APICS Executive Briefing:

Demand Forecasting & Inventory Planning

for Manufacturers & Distributors

by John A. Estep, E/Step Software Inc.

Introduction

These two closely-related subjects contain the solutions to many expensive problems encountered by manufacturers and distributors—problems that executives, managers, and practitioners often believe have no solution. But forecasting and inventory planning are a science, not an art, and a little knowledge of this science can pay big dividends in terms of improved service and reduced inventory and expediting. Without this knowledge, executives and managers cannot hope to guide their people to achieve superior results. The good news is: one does not need to be a statistician or an engineer to understand the general principles of this science. This presentation provides an overview of the concepts and relationships, linking them to typical operational problems and reviewing solution alternatives.

We all know what *Demand Forecasting* is but bear in mind that *Inventory Planning* means figuring out what your inventory *should* be (*not just counting what you have)*. This is an Executive Briefing. That means we will stay at a high conceptual level with almost no formulas.

Forecasters and inventory analysts probably have more power to alter their employer's fortunes—for good or ill—than just about anyone else in the Operations Management area, or even, perhaps, in any area of the organization. Yet this is seldom recognized by others, and that, unfortunately, often includes their own management. So I hope to shed some light on this critical corporate function and perhaps to dispel some myths along the way.

Why Demand Forecasting & Inventory Planning?

Are you in a purely make-to-order environment where your customers are willing to wait your full lead time (and that includes your suppliers' lead times too) for all of your products including any service parts? If so, this material does not apply to you. If not, then you have no choice—you must forecast your demand and plan your inventory (and perhaps also your expediting).

Applicability

Let's look at three examples as a way of gauging applicability.

Service: The service parts operation of a Chicago-area food machinery manufacturer came to us concerned about their service levels. We did a study of a sample of their parts and showed them that we could cut their inventory by a third and raise service to 99%. While we were at it, we also showed them that with their current aggregate inventory they could achieve 100% service. And that's what they decided to do. They removed the inventory where they had too much and plowed it all back into other SKUs where they needed more. Their service promptly climbed right up to 100% *on the same aggregate inventory with which they started*. Wouldn't you just love to be one of their customers? What a commitment to service!

Let's think about this for a minute, however. Two thirds of their current inventory could achieve 99% service. It then took the entire remaining third to get to 100%. Looking at it from a different perspective, however, if they had that two thirds of their current inventory, correctly allocated, and were achieving 99% service, by how much would they then have to increase inventory to get to 100% service? In that case the basis has changed, so for them to go from 99% service to 100% service would require a 50% increase in inventory! Why is that last 1% of service so expensive? We will cover that shortly.

Inventory: Another company, a provider of jet engine service parts, said their service wasn't bad but they thought the inventory required to reach it was questionable. Plus they wanted to improve the productivity of their forecasters. Again we did a study, the results of which exactly matched the results of their subsequent implementation, which was that they raised service slightly while cutting inventory 30%. That 30% inventory reduction was over $300 million of unnecessary inventory which they eliminated over a 2 year period. Their return on investment took a noticeable leap. And during the project they cut their planning headcount 30%. [1]

Inventory and Service: Lastly, a fastener manufacturer in the Boston area questioned both their inventory and service levels. This company had significant seasonality in half of their SKUs. After we educated them and they implemented our software, they raised their service dramatically—to their 95% target—and slashed inventory by 75%. Before beginning they had 4 times as much inventory as they needed! (I'll address how that's possible a bit later in the briefing.) This company occupied two buildings across a street from each other. One was a combined factory and warehouse. The other was strictly a warehouse. They sold the warehouse and converted half of the remaining warehouse space in the factory to production space. When I visited them a few months later they had a huge banner stretching across the entire wall in the planning department that said, "Don't buy anything!"

The science we will be discussing is applicable to a wide range of industries. We've applied it to everything from jet engine service parts to panty hose; from hamburgers to hot sauce to oilfield chemicals; in the following general categories:

1. Process Manufacturers
2. Discrete/Repetitive Manufacturers
3. Service Parts Manufacturers
4. Distributors of products from the above
5. Overhaul/Service shops
6. Electrical Utilities

Essentially this science is applicable to any business with an inventory to manage.

Let me draw particular attention to service parts which, as it turns out, is the toughest business to forecast—primarily because of the large forecasting errors (due to sparse demand) but also because of long lead times.

E/Step Software Inc. Overview

E/Step Software was founded in 1983 to take advantage of the then-new IBM Personal Computer in addressing the problems inherent in transaction-processing based planning systems, which is, that they ignore the analytical front-end functions of Demand Forecasting and Inventory Planning—which are so critical to success.

In our increasingly complex—yet, paradoxically, increasingly math-phobic—society, solving these problems requires high-level statistical tools that can be used by ordinary planners (who are not statisticians).

Our software is called the Finished Goods Series, or FGS. It embodies the science of demand forecasting and inventory planning. Our people have degrees in mathematics and engineering, plus a combined experience of over 100 years in the industry.

Our first few clients were service parts companies; including GE Transportation's locomotive service parts operations which began working with us in 1984. So we did the research to extend this science and we developed a number of tools specifically for the service parts industry. We were subsequently surprised to discover that the same tools used differently were also quite helpful to other industries. We have clients in North, South, and Central America, Europe, Asia, and Australia. Our clients make FGS the front end to their ERP/WMS/OSB and/or Supply Chain systems.

Benefits

We educate our clients in the science of demand forecasting and inventory planning and provide them with the software tools which enable them to achieve the following benefits:

1. Raise forecast accuracy
2. Increase customer service to their targets
3. Reduce inventory by 30 to 50%
4. Reduce expediting
5. Reduce setups or receiving
6. Raise forecaster/inventory analyst productivity

Principal Functional Areas

The goal of **Demand Forecasting** is to predict what products your customers want, where they want them, when they want them, and how much they want—without regard for your ability to meet that demand.

Inventory Planning tells you where you need inventory and where you don't. It manages the tradeoff between service, expediting, and inventory. Essentially it's a price list for service (by part by location)—something very few companies have ever had. Or, if the squeeze is on return on investment (ROI), given an inventory budget, Inventory Planning will tell you what service you can afford (and how often you will need to expedite).

With a demand forecast and an inventory plan in place, **Replenishment Planning** creates a replenishment plan that meets both the demand forecast and inventory plan, takes into consideration what you have on hand, what you already have on order, and makes it all happen despite real-world constraints, such as the need to fill containers, to plan around Chinese New Year, or to send a truck out on a regular multi-day route that stops at several vendors to pick up material, among other constraints.

For companies with multiple satellite warehouses, **Distribution Planning** ensures that stock distributed to a satellite warehouse meets each its demand forecast and inventory plan, while also ensuring that each satellite gets its fair share of the total available inventory.

Science, not Art (or Fad)

Forecasting & Inventory Planning are a *science*, not an *art* or a *fad* or even a *matter of opinion*. As with any science there are concepts and relationships to learn. For example, there is a relationship between service and inventory; but it's more complex than that as lead time, replenishment frequency, forecast error and error distribution plus some others are all involved as well. If you master these concepts and relationships, you can put them to work for you, making your work easier and raising your ROI. But if you are unaware of them or ignore them—or even do them *almost* correctly—they will make your work very difficult, if not impossible.

What kinds of difficulties? Every company we have worked with has had several of these problems. Which of these do you run into?

1. Run out of stock too often?
2. Have too much inventory for many products?
3. Have too little inventory for some products?
4. Have trouble forecasting your demand?
5. Spend too much time expediting?
6. Have a warehouse that is too small?
7. Lack manufacturing capacity?
8. Suffer from low inventory turns?
9. Have competitors who give higher service?

So let's look at a few concepts, starting with inventory planning.

Why do we have inventory?

For lots of reasons...*but they should all be known and quantified.* For example, we have working stock because we can't make or buy just one every time we sell or use one. We have safety stock because of lead times, forecast errors, service requirements, etc., and our customers are unwilling to wait. We may have extra stock to buffer against supply variability, for example because of poor vendor performance, manufacturing problems, or seasonal product availability. Lastly we could require hedge stock because we may know of a future price increase.

Of course the above could occur at any level of the bill of materials, i.e., finished goods, work in process, or raw materials. For this discussion, we will concentrate just on the first two reasons for having inventory: working stock and safety stock, and just at the finished goods level. When we started in business, we concentrated at first on just the finished goods. Eventually, however, we did the research and development necessary to expand the science to address inventory at all levels.

