

A simple definition of a finished-goods item is one that is in a saleable or shippable state within a given environment. This means that the same item will be classified differently based upon its location within the supply chain. For example, castings produced at a foundry would be a finished good at the foundry, yet when the casting arrives at a machine shop it becomes an unfinished good. The machine shop may machine the casting, at which point the machined casting becomes a finished good at the machine shop. Once the machined casting arrives at a final assembly plant, it is once again an unfinished good. However, if the same business owns the foundry, the machine shop, and the final assembly plant, they will likely only consider the final assembled item to be finished goods, and all other forms as unfinished goods.

Sometimes an item can be both a shippable item and an item used to produce other items within the same facility. For example, a bicycle manufacturer may separately sell wheels, forks, frames, pedals, and other components, as well as use them to produce complete bicycles. Should these saleable components be considered finished goods or unfinished goods? Actually, they could be either or both. If the saleable component usage is very small, they will likely be considered unfinished goods. If the saleable component usage is greater, they may be considered finished goods, or you may set up an additional category for saleable components. You may even set up separate item numbers for the saleable portion of component inventory.

What is important to note here is that the finished-goods classification is essentially a management tool. How you classify an item should be based upon how you intend to use the classification in business decisions. You need to develop a consensus within your organization about what management wants to consider a finished good, and then make sure you are consistent in applying this classification.

### **Unfinished goods.**

Once you figure out what your finished goods are, you will likely also have a pretty good idea of what your unfinished goods are. Unfinished goods are items that are used to produce finished goods items. Unfinished goods are often called components, ingredients, raw materials, semi-processed materials, and subassemblies. Each of these names is somewhat descriptive of groups of unfinished goods, however, as with other classifications, these can also be interpreted differently in different environments. Let me give some brief descriptions of each.

**Components.** These would usually be items that are used in an assembly process. A rim, spokes, and a hub, would be examples of components used to assemble a bicycle wheel.

**Ingredients.** These are generally bulk materials that are mixed or otherwise processed to create another product.

**Raw materials.** These are materials that will undergo some type of a physical change. Examples would include bar stock that will be machined, sheet

***These are sample pages from the book *Inventory Management Explained*.  
More information is available at <http://inventoryexplained.com>***

### Forecast bias measurements.

Forecast bias describes the tendencies of forecasts to be high or low. Whereas the mean absolute deviation calculation ignored positives and negatives, the forecast bias uses these values (signed errors) to determine if your forecasts are more likely to be high (a positive bias) or low (a negative bias) and to what level they are high or low. There are several ways to calculate forecast bias.

**Cumulative forecast bias (Sum of forecast errors).** To calculate a cumulative forecast bias you simply add up each period's forecast error for a series of forecast periods. In the previous example where we had forecast errors of -20%, +10%, and -6% for the three periods, we would sum the three values and get a result of -16%.

**Average (mean) forecast bias.** Forecast bias may also be measured as an average (rather than a sum) of forecast errors. In the previous example, that would result in -5.333%

**Exponentially smoothed forecast bias.** And, of course, you can exponentially smooth the forecast bias. This calculation is exactly the same as the exponential smoothing calculation used for forecasting, except you will be using the forecast error in place of sales (demand).

### Tracking signal.

Tracking signals are important parts of many forecasting systems. The basic idea here is to create a calculation that describes the overall health of the forecast relative to trend. The tracking signal can then be used to initiate changes to the forecasting technique or parameters. Though I've encountered several ways to calculate tracking signals, I have found the most useful to be dividing the absolute value of the forecast bias measurement by the multi-period amplitude measurement. This would include either dividing the absolute value of your average forecast bias by your mean absolute deviation, or dividing the absolute value of your exponentially smoothed forecast bias by your exponentially smoothed absolute deviation.

To use the previous example where we had forecast errors of -20%, +10%, and -6% for three periods, resulting in an average forecast bias of -5.333% and a mean absolute deviation of 12%, we would divide 5.333 (the absolute value of your bias) by 12 to get .444. So what does .444 tell us? It's basically a ratio that describes the relationship between the bias and the forecast error. The higher the number (the closer this number is to 1), the more likely it is that you are not adequately accounting for trend. A lower number (closer to 0) suggests that the forecast error is more likely the result of noise or an inadequate seasonality adjustment.

