

10 STEPS TO FINDING PROFIT //

IN LOCALIZED REPLENISHMENT

Most traditional replenishment systems are filled with gaps and constraints preventing retailers from reaching their most profitable potential. Greg Wilson, VP of field strategy with Quantum Retail, will walk through 10 steps to help you address and solve the constraints and limitations within traditional replenishment approaches. Greg will also discuss what to consider when evaluating updated solutions and capabilities to help you become more profitable.

QUANTUM RETAIL TECHNOLOGY



TABLE OF CONTENTS //

3 INTRODUCTION

4 USING THE RIGHT INFORMATION

STEP #1: Start to see the battlefield

6 BEHAVIOR AND THE SHAPE OF SELLING

STEP #2: You cannot change the seasons or the wind

8 THE MYTH OF THE PERFECT FORECAST

STEP #3: I have seen the future and it doesn't work

11 WHAT THE HECK ARE YOU TRYING TO DO ANYWAY?

STEP #4: Never lose sight of the goal

13 LOCK AND LOAD

STEP #5: Lock in the target before loading the inventory

15 WHAT NOT TO DO: MIX/MAX

STEP #6: It's the exception, not the rule

18 WHAT NOT TO DO: SAFETY STOCK

STEP #7: Please mind the gap

20 WHAT NOT TO DO: PRESENTATION AND COUNTER STOCK

STEP #8: Step outside of the comfort zone

22 SCARCITY AND ABUNDANCE

STEP #9: You can't always get what you want

24 NOW, HOW DO WE MAKE MONEY WITH ALL OF THIS?

STEP #10: Become profit aware

26 WHERE DOES ALL THIS GET US? PUTTING IT ALL TOGETHER

10 STEPS TO FINDING PROFIT IN LOCALIZED REPLENISHMENT //

INTRODUCTION



Since writing the [10-step series to more profitable allocation](#), I've been frequently asked about how many of the philosophies that were covered there might apply to replenishment. This series will build on the topics covered there and take a more replenishment-oriented look at gaps in traditional replenishment approaches. We will look for things that you can do to address these gaps within the constraints of traditional approaches, and also discuss what to evaluate if you are considering investing in updated capabilities.

In this series, we will focus on selling-location replenishment, aka store replenishment. However, in today's multi/omni-channel environment, I find "store" too constraining. While we'll cover warehouse/distribution center replenishment in a future series, many—if not all—of the concepts we'll be discussing here will apply to direct channels as well as stores and outlets. Let's start with getting the foundation correct...

- **Greg Wilson**

Vice President of Field Strategy, Quantum Retail

1 USING THE RIGHT INFORMATION //

Step #1: Start to see the battlefield



One of the most important components to getting better results is starting with the best foundation of data possible. In the allocation series, we discussed a few of these. Rather than repeating them here, I'm going to take the easy route for this first topic and suggest that readers refer to the [A 10 Step Guide to More Profitable Allocation](#). If you're not interested in allocation-centric topics, focus specifically on these topics—which also apply to replenishment and getting the foundation of data right:

[#1: Demand, Demand, Demand](#)

[#2: Debunking the Cluster Myth](#)

And finally the summary, [Putting it all Together](#), discusses why allocation (and selling-location replenishment) makes sense to prioritize and focus on improving.

Once that foundation is in place we'll begin with the next topic specific to replenishment. The following concepts should also be considered before we get into the replenishment specific topics.

HOW MANY SELLING-LOCATION FULFILLMENT SYSTEMS DO YOU NEED?

When I started in retail over 25 years ago, most systems were manual. I watched technology evolve within inventory management over many years, operating within the capabilities and constraints of what the technology of the day could support. Since basic products could be monitored over long periods of time, simple forecasting processes became feasible and evolved into replenishment technologies. Fashion, highly seasonal or other products with volatile selling characteristics were difficult to forecast so a different type of system began to rely on users to make guesses at appropriate historical activity to base store fulfillment decisions on; these became allocation systems.

Given the constraints of the technology when these processes evolved, it was understandable that two different approaches were used. However, considering the advances in technology, retail sciences and related capabilities that have evolved since then, it no longer makes sense to require two different systems to answer what is effectively the same question: "How much inventory is needed in the selling location?" If you find yourself in a situation where you're still supporting two systems with two very different processes and sets of capabilities, it may well be time to look into what advancements have occurred in fulfillment technologies over the last decade or two.

HOW CURRENT IS YOUR REPLENISHMENT CAPABILITY?

As mentioned above, technology was being applied to replenishment back in the 70's and 80's when technology was still quite young. Many of us grew up using systems like INFOREM, which was originally designed to forecast commodity items for the grocery industry. This was applied to other areas of retail and systems like E3 evolved to add many of the components and capabilities that originally had to be developed around the INFOREM engine. Since then a number of other similar technologies came around offering improvements in forecasting, criteria management, user interfaces, scalability, etc., but they were generally “a better mouse trap” based on the same process.