Working Stock

Let's look at a picture of inventory level over time:

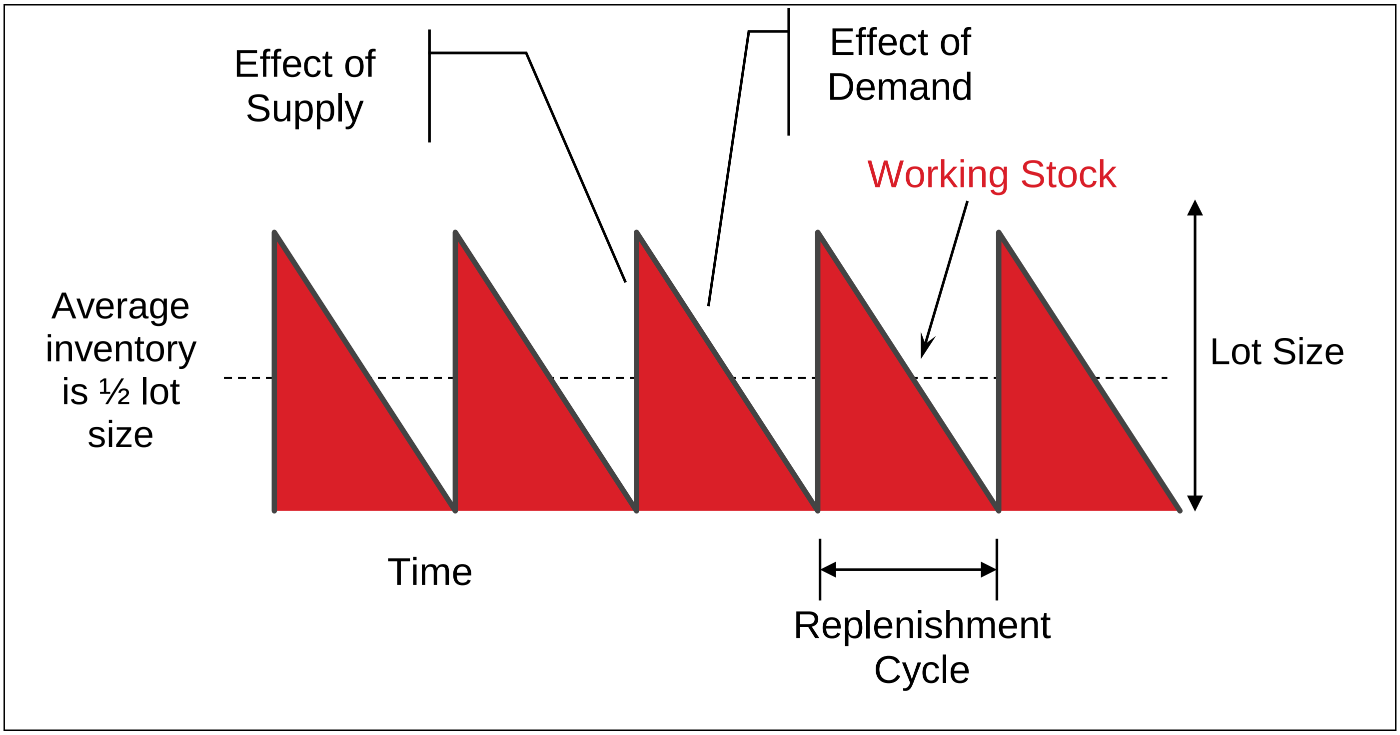

When we receive a new lot, inventory goes up to its maximum. Then over time, demand whittles that down until we reach zero. Ideally, at that instant the next replenishment lot arrives and we start the cycle all over again. Since the inventory oscillates between zero and a full lot size, the *working stock* (or *cycle stock*) is defined as half a lot size. The *replenishment cycle* is the lot size expressed in time. So if a lot size of 100 is enough to last one month, the replenishment cycle is one month. Similarly, the *replenishment frequency* is 12 times per year.

Note that this is an idealized environment, with perfectly smooth demand that exactly matches the forecast. Note also that the replenishment cycle is not the same as *lead time*; the lead time could be longer or shorter than the replenishment cycle.

Sometimes the working stock is the majority of the inventory, but sometimes it's not. To calculate working stock we need to calculate the lot size. This usually involves (but should not be limited strictly to) the *economic order quantity* or *EOQ*. EOQs have been so thoroughly discussed elsewhere that we will not do so here.

Many things go into the lot size/working stock calculation; these should include at a minimum:

1. Annual unit forecast
2. Product cost
3. Setup cost (or lost opportunity cost)
4. Carrying cost (or management's willingness to invest in inventory)
5. Forecast error *(not demand variability)*
6. Shelf life
7. Minimums, Increments, Maximums

Some inventory planning tools consider none of these. Many consider just the first 4. If you don't consider the forecast error, you can end up with too much inventory! Especially where your forecast errors or lead times, or both, are large.

Safety Stock

Now let's go back to the picture of inventory over time, but this time adding in some real-world effects.

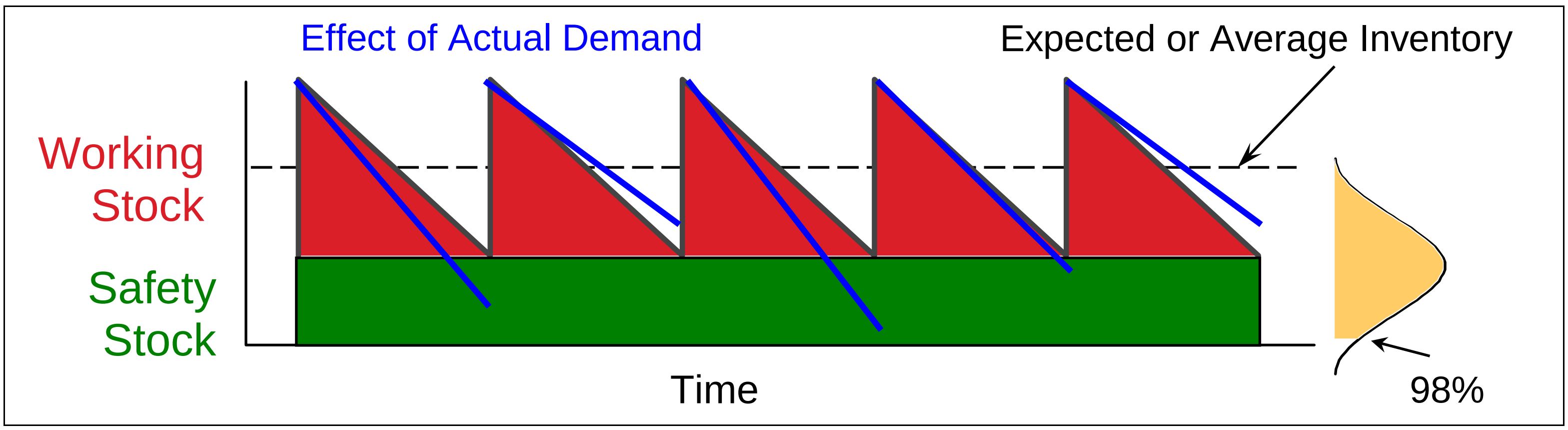

Expected Inventory = Working Stock + Safety Stock

As before, we see the effect of lot sizes idealized, but now, in the real world, demand is higher or lower than forecast. We might even run out of stock. So we add safety stock carefully calculated to provide 98% service; or whatever our service target is. Now our average or expected inventory is equal to the working stock plus the safety stock. Safety stock protects against forecast error during the lead time. It provides control over service.

Raising Service While Reducing Inventory

So how is this possible? Let's consider a very simple example.

Suppose you need to determine how much safety stock is required to achieve a 98% level of customer service for a company with just two parts, A and B:

Part	Cost	Forecast Usage/Month
A	$ 1.00	2,000
B	1.00	2,000

Since each part costs $1.00, we can work with dollars or units interchangeably. Most companies, when faced with this question, handle the situation by educated guessing. If service is too low, say 90%, and finished goods safety stock last year averaged six weeks, they may decide to increase it this year to perhaps two months (which is quite possibly the most frequently used safety stock method in the world today). *In reality, there is not enough information provided about this situation to make any decision.*

Let's provide some additional information and see if we can do better:

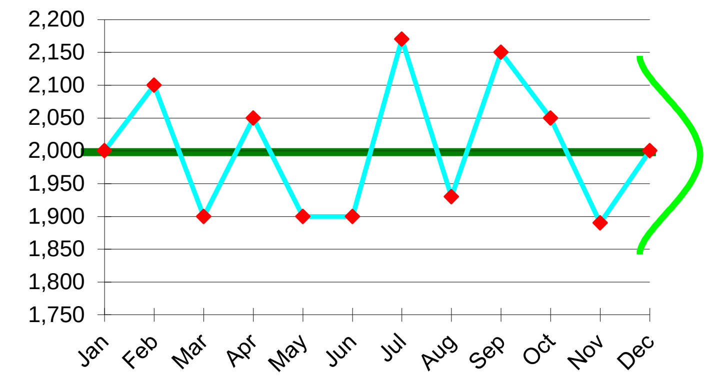

Product A's demand is clustered tightly around the forecast:

- The average is 2,000
- The high is 2,180 and the low is 1,900
- The forecast error (determined statistically) at 98% confidence is 200 units.

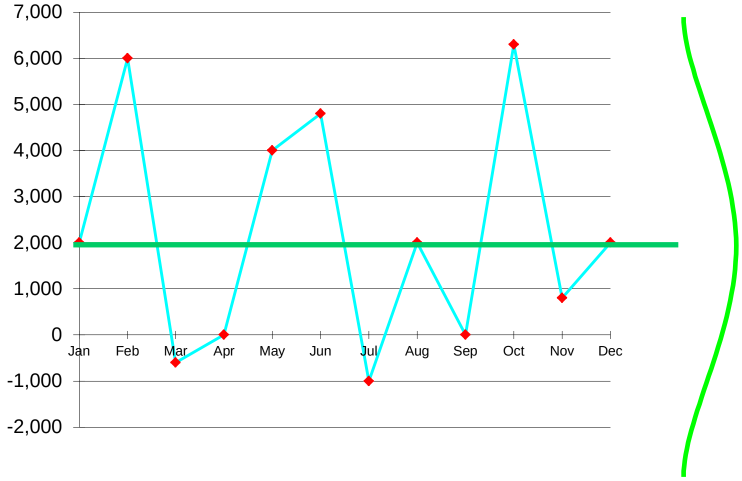

Product B's demand is often far above or below the forecast:

- The average is 2,000
- The high is 6,500 and the low is -1,000 (with returns)
- The forecast error (determined statistically) at 98% confidence is 6,000 units.