In the next chapter, we will cover the use of a tracking signal in adaptive smoothing.

***These are sample pages from the book Inventory Management Explained.  
More information is available at <http://inventoryexplained.com>***

|   | A                       | B                   | C |
|---|-------------------------|---------------------|---|
| 1 | <b>Historical Data</b>  |                     |   |
| 2 | <b><u>Predictor</u></b> | <b><u>Sales</u></b> |   |
| 3 | 1054                    | 54                  |   |
| 4 | 1502                    | 75                  |   |
| 5 | 2007                    | 98                  |   |
| 6 | 1375                    | 60                  |   |
| 7 |                         |                     |   |

**Figure 3K.** Set up spreadsheet as shown.

Set up your spreadsheet as shown in Figure 3K.

In Excel 2003: Click Tools>Data Analysis>Regression.

In Excel 2007: Click Data>Data Analysis>Regression

You'll now see the Regression popup window. Click the field labeled "Input Y Range:" and then click and drag through the values for the variable to predict.

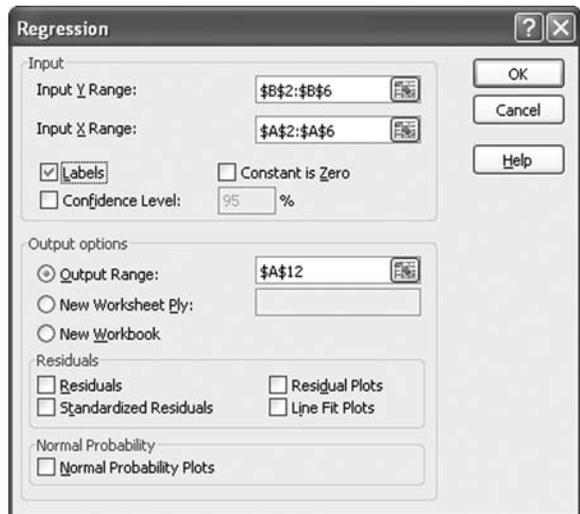
Next click the field labeled "Input X Range:" and click and drag through the values for the predictor variable.

If you included the column headings as you clicked and dragged through the values, you need to click the "Labels" check box.

Now click the "Output Range:" check box, and then click in the field next to it. Now click a cell in an open area of your worksheet (we'll click cell A12).

The Regression popup window should appear as in Figure 3L.

Click OK.



**Figure 3L.** Regression popup window

***These are sample pages from the book Inventory Management Explained. More information is available at <http://inventoryexplained.com>***

the percentage change between December and January revenue for your business.

Let me once again mention that many of you will probably never need to break down trend into its various elements for your forecasts. I felt it was important to cover this for those who may need to go to this level, but also to help the rest of you to develop a better understanding of the complexities of trend.

## **Exponential smoothing revisited.**

---

We've already covered simple exponential smoothing, but forecasters have taken exponential smoothing far beyond simply smoothing historical demand. Terms like double exponential smoothing and triple exponential smoothing are somewhat misleading since they seem to imply you are somehow re-smoothing the historical demand over and over again, when in reality you are applying the smoothing technique to additional elements of the forecast (though some re-smoothing does occur in some methods). So rather than focusing too much on catchy names for particular variations, I'm going to show the logical progression in using smoothing in forecasts. By the way, these same concepts can be applied to other smoothing techniques (moving averages or weighted moving averages).

### **Smoothing trend.**

So you have your simple exponential smoothing calculation all put together, but you realize your forecast always lags in trend. What do you do? You add trend to the calculation. Since trend is a gradual change in demand over time, we can look to quantify trend as the change in demand from one period to the next. However, since actual period-to-period demand variability tends to consist of a lot of noise, we need to find a way to eliminate as much of this noise as possible. And how do we do that? Well, we smooth it, of course. But, instead of taking the difference in actual demand from one period to the next and then smoothing it, how about we take the difference in the forecasts created from one period to the next (which is already smoothed to some extent) and then smooth that? Well that would make it double smooth wouldn't it? Hmm, I wonder what would be a good name for something like that?