If we think about advancements in technology that have happened since the arrival of those original systems, which defined how replenishment is executed, it's eye opening. The Motorola brick phone was the pinnacle of mobile communication. The Internet was in its infancy and unknown to most people. Portable music was supported by that miracle of technology, the Sony Walkman. A lot has changed since then not only on those fronts, but in the capabilities of modern replenishment systems as well. Many technology vendors found themselves married to a technology platform that made advancing cumbersome, if not impossible. Some engaged in re-platforming exercises that took years and didn't necessarily have the desired improvements, some outright failed.

If you're using technology that was developed before or shortly after the turn of the millennium, there's likely to be a lot of opportunity available to you if you spend some time looking at what has since changed in replenishment.

HOW FRESH ARE YOUR CRITERIA?



Traditional replenishment systems required a variety of user inputs to achieve their true potential. Exercises around creating fresh seasonality profiles and updating service goals or model stock and min/max criteria are common to maintaining those systems. If these criteria are left alone, results begin to deteriorate. Similarly, if the processes to update them are flawed or compromised, the results will suffer as well. There have been whole solutions built and sold around “refreshing” or “optimizing” this information. Given the capabilities of today's technology, having to run separate, standalone utilities to get another system to perform just doesn't seem adequate any more. If you find yourself planning these exercises or running additional systems to update criteria like this, it may be a sign that it's time to refresh your knowledge on what advancements have been made in replenishment technologies.

Images courtesy of watcharakun, Salvatore Vuono/FreeDigitalPhotos.net.

2 BEHAVIOR AND THE SHAPE OF SELLING //

Step #2: You cannot change the seasons or the wind

One of the key criteria in traditional replenishment is a seasonality profile. Seasonality refers to the behavior of an item over time. The goal of seasonality profiling is to identify how time of year influences an item's demand. This influence can then be subtracted from observed selling to get a baseline for current activity. An example is that an item may sell an average of one unit per week in a given location, but that moves to three units per week during a particular seasonal time such as holiday, back to school, and so forth.

The importance and impact of seasonality is often underestimated. Having an incorrect seasonality profile actually has a double negative impact on the quality of the forecast being generated. This is due to how the profile is actually applied within the process of generating a forecast. The seasonal profile is used to “de-seasonalize” observed activity as it happens. This removes the impact of time of year such as holiday, leaving us with a “baseline” of demand – or what the average sales are regardless of the time of year. The reason the baseline is established is so we can then use it to apply forward seasonality to appropriate time periods we're trying to forecast. So the baseline we have today impacted by the seasonality for the period we're trying to forecast results in the actual forecast that is generated.

If we start with a flawed seasonality, our baseline will be wrong. We make matters worse when we use a flawed seasonality and the wrong baseline to create a forecast, which is now “double-flawed”. The odds of deriving a quality forecast with a poor seasonality profile are very low.

Then there is the fact that seasonality is just one level of behavior that can be profiled. If you work with shorter life or seasonal merchandise, an understanding of product lifecycle (how a product sells when it's new, vs. once established vs. near end of its full priced life) is important. If you have opportunity to forecast and/or replenish multiple times per week or per day, an understanding of how different days relate to the week or how times of day relate to the day will be important.

WHAT CAN YOU DO ABOUT IT?

The simple answer is to create and maintain better seasonality profiles, something easier said than done. There are a variety of tools available that can assist, but a couple of core concepts will ensure that you are at least consistent with the quality of your profiling.



First – if at all possible, execute your profiling activity on a base of historical demand, not just historical sales. If you haven't already, refer to the Allocation series [Strategy #1: Demand, Demand, Demand](#).

If you have short life products you should try to focus on product groups with similar lifespans in an effort to make the seasonality profile serve double-duty as a lifecycle profile as well.

Second – Make sure you're profiling at an appropriate level. We discussed how to go about this in the allocation series in [Strategy #3: You're Biased, But it's Not Your Fault](#).

Finally – revisit seasonality frequently. BARE MINIMUM is annually, but seasonally is a good practice.

WHAT SHOULD YOU CONSIDER IN THE FUTURE?


Since the introduction of traditional replenishment solutions, there have been a variety of explanations offered solely to support creation or "optimization" of seasonality profiles. In reality, technology has evolved to the point that a well-conceived system can now understand what appropriate product levels and groups are that should be used to learn about characteristics of behavior such as seasonality and can update them automatically when appropriate. Modern systems don't stop at seasonality behavior, however. They can also automate profiling of different "dimensions" of behavior such as lifecycle and time of day uniquely. They then automate the process of applying those to the forecasting process as appropriate. This means that replenishment users can now focus on activity that has a more strategic benefit rather than manipulating data as they've had to with traditional solutions of the past.

Solutions which continue to depend on a user profiling exercise or a periodic refresh or optimization of profiles and related criteria are advertising the fact that they haven't leveraged the capability that today's technology has to offer.

Image courtesy of digitalart/FreeDigitilPhotos.net.

3 THE MYTH OF THE PERFECT FORECAST //

Step #3: I have seen the future and it doesn't work



Any discussion about forecasting should start with a definition. In retail we regularly refer to everything from a high level plan based mainly on intuition, to complex statistical algorithms using modified historical activity and a variety of parameters at low levels of detail, as a forecast. As I see it, anything that estimates future activity in any way is a forecast. What really distinguishes these are the methods used, inputs considered, levels of detail, and quality of results.