So we can add the 98% confidence limits to our table:

Part	Cost	Forecast Usage/Month	**98% Confidence Limit**	2 Month Safety Stock
A	$ 1.00	2,000	**+/- 200**	$ 4,000
B	1.00	2,000	**+/- 6,000**	4,000
Total				$ 8,000

Now parts A and B no longer look identical in all respects—in fact their forecasts are very different. The forecast for part A, with 98% confidence, is between 1,800 and 2,200 pieces per month. B's forecast, with the same confidence, is between net returns of 4,000 (zero, if returns are not possible) and 8,000, but with a most likely value of 2,000. Clearly, different actions are justified for the two parts! In fact, since the given forecast confidence interval just happens to match our desired service level, the correct safety stock is simply the same value as the confidence limit.

Part	Cost	Forecast Usage/Month	98% Confidence Limit	2 Month Safety Stock	Correct Safety Stock
A	$ 1.00	2,000	+/- 200	$ 4,000	$ 200
B	1.00	2,000	+/- 6,000	4,000	6,000
Total				$ 8,000	$ 6,200

This safety stock achieves our objective of 98% service, but at a cost of almost 25% less than the two month safety stock. What happened to service in the meantime? Consider the two parts individually.

Part	Cost	Forecast Usage/Month	98% Confidence Limit	2 Month Safety Stock	Service with 2 Month SS	Correct Safety Stock	Service with Correct SS
A	$ 1.00	2,000	+/- 200	$ 4,000	~100%	$ 200	98%
B	1.00	2,000	+/- 6,000	4,000	~88%	6,000	98%
Total				$ 8,000	~94%	$ 6,200	98%

For part A, the two-month supply is 20 times more than necessary. So the service is something higher than 98% (most likely 100%) and now drops to exactly 98%. For part B, the two-month supply is only two-thirds enough to reach 98% customer service. The information provided is insufficient to tell precisely, but it is safe to assume that service was substantially below 98%; we'll estimate it at 88%. This means aggregate service was only about 94%. So you have a choice of 94% service at a cost of $8,000, or 98% service (the target) for $6,200. Tough decision!

We've learned 2 important lessons here:

1. It is possible to raise service and lower inventory simultaneously. *All that requires is that you have more than 1 part in inventory!*
2. Fixed time supply guesses do not work!

Four Methods for Computing Safety Stock

With fixed time supply you pay too much in inventory for service that falls short of your objectives. Safety stock, whose purpose is to cover errors in the forecast, must be set according to the *size of the errors*—not according to the *size of the forecast*. Why do most companies use fixed time supply guesses? Because they are easier to describe in detail and easier to program. Unfortunately they do not do the job. Besides, most practitioners are using the software, not writing it; and they don't need to know the details.

The reason this works with such dramatic result is that most forecast errors are distributed in such a way that there are usually a large number of small errors and a small number of large errors. Setting safety stocks using a fixed time supply results in too much inventory for most items (the ones with the small errors) and too little inventory for a few items (the ones with the large errors).

As it happens, fixed time supply guesses are not the only things that don't work. When I started my APICS Certification classes in 1978, I learned the Z Factors method for calculating safety stock. Now I promised to avoid formulas as much as possible so I won't put up any equations, but let's think about it. To compute safety stock with this method, you look up your service target in the Z Factor table to get the corresponding safety factor and then you multiply that by the forecast error. That's your safety stock. Here are some typical service levels and Z factors:

Service Level (% of Order Cycles without Stockout)	Z Factor (Number of Standard Deviations)
50.00%	0.00
84.13%	1.00
90.00%	1.28
95.00%	1.65
97.72%	2.00
98.00%	2.05
99.00%	2.33
99.86%	3.00
99.99%	4.00

What is not taken into account in this method is your working stock. Working stock is inventory too, so don't you get service out of it? Of course you do! Yet the Z Factor method doesn't take your working stock into consideration. In fact, the Z Factor method implicitly assumes that your replenishment frequency and your forecasting frequency are the same. So if you forecast monthly, for example, your safety stock will correctly achieve your service target only if the product is also replenished monthly. In my experience the two are rarely the same. So that means you are usually getting service that is less than you need, or you are paying for service—with inventory—that is more than you want.

As an example, imagine you have a product you forecast monthly and replenish monthly. You compute your safety stock to achieve a 99% service target using the appropriate Z Factor, which is 2.33. With monthly replenishments there are 12 opportunities per year to stock out. Now suppose management decides to replenish this part once per decade instead of once per month. Now how much safety stock do you need? Probably none! After all, you can run out only once every 10 years!

On the other hand, suppose management decides to replenish the part once per day. Now some of you might say, "We've got stock arriving every day—we can never run out!" That's not the way it works. Every peak on that sawtooth curve occurs at the same time as a valley. So you have an opportunity to run out of stock every day, or about 260 times per year. You will need much more safety stock to hit your 99% target if you replenish daily than if you do so monthly, because *your probability of stocking out is much greater*.

With that in mind, let's compare some methods for computing safety stock. Here is data taken from a study we did for a client some years ago. The sample contained data on 27 SKUs and the service target was 99% dollar fill. Here are summary statistics across all 27 SKUs:

Safety Stock Calculation Method	**Minimum % Service**	**Maximum % Service**	**Aggregate % Service**
Z Factors guess	99.66	99.97	99.80
Fixed Time Supply guess	83.46	100.00	99.00
Equal Service	99.00	99.00	99.00
Optimized $ Fill	96.00	99.50	99.00

For the first method, the **Z Factors guess**, we multiply the forecast error by 2.33 for every SKU and that becomes the safety stock. Then our software, knowing the working stock (which is the same for all 4 methods) and the resulting safety stock, computes the percentage dollar fill service for each SKU. Of the 27 SKUs, the lowest service is 99.66%, well above our target (which is not good, because we have to pay for that). The maximum service over the 27 SKUs is 99.97%, and the aggregate service 99.8%. So in this case, Z Factors give us service much higher than we want or can afford. (Remember that first client cited above that needed 50% more inventory to go from 99% to 100% service? That's where the Z Factors method is putting us here.) Depending on the replenishment frequencies, service could also have been much lower than we want—which is why this method is considered a guess.

The **Fixed Time Supply guess** is both better and worse: better in that the aggregate service hit our 99% target; worse in that service on individual SKUs ranged from 83.46% to 100.00%. So this is just a guess too.

The **Equal Service** method is smart enough to recognize that working stock provides service too. To be able to do that, it uses a bunch of probability functions as well as a little gem called the partial expectation function. Without these it cannot take the service provided by the working stock into account (*but thankfully, the user doesn't need to understand them!*). So Equal Service computes only the safety stock needed (*in addition to the working stock*) to achieve the service target for each SKU. That means every SKU gets exactly 99% service.

Lastly, the **Optimized $ Fill** method is similar to, but even smarter than, the Equal Service method. The algorithm (a) relaxes the requirement that every SKU be exactly on target, (b) still requires that the aggregate service hit the target, and (c) optimizes the distribution of service levels among the 27 SKUs so that the dollar value of the resulting aggregate safety stock is the minimum possible. That means it pushes up the service on some SKUs while lowering it on others until it finds the safety stock with the lowest aggregate cost. Because of the requirement to meet the aggregate target and because there is a lot more room below 99% than there is above it, for every SKU whose service is pushed down, the optimization typically must raise the service for several SKUs. This means that most customers will usually see higher than 99% service. In this example, the method pushed down service to 96% for the lowest SKU and raised it to 99.5% for the highest, while still maintaining 99% aggregate service.

Now let's look at the inventory costs associated with these 4 methods for computing safety stock.

Safety Stock Calculation Method	Average Safety Stock in Days	Total Safety Stock $	$ Savings Over Z Factors Guess	% Savings
Z Factors guess	12	$ 187,605		
Fixed Time Supply guess	11	$ 166,982	$ 20,623	11.0%
Equal Service	8.7	$ 136,530	$ 51,075	27.2%
Optimized $ Fill	8.5	$ 133,996	$ 53,609	28.6%

Unfortunately, the Z Factors guess is the most expensive of all 4 methods by far; costing $187,605 in safety stock to provide service that is higher than our desired target of 99%. At high service levels the inventory/service curve is very steep—practically vertical (see below)—so raising service there is extremely expensive. The Fixed Time Supply guess set safety stocks for every SKU to an 11 day supply as that was what would achieve our 99% target. This is the second worst method for computing safety stock, here costing $166,982 to achieve 99%. Still, it required $20,623 (11%) less inventory than the Z Factors guess.

However, if our goal is to provide exactly 99% service, the Equal Service method (which computes safety stock only to the extent that it is needed above the working stock to achieve the target service) is the only method that actually does what we want. It achieves this for just $136,530 in inventory, $51,075 (27.2%) less than is required by the Z Factors guess.

Lastly, if for budget reasons we need to squeeze every dollar we can out of inventory without sacrificing service, the Optimized $ Fill method is a good one to use. It saves an additional $2,534 of inventory, yet still achieves the aggregate 99% target. True, some SKUs will have lower service (and some higher), but your customers' vendor performance systems will likely report only the aggregate service achieved so your customers should be happy. Compared to the Z Factors guess, this saves $53,609 or 28.6%. That's using the Z Factors guess as the basis, but the better way to look at it is using Optimized $ Fill as the basis. Thus you can achieve 99% aggregate service for only $133,996 in inventory (the absolute lowest inventory possible) using Optimized $ Fill; but if you want to use the Z Factors guess you will need to add 40% more to your inventory. Not many companies would agree to that!

Yes, the last two methods are more complicated to program than the first two, but that's why you buy software, rather than writing your own. (Just as you buy light bulbs rather than manufacturing your own!) Once the logic is in place, the last two methods are actually easier to use and understand than the first two.