In other words, we are taking the results of our simple exponentially smoothed forecast, then looking at the difference between the previous period's forecast and the next period's forecast, and then applying a smoothing factor to that to get a smoothed trend, which we can now add to the forecast to better account for trend. Essentially we are running two separate-but-related exponential smoothing calculations (one for the base demand forecast, and one for the trend forecast) and then combining the results for our finished forecast. So if our basic exponentially smoothed forecasts (note this is forecast, not actual demand) for the last two periods were 50 and 55 respectively, and our next period's basic exponentially

***These are sample pages from the book *Inventory Management Explained*.  
More information is available at <http://inventoryexplained.com>***

you may need to consider when calculating safety stock.

### **Oversimplified use of statistical methods.**

---

There are a couple of very similar statistically-based safety stock calculations that you will likely encounter in many inventory books or courses. The basic logic behind these calculations is sound (sort of), however, these calculations and the associated logic is generally communicated in an oversimplified manner and therefore doesn't provide enough information to properly use them.

The two statistical safety stock methods you encounter most often are actually variations of the same basic statistical model. That is, they apply the Normal Distribution statistical model to safety stock. The inputs to the calculations include a calculation of the Standard Deviation or the Mean Absolute Deviation (MAD) of historical demand, a desired service level, and a table that associates a service level with a multiplier. You just calculate either the standard deviation or the MAD, then multiply the result by the multiplier from the table that corresponds with your desired service level (each model has its own table). The result of this calculation is a safety stock level that should meet your desired service level. Or is it?

The answer to this question will become obvious as we progress through this chapter. In fact, the remainder of this chapter is based entirely on this single model, but rather than taking the mucho popular keep-it-simple approach, I'm actually going to try to explain the logic of the model and how that logic can be adapted to inventory management. In addition I will provide some of my own little adjustments (mostly cheats) and try to explain the logic behind them, and in the process hopefully get you to the point where you can put together a safety stock calculation that works for you. Be forewarned though, things are going to get really complicated here.

### **Appropriate use of statistical methods.**

---

I guess it makes sense to start out by addressing the question of why we would use a statistical model to calculate safety stock. To risk oversimplifying things a bit myself, statistics are used to interpret and explain data, but more importantly for our purposes, statistics can be used to assign probability based on a sampling of data. Since what we want to do is use our historical demand (a sampling of data) to help us determine how much inventory it would take to compensate for demand variability (the probability that a certain quantity of safety stock will cover our needs), it would seem that the use of statistics would be a perfect fit.

Let me just throw in a disclaimer here. My knowledge of statistics is rather limited (I am in no way an expert), and what I am trying to do here is share with you what I think I know about statistics. And more specifically, what I think I know about the statistical model we will be using to calculate safety stock. It's

***These are sample pages from the book *Inventory Management Explained*.  
More information is available at <http://inventoryexplained.com>***

## Back to the calculations.

Let's go step-by-step setting up some order quantity calculations in a spreadsheet.

The basic EOQ calculation is really simple.

| B7 |                       | fx | =SQRT((2*B2*B5)/(B3*B4)) |
|----|-----------------------|----|--------------------------|
|    | A                     | B  |                          |
| 1  | <b>Basic EOQ</b>      |    |                          |
| 2  | Order Cost            |    | \$13.00                  |
| 3  | Carrying Cost         |    | 20%                      |
| 4  | Unit Cost             |    | \$5.25                   |
| 5  | Annual Demand         |    | 300                      |
| 6  |                       |    |                          |
| 7  | <b>Order Quantity</b> |    | <b>86</b>                |
| 8  |                       |    |                          |

**Figure 7D.** The basic EOQ calculation

In cell B2 enter your order cost. For this example, we will use \$13.00.

In cell B3 enter your carrying cost. For this example, we will use 20%.

In cell B4 enter your unit cost. For this example, we will use \$5.25.

