturing, Dowty Aerospace, and AAR Corp.

Prior to founding Forparco, Mr. Boyd spent 18 years with Batco, a heavy-duty truck service parts company in Dallas, Texas, where he served in a variety of positions, most recently as Vice-President of Operations. Mr. Boyd was responsible for selecting and implementing FGS to manage Batco's inventory of 60,000 SKUs. He used FGS to maintain their already high service level, while cutting inventory by 48%.

Mr. Boyd may be reached by phone at (469) 955-7725 or by email at [Forparco@aol.com](mailto:Forparco@aol.com).