<u>Your Opportunity</u>

If we want to look for a silver lining in all of this, I suppose it is that most companies do not use the Z Factors guess; they use the Fixed Time Supply guess instead. If you are doing either, you can raise your customer service (which improves your ability to compete) and raise your ROI at the same time.

We spend a good portion of our time educating forecasters, managers, and executives not to manage their inventories based on days' supply or Z Factors. You can see why. You never want to guess or use inaccurate methods when computing safety stocks. That can kill your ROI and your customer service, and the resulting expediting can create chaos.

Bear in mind that although computing safety stock correctly is quite complex, and involves a lot of statistics; the analyst doesn't need to see (or understand) them. They are best hidden in the software.

Factors Affecting Safety Stock & Service

To compute safety stocks correctly, your calculation needs to include all of the following:

1. Forecast
2. Lead time
3. Forecast error (*not demand variability*)
4. Forecast error distribution
5. Product cost
6. Replenishment frequency (i.e., working stock)
7. Service target
8. Median customer transaction size
9. Selected optimization strategy
10. Supplier lead time variability [2]
11. Supplier quantity variability [2]

Most forecasting tools consider fewer than half of these. Some include just the first one or two. As I said earlier, if you ignore any of these, they will make your work much more difficult, if not impossible.

If you don't include all of these items, then you're just guessing. Let's consider *forecast error distribution* for example. Most textbooks state the assumption that forecast errors are distributed normally, but then leave you with no alternative when that's not the case. Yes, for most companies, most forecast errors are distributed normally, *but not all of them are*. If you assume the wrong error distribution then either your service is too low or your inventory is too high.

Unfortunately for service parts companies and for any products with sparse demand, typically most forecast errors are distributed exponentially, not normally (*but again, not all of them*). If you incorrectly assume the normal distribution you will not reach your service targets. The solution is to use the right error distribution for each individual SKU.

While we're discussing safety stock…recently a client sat in on a briefing of their ERP vendor's forecasting and inventory planning capabilities. I won't embarrass the vendor by naming them…although one would expect from their name that they would know something about predicting the future. When the subject of computing safety stock came up, the presenter asked them (and I quote), "What is the equation you want to use to calculate your safety stock?" And it wasn't like he gave them a list of the possibilities. He was expecting the client to know what to use, because he—the ERP vendor—did not!

The Price for Service

If you do compute safety stock correctly, you end up with a price list for service that always has this general shape:

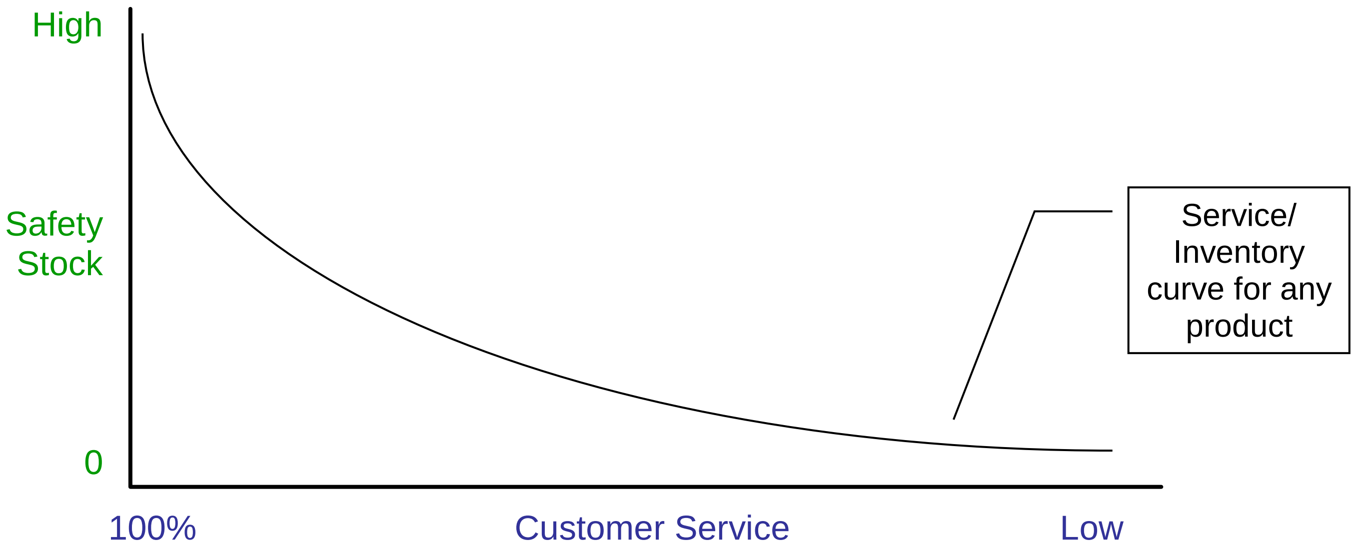

Note here that safety stock runs from low (bottom) to high (top), while service runs from high (left) to low (right). While this tells us what we already knew (good service is expensive, poor service is cheap), notice that as service approaches 100% the cost to achieve it increases exponentially, and as service drops below 90% the cost flattens out. Essentially, on the high service end, very small changes in service are accompanied by very large swings in cost. (Remember the company above where going from 99% service to 100% caused a 50% increase in inventory? This is why. It's not unusual to see it double or triple.) Similarly, on the low service end, very large changes in service are accompanied by very small swings in cost. So on the low end, you can often afford better service than you thought.

Now if you're guessing—perhaps using a fixed time supply or Z factors—you're not actually on the curve. You are way above and to the right of the curve.

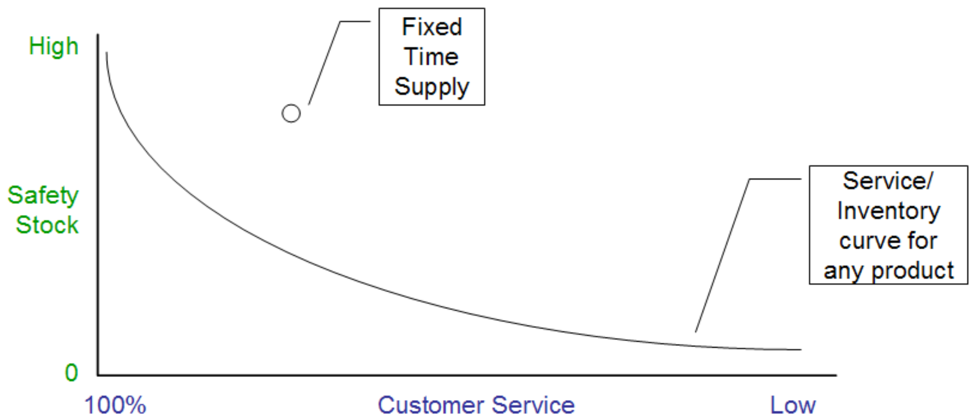

But you could be on the curve—anywhere you want to be. So what do you do? One approach—if inventory is considered to be satisfactory but service is too low—is to take inventory out where you have too much, but then plow it all back in where you have too little. This moves you horizontally over onto the curve, raising your service while maintaining your current inventory level.

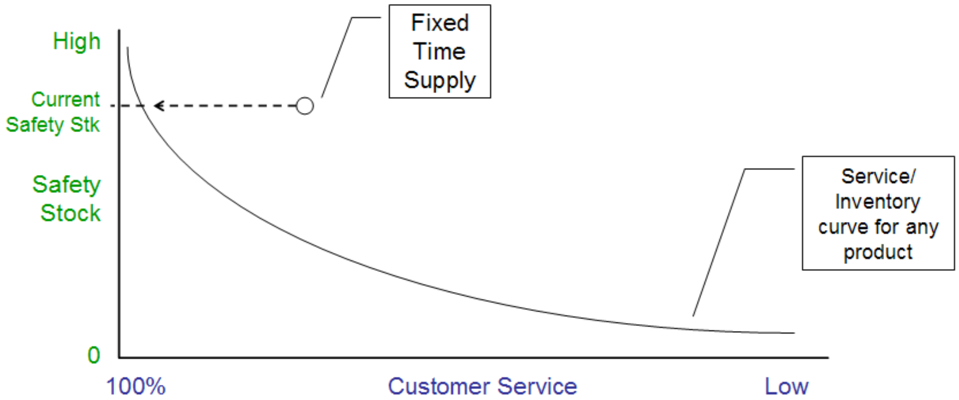

Another approach—if service is good but inventory is too high—is to keep your current service level by removing inventory where you have too much, adding some back in where you have too little; but you will remove more inventory than you add, moving you vertically straight down to the curve.