In cell B5 enter your annual demand. For this example, we will use 300

In cell B7 enter =SQRT((2\*B2\*B5)/(B3\*B4))

The result is an EOQ of 86 units. That's all there is to it. You now have a working EOQ calculation and can change the values in cells B2 through B5 to reflect your own examples.

Now let's try testing our output a little using another calculation.

In cell B9 enter 86 (the result of our EOQ calculation)

In cell B10 enter =(B3\*0.5\*B9\*B4)+(B2\*B5/B9)

What you've just entered in cell B10 is the calculation of the expected total annual cost given the order quantity in cell B9. If you recall when we first started solving the EOQ formula, we had the following calculation that represented annual carrying cost on the left side as being equal to annual order cost on the right side.

$$c \cdot 1/2x \cdot u = p \cdot a / x$$

We can simply move things around a little to get a calculation for total annual cost. So instead of having annual carrying cost equal to annual order cost, we have a calculation that represents total annual cost as being equal to annual carrying cost plus annual order cost.

$$\text{Total annual cost} = (c \cdot 1/2x \cdot u) + (p \cdot a / x)$$

**These are sample pages from the book *Inventory Management Explained*.  
More information is available at <http://inventoryexplained.com>**

The truth is, companies that have historically used min-max probably did not use a cost-based order quantity calculation. Instead, they used a max because it represented the physical amount of shelf space they had for that item, or because it represented some arbitrary quantity they felt was the most they wanted to have, or it represented a quantity that was equal to the demand for a fixed period of time. If we go back to the “meat guy” example from the section on periodic review (earlier in this chapter), where we had a fixed delivery schedule (once a week) and placed an order each Monday for a Thursday delivery, I may set my max to represent a 10-day supply plus safety stock (the quantity I would need to get me to my next delivery). This way, I would be planning on ordering this item each week, and planning on ordering just enough to get me through until my next opportunity for ordering. Oddly enough, this scenario would result in a min-max where the min and max are the same quantity.

I think the meat example is a good example of where min-max is applicable. You’ve got a group of short-lead-time items with a regular delivery schedule, where it doesn’t make a lot of sense to order more than the quantity to get you through to the next delivery (order costs in a scenario like this are generally very low, so a small lot size makes sense).

### **Kanban, Two-bin, and that whole Push-Pull thing.**

---

If you’ve worked in manufacturing, and especially if you’ve been exposed to “Lean Manufacturing”, “Just-In-Time”, or the “Toyota Production System”, you have probably encountered the term Kanban. Kanban is a Japanese term that means signal or card or signal card or sign or signboard or . . . hell, I don’t know, it’s not like I speak Japanese.

But before we get to Kanban, I want to briefly cover an ancient, but still used, reordering method known as the two-bin system. And the really nice thing about the two-bin system, is that the name pretty much says it all. Each item has two bins (or other type of container), and when one bin is emptied, it is sent somewhere to be refilled. This would most commonly be used for internal replenishment where the empty bin would be sent either to a production area where the needed item is manufactured, or to a reserve storage area where bulk quantities of the item are stored. In reality, we should call this a multi-bin system since you can just as easily use more than two bins if you want.

And now back to Kanban. The reason I briefly covered the two-bin (multi-bin) system first, is that Kanban is essentially the same thing as a multi-bin system. The only difference is that sometimes, instead of sending an empty bin somewhere for replenishment, Kanban systems may just use a card or other “signaling” device such as a light, colored plastic marker, golf ball, paper airplane, snowball, whatever. OK, I was kidding about the paper airplane and snowballs, but there have been a wide assortment of “signals” used in Kanban systems. Regardless of the signaling device used, Kanban still comes down to an empty bin (or other type of container) triggering the event. For example, when using “Kanban cards”