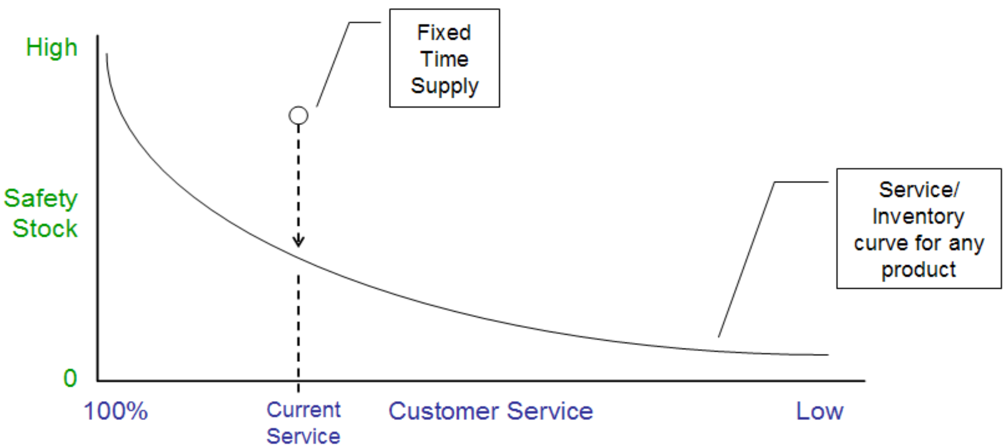

Now you can do anything you want, but what we like to recommend—at least initially—is to do both: raise the service *and* reduce the inventory:

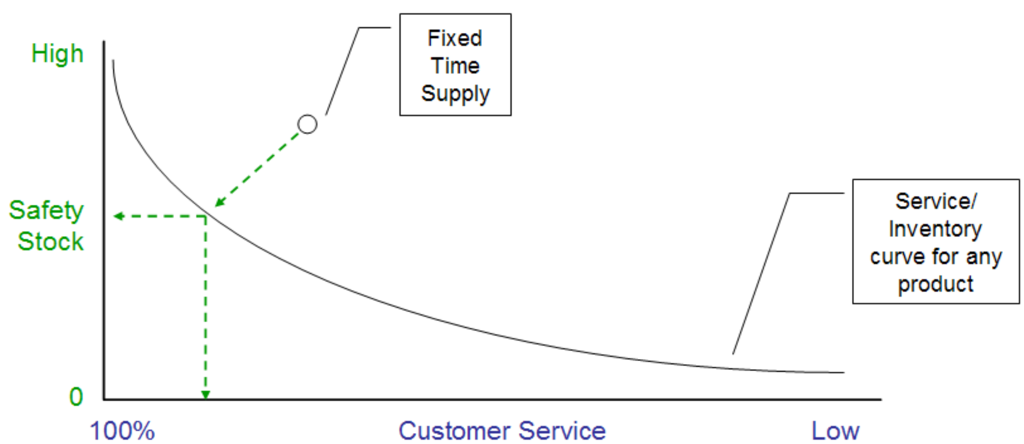

This demonstrates to the entire organization that you now have both service *and* inventory under proper control.

Inventory Planning Key Thoughts

1. It is possible to *know* the relationship between inventory and customer service.
2. You can determine the inventory required to meet your service target.

3. You can determine how much service an inventory budget will buy.
4. You don't have to guess.
5. These same concepts apply to all levels in the bill of materials—not just to finished goods. [3]
6. If you are guessing today, applying this science to your business means you can raise service while simultaneously cutting inventory and expediting, raising your ROI.

Forecasting Fundamentals

1. *Use **demand data**, not sales or shipment data.* You want to base your forecasts on data that tells you what your customer wanted and when he wanted it—*not* on what you had available to ship. So drive forecasts with the requested quantity and date, not the actual quantity and date. Otherwise, if you're in a hole, your forecasts will keep you there.

2. *A forecast is a range, not a number.* So you don't say the forecast is 1,000. Instead you say the forecast is 1,000 +/-100 (at 98% confidence, if 98% is your service target). *And it's time-phased.*

3. *A tighter range means the error is lower.* A smaller error means you need less inventory to reach your service target. For example, 1,000 +/- 75 rather than +/- 100.

4. *A wider range means the error is larger.* A larger error means you need more inventory to reach your service target. For example, 1,000 +/-250.

5. *Actual Demand = Forecast + Error*

   If it's not in the forecast, it's in the error (and must be protected with inventory). But if it is in the forecast, it's not in the error. You may have noticed earlier I indicated that *forecast error is not the same as demand variability*. Your demand can vary widely, but if the forecast accurately models that variability, you don't need safety stock to protect it. And since you also model the forecast errors (via the error distribution), as well as the demand, you can still reach your service target.

6. *Measure forecast error using inventory $, not percentages (MAPE) or obscure statistics like R^2.*

   Do you care about large percentages or large inventory dollars? Percentages lead you on wild goose chases. For example, on which of these parts should you spend time looking for ways to reduce the forecast error?

	% Error
Part A:	10%
Part B:	200%

   You're thinking you want to direct your attention to part B, right? Are you sure? If you think this is a trick question, let me say that it's one forecasters encounter every day. Here is some additional information (assuming your service target is 98%):

	% Error	Inventory $ Required to Achieve 98% Service
Part A:	10%	$10,000
Part B:	200%	$10

   Now where do you want to spend your time? If you spend all your time on part B, the most you can save is $10, while part A is a $10,000 savings opportunity. We say the 10% item is 1,000 times worse than the 200% item. Don't let MAPE (mean absolute percentage error) send you off on wild goose chases!

Four Ingredients of a Statistical Forecast

A statistical forecast can have any or all of 4 ingredients:

1. Level

2. Trend
3. Seasonality
4. Selling Day Adjustment

Level: The first is a Level model—a constant rate of demand per period (week, month, quarter, year, etc.). This is demand that is neither growing nor declining, nor does it have any seasonal variation.

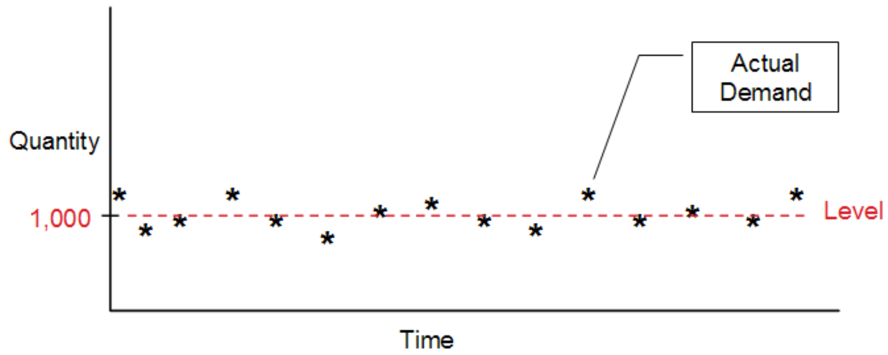

Trend: The second ingredient is a trend, which is the rate of growth or decline per period. Here the trend is positive, which means demand is increasing.

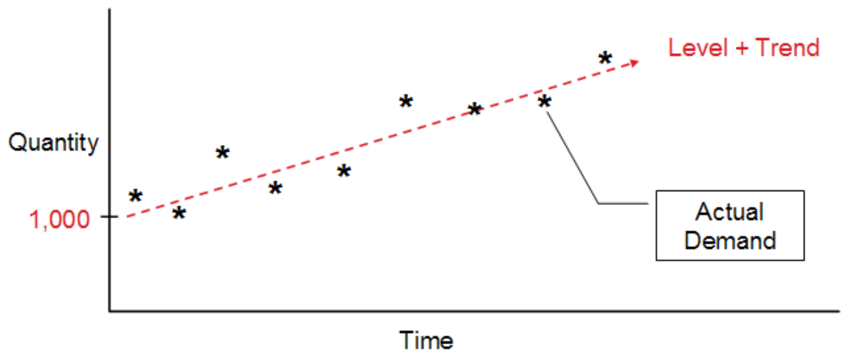

In this example the trend is negative, which means demand is declining.

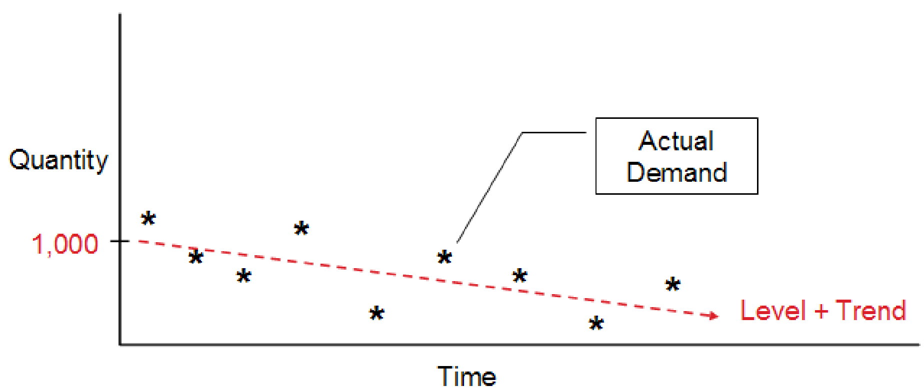

Seasonality: Many products exhibit seasonal patterns of demand. This is a pattern that repeats each year at the same time. This could be an annual pattern, a quarterly pattern, a weekly pattern, or any other pattern *or combination of patterns*.

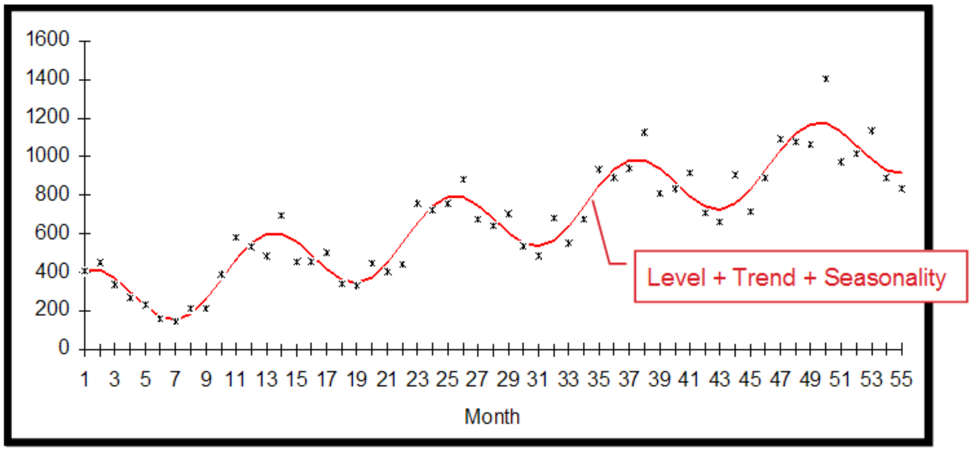

Let me just add parenthetically: some companies, with annual or quarterly sales goals (*or both*), exhibit *management-induced seasonality*—shifting demand from one period to another. This almost always increases the forecast error (which increases the inventory required to meet your service goals, which reduces your ROI) while merely shifting demand from one period to another without increasing overall revenue. The same is also often true of promotions. Anything you can do to reduce or eliminate management-induced seasonality will be greatly to your benefit.

Since forecasters are not statisticians, it's important *not* to make the user specify the seasonality. Instead, the software should be able to determine automatically whether or not there is statistically-significant seasonality, and whether it is annual, quarterly, monthly, weekly, etc., or some combination of these.

The most-frequently-used method for computing seasonality, the **Base Index** method (also called the **Seasonal Index** method or **Winter's** method) has been obsolete since at least 1963, but it is still being used because it is easy to program and to teach programmers how to program it. Unfortunately it has several glaring weaknesses which make it unsuitable for reliably computing seasonality. The latter is especially true for service parts companies or any companies with sparse demands.

Here is a comparison of the Base Index method (lower window) with the far more accurate and statistically rigorous Fourier Seasonal Profiles (also called *harmonic analysis*) method (upper window).

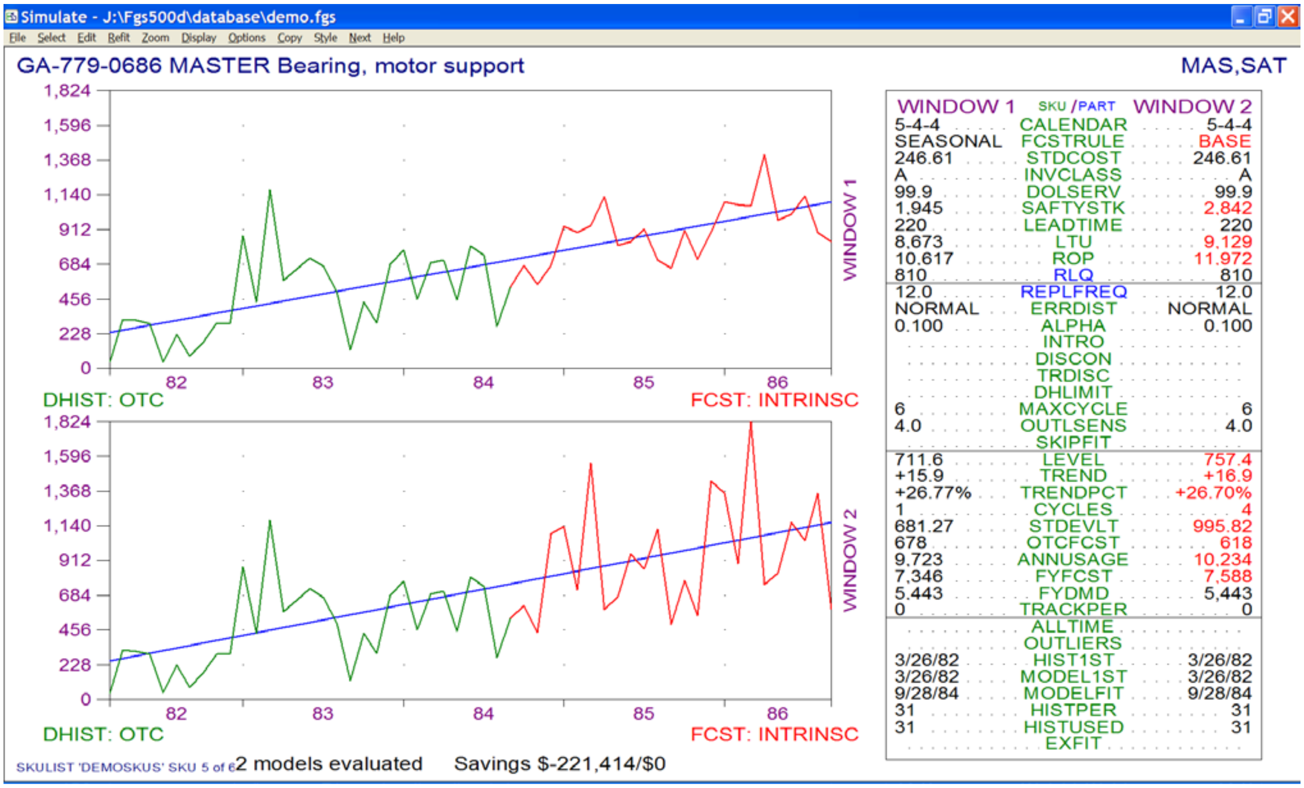

For this service part you can meet your service objective of 99.9% service using the Base Index method of seasonality so long as you as you are willing to spend an additional $221,414 (46% more) to cover the increased forecast error caused by using this obsolete method. Why would you want to do that?

We include the Base Index method in our software just so users can see how poorly it performs (*and because it's so easy to program*), but we recommend the statistically correct Fourier method. Both are equally easy to teach to end users who don't need to know how seasonality is computed any more than a driver needs to know how an automatic transmission decides when to shift gears.

Selling Day Adjustment: In my Operations Research days when I was just out of college, I was asked to forecast the number of insurance claims expected to be filed monthly, for staffing purposes. I made the obvious discovery that more claims were filed in months containing more business days—hardly rocket science!

The same is true when forecasting demand. A month with more selling days will have more demand than one with fewer selling days, unless that happens to be a seasonal peak. Think of a 5-4-4 calendar, for example. One month each quarter will have 25% more selling days than the other months.

Some systems recognize that August has more selling days than February. However it's equally important to recognize that in some years August has more selling days than it does in other years, because in some years August has 5 weekends and in other years 4 weekends (if weekends are non-selling days). The same thing applies to holidays. The alternative to computing a precise selling day adjustment is to maintain additional safety stock to protect against selling day fluctuations, and you don't want to do that. Selling days are known exactly and for years (*actually centuries!*) in advance.

So to recap, a forecast model consists of any combination of the four ingredients: Level, Trend, Seasonality, and Selling Day Adjustment. (The latter is always present.) The software should automatically determine which is appropriate for each SKU.

Surprisingly, many tools use only a level, for example a moving average or single exponential smoothing. That's fine if your products are neither growing nor declining, and exhibit no seasonality. And your forecast periods are all the same length. I've never seen such a situation, however, and I've been in this business for over 40 years.

Marketing Intelligence

There is an optional 5th ingredient to the forecast—and that is a manual override (what we call *marketing intelligence*). You don't want to do this on every SKU. In fact, if you're entering marketing intelligence on more than about 10% of your SKUs, you're most likely wasting your time plus making the forecast worse. Use this tool only when the forecaster knows something that is not evident in the demand history. If you're just guessing, you will make the forecast worse most of the time.

Never use marketing intelligence without afterwards evaluating it to see whether it helped or hurt, and by how much (in dollars). Without this feedback loop, your forecasters will never learn to use this capability wisely, and your forecast errors will most likely increase.

Two Forecasting Processes: Model Fitting and Periodic Forecast Revision

It is important to understand the difference between the *initial model fitting process* and the *periodic forecast revision process*. The former gets us to a good starting point, while the latter is important for keeping the model up-to-date and warning us about suspected changes in the market.

Most products have demand that is relatively stable. *Stability* here means that if a forecast model truly represents the underlying demand for a part, then that model will be effective over some period of time. You would not expect that one model would work one month, and a totally different model is required next month. There are certainly changes, but they are most often *incremental* changes, not *fundamental* changes. For example, you would not expect to see a product with a level, trend, and annual seasonality this month, go to just a level and trend next month, and then to semiannual seasonality the month after that.

The smoothing and error tracking that occurs in the forecast revision process has two purposes. The first is to make those incremental (not fundamental) changes which keep the model up-to-date with reality. This handles situations such as a trend which is gradually flattening or seasonality which is becoming less conspicuous. The second purpose of the revision process is to identify those SKUs where the chosen model is suspected of no longer working (i.e., identifying fundamental, not incremental, changes). This means that you are alerted when the pattern of demand changes. Knowing that, you can investigate the cause of the change, be it new competition, product changes, or whatever. The fact that a change has occurred or is suspected of having occurred is the key. This is counter to the mistaken belief that one should try every model on every SKU every month.

The *model fitting* process starts with the historical demand. It assumes initially that history will repeat itself (but identifies where it doesn't). It finds the mathematical model that best fits the data, computing the level, trend, and seasonality, as well as the initial forecast error. It employs many heuristics (some call it "expert system logic") to detect the various exceptions that happen in the real world, respond to them appropriately, and present them to the user for review. This includes identifying and handling outliers, pattern changes, recommending a different forecast calendar, etc.

The *periodic forecast revision process* applies Statistical Process Control (SPC) techniques to the *process of generating an accurate forecast*. This process starts by importing the just-completed period's demand and checking it for reasonableness. That way you catch errors before they perturb your forecasts. Then, using the most-recent demand (and a bunch of linear algebra), it revises the model (or tweaks it) to keep it current. These are slight changes, not fundamental changes. Then it alerts you to those products for which the market has possibly changed and evaluates the marketing intelligence. Any other exceptions (such as user-defined exceptions) are then identified and presented to the user for review.