***These are sample pages from the book *Inventory Management Explained*.  
More information is available at <http://inventoryexplained.com>***

| <b>Net Requirements and Planned Orders</b>                                    |    |    |    |    |    |    |    |    |   |    |    |
|-------------------------------------------------------------------------------|----|----|----|----|----|----|----|----|---|----|----|
| RM001 Rim (beginning inventory = 10 units, lot size = 25, lead time = 1 week) |    |    |    |    |    |    |    |    |   |    |    |
| Period (weeks)                                                                | 1  | 2  | 3  | 4  | 5  | 6  | 7  | 8  | 9 | 10 | 11 |
| Gross requirements                                                            |    | 15 |    | 20 | 5  |    | 10 | 30 |   | 5  | 15 |
| On hand                                                                       | 10 |    |    |    |    |    |    |    |   |    |    |
| Scheduled receipts                                                            |    |    |    |    |    |    |    |    |   |    |    |
| Net requirements                                                              |    | 5  |    | 20 | 5  |    | 10 | 30 |   | 5  | 15 |
| Planned order receipts                                                        |    | 25 |    |    | 25 |    |    | 25 |   |    | 25 |
| Planned order releases                                                        | 25 |    |    | 25 |    |    | 25 |    |   | 25 |    |
| Planned on hand                                                               |    | 20 | 20 | 0  | 20 | 20 | 10 | 5  | 5 | 0  | 10 |

**Figure 9E.** Here we show how we net the requirements and create planned orders for one of our components from the previous example. We start with 10 units on hand in week #1, therefore our gross requirement of 15 units in week #2 results in a net requirement of 5 units in week #2. Since the lead time for this component is one week, our planned order receipt in week #2 results in a planned order release in week #1. Our planned on hand at the end of week #2 will then be 20 units. This is the result of subtracting our net requirement from our planned order receipts for that period.

The planning process then continues through subsequent weeks using the same logic. We currently don't show any scheduled receipts for this item, but once we convert the first planned order to an actual order, the planned order release in week #1 and the planned order receipt in week #2 will disappear, and we will show a scheduled receipt in week #2. Scheduled receipts are then included in the calculation for net requirements.

our lead-time for RM001 is one week, MRP will create a planned order release in week 1.

Just to clarify this whole planned-order-release/planned-order-receipt issue, the quantity of 25 shown in periods 1 and 2 represents a single planned order that has a required date in week 2, and therefore must be released in week 1 due to the lead time of one week. If you were to inquire on this item in a typical MRP system, the planned order quantities would be associated with the planned receipt date, not the planned released date. The planned release date is purely used to signal to the planner (via an order report, ordering screen, or action messages) to place the order.

If you look at the last line of the example, we show a planned on-hand balance of 20 units at the end of week 2. This is calculated by subtracting our net requirement of 5 units from our planned order quantity of 25 units. As we go into week 3, we have no gross requirements, therefore no net requirements; and since we don't show any scheduled receipts, we will maintain our planned on-hand balance of 20 units through week 3. Week 4 shows a net requirement of 20 units, so since we have a planned on-hand balance of 20 units at the end of week 3, this will cover our net requirement for week 4 and result in a planned on-hand balance of 0 for the end of week 4.

***These are sample pages from the book Inventory Management Explained. More information is available at <http://inventoryexplained.com>***

preparation”, another for “wheel assembly”, another for “wheel dishing”, another for “truing and tensioning”, and another for “final inspection and packaging”. I could then set up the system to require that I report production at each step in the routing because this allows me to track my costs at a greater level of detail. But in the process of doing this, I have added a whole bunch of additional production-reporting requirements that not only consume additional labor, but also may force me to change my wheel-assembly process to more accurately track these steps as separate operations. While this may have made sense to systems guys and accountants looking at flow charts, it creates a ridiculously cumbersome process on the shop floor. Similar problems can be created when the realization that you can now have multi-level bills is so exciting that you start creating multi-level bills where there isn’t a real need for them.

## **DRP and Multi-plant MRP.**

---

The “multi-plant” part of Multi-plant MRP is essentially the same functionality that makes up Distribution Requirements Planning (DRP). In other words,  $MRP + DRP = \text{Multi-plant MRP}$ . To keep things simple here I’m just going to refer to this functionality as DRP, but will show how it integrates with MRP.

We have already covered how the main objective of MRP is to use the demand for higher-level items to drive the demand for lower-level (dependent) items. Similarly, the main objective of DRP is to use demand at downstream facilities to drive demand at upstream (supplying) facilities. The logic used in DRP is really just an extension of MRP logic.