One problem with not having the periodic forecast revision process is that you forego the learning benefit that takes place. As GE Silicones reported at APICS 1991 [4] (and their experience was typical), during the year after the initial model fit, the forecast error—and therefore the safety stock—reduced another 30% on top of the initial reduction. Another problem with having no forecast revision process—which means you must refit models every period—is that everything changes every period. This makes management by exception impossible. (Note also that the Base Index method for computing seasonality cannot be revised, but can only be refit—which increases the forecast error and again, rules out management by exception.)

Very few forecasting tools have both of these processes. Most have just the model fitting process, and many omit critical pieces (such as heuristics and exception identification).

Persistent Forecasting & Inventory Planning Myths

Almost every client we've worked with has told us up front, "You can't forecast our business." We subsequently discovered what they were actually saying was "*They* couldn't forecast their business!" *We* could forecast it just fine. That's because they believed one or more of these popular and persistent demand forecasting and inventory planning myths. These myths have bedeviled forecasters and forecasting systems for decades, some going back to the 1950s—and largely still do today. Many of them, I am embarrassed to say, are still taught by APICS. If your procedures or systems incorporate *any* of these myths you are dooming your operations to mediocrity—at best.

Some of these I've already mentioned above.

1. **You can't forecast our business!** *(see above paragraph)*
2. **A forecast is 1 number (e.g., average monthly demand).** *(see Forecasting Fundamentals item 2)*
3. **Safety stock provides service, but working stock does not.** *(see Four Methods for Computing Safety Stock)*
4. **Forecast error = demand variability.** *(see Forecasting Fundamentals item 5)*
5. **Forecast error requires MAD or MAPE (percentage) or R^2.** *(MAPE: see Forecasting Fundamentals item 6)*

   The correct measure of forecast error is the standard deviation. But in the 1950s, Robert G. Brown was trying to calculate safety stocks using an IBM 602A Calculating Punch (a more reliable version of the IBM 602) and it took 45 minutes—an enormous amount of time—to compute a square root (which is required when computing the standard deviation). Bob told me a person could actually walk to the canteen, have a cup of coffee and return before his computer finished calculating a single square root. It was so slow, in fact, that it would take 6 weeks to recompute the monthly forecast errors!

   The IBM 602 Calculating Punch (1946)

   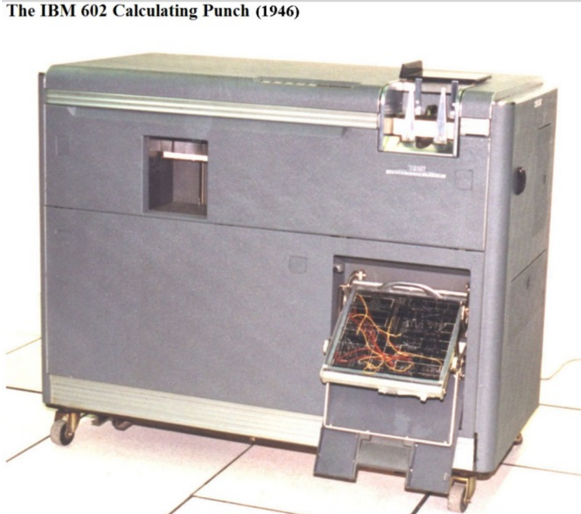

   That was obviously impractical, so he invented the Mean Absolute Deviation or MAD as a *temporary* work-around. He did some statistical analysis at the time and convinced himself that the standard deviation was approximately equal to 1.25 times the MAD.

   Since then, two things have happened: (1) computers have long since been able to compute square roots in less than a microsecond, and (2) Brown ran subsequent studies (once he had access to a faster computer) that showed the standard deviation could be anywhere between one and two times the MAD, making it wildly inaccurate and therefore completely inappropriate for calculating safety stocks. *Today many many thousands* *of companies keep a total of many billions of dollars in unnecessary inventory because they are computing incorrect forecast errors and safety stocks using an algorithm that was necessitated by technology introduced in 1946 (70 years ago)!*

Brown's comment was that when viewed from the psychologist's stimulus-response model, a response that continues long after the stimulus has ceased is the definition of a neurosis, and that therefore *MAD is neurotic!*

The only companies today with an excuse for using MAD are those who are still calculating their forecast errors on the IBM 602A Calculating Punch, and even then, their results will be wrong.

If you don't believe this is still a problem...as recently as just 2 or 3 years ago a reader wrote to the advice column of a national trade magazine asking how to compute safety stock. The answer—in a stunning triple threat to his operations—was that he should compute the *MAD*...of his *demand variability* (*not forecast error*)...times 1.25...times the *Z Factor*! We need a Hippocratic Oath for Operations Management!

Can we finally bury MAD and remove it from all textbooks and software? Not just because it's obsolete, but because it doesn't work and it never did.

6. **Standard deviation = 1.25 x MAD.** *(see item 5 above)*
7. **Replenishment frequency = forecasting frequency.** *(see Four Methods for Computing Safety Stock)*
8. **All forecast errors are normally distributed.** *(see Factors Affecting Safety Stock & Service)*
9. **All forecast calendars must be monthly.**

Some years ago an article in a monthly trade journal stated that all companies should forecast all their products on a monthly calendar. The reasoning given was that forecasting weekly was too frequent for some companies, while forecasting quarterly was too infrequent for others. Therefore everyone should use months!

Let's put that statement into the proper perspective by applying the same logic to everyday life—one's shoe size, for example: Here is an equivalent statement: "It has been observed that size 5 shoes are too small for some people, while size 15 shoes are too large for others. Therefore it is decided that everyone should wear size 10 shoes." The obvious response from every sane person would be, "Why not let each person wear whatever size fits best?"

The same is true with forecasting calendars, i.e., put each SKU on the forecasting calendar that works best. Computers are fast enough today that it's not necessary to treat all SKUs alike; instead we can treat each one *appropriately*. As a general rule, fast-moving SKUs are more accurately forecasted using more frequent calendars, such as weekly or biweekly calendars. Slow-moving SKUs generally do better on less frequent calendars, such as bimonthly, quarterly, semiannual, or even annual (one forecast period per year).

Here's an example of a service part with monthly demand ranging from the single digits to triple digits. Switching it to a quarterly forecast calendar allowed the model to notice the annual seasonal pattern, yielding a forecast whose error is reduced by almost 40%. The safety stock reduces by the same percentage, a savings of $17,300. Since this is now forecasted 4 times per year instead of 12, the workload is reduced by two thirds! There are few situations in life where the best solution is also less work, but this is one of them. According to Jeff Beck, former forecasting/inventory analyst at GE Aircraft Engines, using selectable forecast calendars in their service parts operation resulted in a 43% workload reduction while also reducing their forecast errors. [1] and [5]

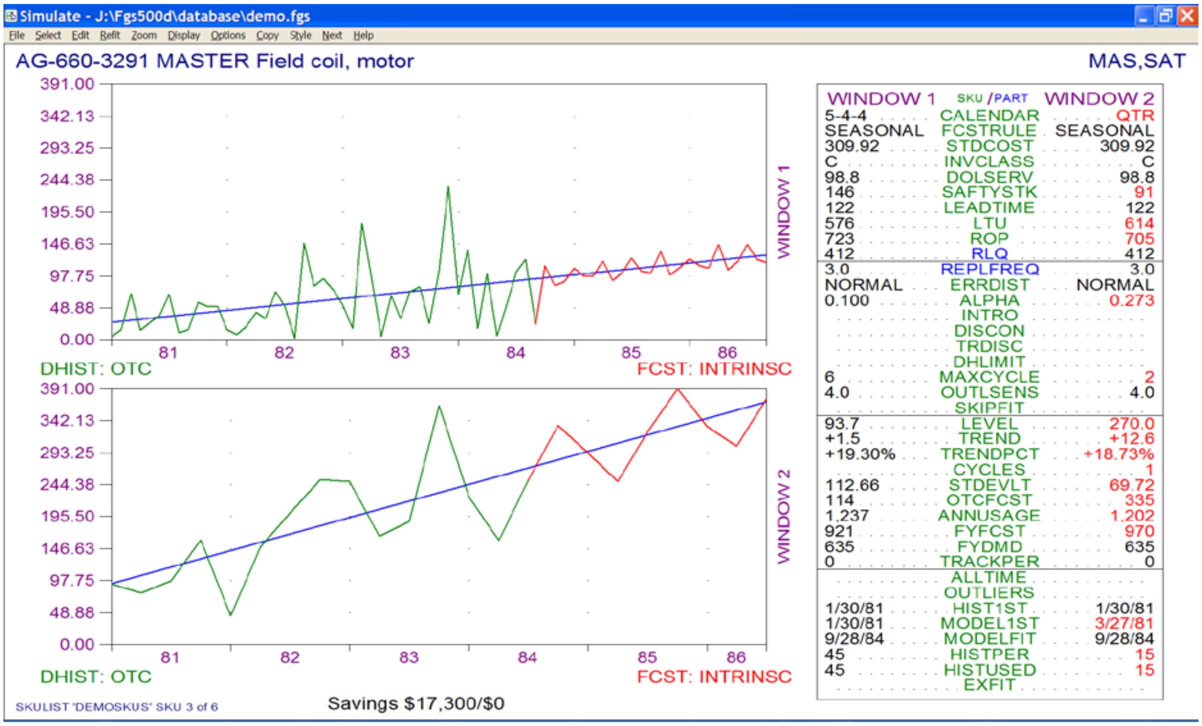

However, in order to do this, you cannot lump up your demand history to months! For example you cannot forecast an item on a weekly calendar if you loaded your demand history in monthly buckets. To do this properly, you should load your demand history at the individual transaction level—just like it's stored in your ERP system. Then your forecasting system can automatically sum it to weeks, months, quarters, etc. internally as it needs to compute a model on the best calendar, leaving it at the transaction level in the database in case you later decide to change the forecasting calendar.

This *should* make interfacing to ERP systems easier, but the same ERP vendor that asked their client how to compute safety stocks, required that demand history be lumped up to months to import it into their forecasting module. And they required *shipment history*, not *demand history*!

One additional advantage to less frequent calendars is that they are often accompanied by a change from the exponential error distribution to the normal error distribution, and that requires less safety stock for the same service target.

10. **Seasonal (base) indices compute accurate seasonality.** *(see Seasonality under Four Ingredients of a Statistical Forecast)*

11. **Forecast model fitting and revision processes are identical.** *(see Two Forecasting Processes: Model Fitting and Periodic Forecast Revision)*

12. **We must override every forecast.** *(see Marketing Intelligence)*

13. **We must not override any forecast.**

This is the opposite extreme. We often see it when the IT department owns the forecasting process and no one in the company pays much attention to the forecasts. We call it the *black box* approach: no one reviews the forecast exceptions. While you can do this if you want, it is an extremely expensive method. For example, one of our clients discovered that by using management by exception with the forecasting process, they were able to save $10,000 of inventory for every hour they spent reviewing exceptions. [4] (Coincidentally, the analysts performing the exception reviews were paid considerably less than $10,000 per hour.)

14. **Use fixed time supplies or Z Factors to compute safety stocks.** *(see Four Methods for Computing Safety Stock)*

15. Forecasters need no education on forecasting.

As I mentioned at the beginning, demand forecasters and inventory planners have a huge impact on their employer's fortunes—yet this is seldom recognized by others. And that, unfortunately, often includes those responsible for making sure that practitioners are properly educated and trained to do their jobs.

It is absurd to have to state this, but forecasters and inventory analysts need to be educated in order to do their jobs. This is not something they are born with, nor something they learn in college—even those with degrees in operations management. Nor do they learn it as part of their APICS certification. It doesn't take a lot of money or time—we educate our users in just eight days, divided into two workshops. Many times we have seen users return from the workshops and in the first day back on the job reduce inventory (while maintaining service) by more than the cost of the workshops plus their travel expenses!

Some managers will say they cannot afford to have a forecaster away from the office for a week. (One wonders what they do when vacation time comes around.) Yet we find that educated users use their time far more efficiently than uneducated users.

To illustrate my point about the importance of education, let me give you an example. The service parts operation of a manufacturer of major consumer appliances approached us. Our analysis showed significant improvements were available.

We educated their manager and users and they were off and running—with good results beginning to show. But then disaster struck. The manager retired and the lead user moved away. We encouraged them to train their replacements, but to no avail. One by one the rest of the users left as well. Their replacements were also not trained—there was no budget for training.

Things stumbled along for a while with just minor improvements but then they began going downhill. Every time management turned over we encouraged them to let us train them and their users. The answer was always no. There was no budget. Then one day a sharp gal called me up and said, "I've just been handed this group—and the first thing we need to do is to train everyone!" After I picked myself up off the floor, we did just that. And, they started quarterly visits by our instructor to review various challenges they faced.

The results were entirely predictable: service climbed; inventory and expediting fell. After a year or two they were awarded the Sears award for the top service by a vendor. This is a big deal! The Sears award is very hard to win! The manager told me they wanted to erect a statue of our instructor outside the entrance to their corporate headquarters. Since then they have won the Sears award for 3 years in a row! That's never happened before! The client tells me their service is higher than it's ever been, and their inventory is at an all-time low.

You never want to scrimp on education!

A closely related myth is that clerical people can handle this job. The reason that forecasters and inventory planners are called analysts is that their tasks are *analytical*, not *clerical*. They will be successful only if they are analytical people. Yes, analytical people are more expensive than clerical people; but I've seen a forecaster save many times the difference in salary with a single decision on one part. A good way to select such people is to hire those with degrees in a quantitative field.

My ERP System Already Does All of This, Right?

Unfortunately, no, it doesn't.

Why not? Consider this: nobody can be an expert at everything. ERP vendors know this so they ask themselves, "Who buys ERP systems?" The answer is usually, "The CFO". So the vendors make certain the financial functions are as good as they can possibly make them. The rest of the functionality is much less important—often just barely enough to allow them to check off the box on the clients' checklists. This is the *wide-but-not-deep* approach to systems development. Our clients tell us the forecasting and planning functions are by far the worst facilities in their ERP systems.

We saw one ERP system that offered both "Demand Forecasting" and "Inventory Planning". The first was a place in the database where the users could type in any forecast they wanted. Similarly, the Inventory Planning "module" was a field in the database where the users could enter their desired safety stock.

ERP systems and Forecasting & Inventory Planning systems are fundamentally different. ERP systems are *transaction processing* systems. Forecasting & Inventory Planning are *analytical* systems. Why does this matter? The two require developers with different educations and different skills. They also require different organizations, different tools, and even different software development processes.

Transaction processing systems are concerned with the past and present, keeping track of ins and outs. The most complicated math they use is multiplication and division. Forecasting & Inventory planning look to the future. They deal with what should be or could be—not just what is today (i.e., the status quo). Such analytical systems involve complex probability and statistical analyses, advanced calculus, and engineering optimizations.

No organization that excels at one can excel at the other. *But that doesn't mean their systems can't exchange data with each other.*

The Single System Syndrome

Many companies suffer from the mistaken belief that having just a single system will make their lives simpler and better. This means that all their software must come from a single vendor. Our clients call this the dreaded *single system syndrome*.

This approach doesn't work in any other aspect of life. For example, can you imagine buying all your clothes from a single vendor? Florsheim doesn't even make underwear! And it's probably a good thing they don't! It's surprising that *anyone* would expect this approach to work with software.

Most companies suffering this syndrome probably do it out of fear of interfacing, yet there are many good standard interfacing tools available. Our system has them built in. (Our users usually say afterwards that interfacing was a non issue.)

The Single System Syndrome dooms companies to mediocrity, or often outright failure. In such a situation, the best they can do is to hope that their competitors do likewise. Let me illustrate this with one heart-breaking example.

Many years ago we had a client that installed our software and trained one user. But to reach this point we first educated management in the science of demand forecasting and inventory planning. By the end of the first year after implementation, inventory dropped 50% while service rose 50% to the high 90s. Expediting—which had been debilitating before—dropped to almost nothing. The results were so dramatic that the forecaster was given a management award which included a $15,000 bonus—and this was over 25 years ago when $15,000 was a lot of money! Plus it happened at a company that was not known for being excessively generous to its employees.

Now fast forward 15 years: management had turned over several times and since the company refused to educate their new managers in this science, they consequently no longer understood the science of demand forecasting and inventory planning. They fell prey to the Single System Syndrome and declared that all software had to come from their ERP vendor. So they installed their ERP vendor's advanced planning and scheduling module (at a cost that was 40 times higher than our software).

What happened? Things should have gotten 40 times better, right? Not quite. Service fell dramatically; inventory doubled and then continued to grow. Expediting exploded. The interesting thing was that the user time required to operate the system rose 1,000% from 2 days per month to 20 days per month! Rather than admit their error, the company stumbled along producing poor results. Eventually their results were so bad that their parent company sold them off. This company has paid literally hundreds of millions of dollars unnecessarily because of this misguided policy. And they will continue to do so in the future.

When I look at corporate websites for manufacturers and distributors I see words like "excellence" and "high quality" and "unmatched service" appearing in their mission statements and goals. I've never seen a company say they strive for *mediocrity* (or worse). Yet mediocrity, *at best*, is exactly what the single system syndrome achieves. Achieving excellence, high quality, and unmatched service requires a good working knowledge of the science of forecasting and inventory planning. And that requires education. It also requires tools which embody the science of forecasting and inventory planning.

About the Author

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, statistics, operations research, and electrical 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 demand forecasting and inventory planning software.

E/Step Software Inc.
12015 Summitview Road
Yakima, WA 98908

Tel. 509-853-5000
Email: John.Estep@EstepSoftware.com
Web site: www.EstepSoftware.com

References

[1] Beck, J., "SPC and Selectable Forecast Calendars Work Magic for GE Aircraft Engines' Service Parts Operation," *APICS 38th Annual International Conference Proceedings*, APICS, 1995, pp. 477-481, available at http://www.estepsoftware.com/BeckArticle.htm.

[2] Boyd, N. & J. Estep, "Managing the Impact of Supplier Lead Time & Quantity Variation," *APICS 49th Annual International Conference Proceedings*, APICS, 2006, available at http://www.estepsoftware.com/papers/ManagingTheImpactOfSupplierLeadTimeAndQuantityVariation.pdf

[3] Boyd, N. "Safety Stock: Finished Goods, Intermediates, or Raw Materials? That is the Question", *APICS 48th Annual International Conference Proceedings*, APICS, 2005, available at http://www.estepsoftware.com/newsletter/SSInLean.pdf.

[4] Duncan, R. and J. Estep, "Simulate Your Way to Exceptional Forecasts," *APICS 34th Annual International Conference Proceedings*, APICS, 1991, available at: http://www.estepsoftware.com/articles/SimulateYourWayToExceptionalForecasts-APICS1991.pdf

[5] Estep, J.A., "A Simple Technique for Reducing Forecast Errors," *APICS 30th Annual International Conference Proceedings*, APICS, 1987, pp. 291-293.

[6] Greene (Ed.), *Production and Inventory Control Handbook* (2nd ed.), APICS, 1987, chapter 29.