We’ll start with a simple example. Let’s say I have a large centrally located distribution center (DC) that I use to supply several smaller regional distribution centers. Without a system like DRP, I would probably end up planning inventory separately at each DC. For example, each regional DC would have its own inventory planning system and place orders on the central DC when they needed replenishment. The central DC would have its own planning system and essentially treat the regional DCs as though they were separate entities (the same way it would treat a customer).

Subsequently, each regional DC and the central DC are likely forecasting demand in isolation of each other and there is no visibility between the DCs on inventory levels and actual demand. This is known as a “disconnected” supply chain, and while you may have no choice but to be disconnected with external customers and suppliers, there is no reason to be disconnected from your own internal customers and suppliers.

This is where DRP comes into play. Like MRP, DRP also calculates gross requirements and net requirements, and then uses the net requirements combined with lot-size rules to create planned orders. But where MRP would use planned orders to drive gross requirements of components based on a bill of materials,

***These are sample pages from the book *Inventory Management Explained*.  
More information is available at <http://inventoryexplained.com>***

buck by focusing greater effort on the items that have a greater impact on the business.

It's not just about focusing efforts though. It's also about catering inventory management methods to groups of items that share similar characteristics. For example, being a "fast mover" is a characteristic, but there may be other characteristics that are common to fast movers. You may recall in the chapter on safety stock I mentioned that relative variability is often lower with fast movers than with slow movers. So given that variability is a key factor in how we calculate our forecast and safety stock, it may not be a bad idea to explore using different parameters or even different techniques for these calculations based on the item's fast-mover designation.

When you really start analyzing your inventory, you will likely find that a simple designation of "fast mover" or "slow mover" is not adequate. That's where ABC comes in. You can use ABC to rank your items into as many categories as you like. So rather than fast movers and slow movers, you have A items that represent your fastest movers, B items the represent your second fastest movers, and so on. And despite the term ABC, you are not limited to just three (A, B, and C) categories. In fact, I generally find that I need to go to E or F (A-B-C-D-E-F) in order to adequately break down a specific inventory characteristic.

| <b>ABC by Sales Dollars summary</b> |            |              |      |           |             |                          |
|-------------------------------------|------------|--------------|------|-----------|-------------|--------------------------|
| ABC                                 | % of Sales | Cumulative % | SKUs | % of SKUs | Sales \$    | Average Sales \$ per SKU |
| A                                   | 50.00%     | 50.00%       | 176  | 3.2%      | \$1,000,000 | \$5,681.82               |
| B                                   | 30.00%     | 80.00%       | 952  | 17.1%     | \$600,000   | \$630.25                 |
| C                                   | 15.00%     | 95.00%       | 1586 | 28.5%     | \$300,000   | \$189.16                 |
| D                                   | 4.00%      | 99.00%       | 1230 | 22.1%     | \$80,000    | \$65.04                  |
| E                                   | 0.99%      | 99.99%       | 987  | 17.7%     | \$19,800    | \$20.06                  |
| F                                   | 0.01%      | 100.00%      | 643  | 11.5%     | \$200       | \$0.31                   |
| Totals                              |            |              | 5574 |           | \$2,000,000 |                          |

**Figure 10D.** Example of the summarized output of an ABC breakdown.

In Figure 10D you can see an example of an ABC breakdown. The example is based on ABC by Sales Dollars (based on annual sales), and "A" items represent 50% of sales dollars, "B" items represent the next 30% of sales dollars, and it works its way down to "F" items, which represent the bottom 0.01% of sales dollars. The "Cumulative %" column just makes it easier to see cumulative level breaks. For example, the combination of "A" and "B" items account for 80% of sales dollars. The "SKUs" column shows how many items fall into each category. This is where you really start to get a feel for the Pareto Principle. In the example we have 5,574 SKUs, but only 176 SKUs (3.2% of SKUs) account for 50% of our sales, and another 952 items (17.1% of SKUs) account for another 30% of our

***These are sample pages from the book *Inventory Management Explained*.  
More information is available at <http://inventoryexplained.com>***