When we discuss forecasting as it relates to localized replenishment, we ultimately need some estimation of activity that will ultimately drive inventory movement to the specific location where it is needed. While various levels of detail are often applied to solving this problem, it is a SKU/location level estimate of how much inventory is needed at any given time that we will be referring to here.

Notice that it is an inventory value, not a demand value that we fundamentally need. Demand is just one of many mechanisms that enable us to derive the inventory needed. Understanding demand is important to getting the right answer, but where retailers have frequently failed is in becoming overly focused on the pursuit of the “perfect demand forecast”. To measure our progress in this, we rely on metrics such as forecast accuracy, which compares the forecast to what actually happened.

While forecast accuracy is an important and valuable metric, the potential flaw is that measuring forecast accuracy at the level where it's essentially needed – SKU/location – it's so volatile and variable for the vast majority of products that it approaches being irrelevant. In reality, getting a consistently perfect forecast at SKU/location level would require the likes of psychiatric evaluations of all potential shoppers to determine if they're in a frame of mind to pull the trigger on purchasing a given size or color or shape of product in a given location on a given day. It's just not feasible.

To address this reality, we escalate hierarchies to move from the detail level akin to psychology, to an aggregate level that is more consistent and predictable as is done with sociology. Then we measure the quality of our forecast there. At some point we can normalize the inevitable noise that exists in the low levels into something a bit more consistent and predictable. If we measure forecast accuracy at those higher levels we can get more consistent measurements and claim to be accurate forecasters. However, doing that doesn't solve the problem at the detailed level where we ultimately need to execute.

We still need to get individual SKUs to individual locations. So to move from the aggregate forecast down to the execution level, we compromise on a variety of spreading mechanisms that ultimately get us back to the volatility that we didn't want to measure before. To make matters worse we've used an aggregate demand forecast and either slapped a guess at an "ideal" service level for groups of stores then spread that down, or spread our interpretation of demand to stores then applied a similarly defined service level target onto it to derive the answer to the critical question of "How much inventory do I need?" Worst case scenario is in situations where deriving a detailed demand forecast is so difficult we just give up and throw a min/max or model stock out there and replenish around those typically static and sub-par guestimates resulting in missed sales and markdowns. I frequently see retailers where this represents 70, 80 or even 100% of their replenishment strategies!

Historically, most effort on improving these processes has been focused on forecasting and applying more sophisticated calculations to compensate for volatility, variability and other factors that make forecast accuracy challenging. There are trends of having multiple forecasts come together to get to the "perfect demand forecast". But there is another approach to consider...

What if we concede and admit that the perfect forecast for purposes of driving localized replenishment is effectively the equivalent of "unobtainium". Then we can begin to focus on the question that really matters – "How do I get to the ideal level of inventory to achieve the product's goal?" for each location. Doing this enables us to utilize a forecast with known weaknesses, but consider these and marry that with a deeper understanding of how inventory plays into achieving a product's goal. Focusing on the inventory side of the equation can enable us to free up inefficient time spent on forecast adjustments and apply it to effective improvements on deriving more effective inventory support at detailed levels.

Consider this: at the detail level we need to execute at, a 5, 10 or even 20% improvement on forecast accuracy can often result in *little or no change* to the inventory answer – especially if there isn't much sophistication on how that inventory is derived. However a 1, 2 or 3% increase in availability or similar reduction in markdowns or carrying costs due to getting the inventory answer right can quickly add up to big profits.

WHAT CAN YOU DO NOW?

The answer here depends on what your current process is.

If you're spending all of your time manipulating forecasts so they result in inventory levels you're comfortable with, it's probably time to consider making changes to how inventory levels are derived. Rather than manipulating the forecast, what options do you have to define more effective inventory levels relative to a lesser-managed forecast? Can you put more time into getting to a lower level of detail on setting service level targets? If so, what impact does that have on your results vs. manipulating the forecast?

If you're driving a large portion of your business by min/max or model stocks, consider revisiting these more frequently or potentially analyzing them at lower levels of detail. We'll discuss strategies on how to do this in another part of this series.

Learn more about understanding your forecast [here >>](#)

WHAT SHOULD YOU CONSIDER IN THE FUTURE?

Rather than focusing on the quest for the perfect forecast, consider products that have put effort into solving the inventory side of the equation with more sophistication. If you can get a consistently accurate forecast with understood rates of error using little or no effort and redirect that to focusing on determining ways to achieve ideal inventory positions by SKU/store, there may be much more upside opportunity to be had.

4 WHAT THE HECK ARE YOU TRYING TO DO ANYWAY? //

Step #4: Never lose sight of the goal

So you've spent time getting your demand forecast tight, you've analyzed and clustered stores based on similar items then assigned service targets by cluster. One question: What are you trying to achieve? I've asked this question of many retailers and I get everything from a detailed explanation of service level to glazed-over looks saying, "What do you mean?"

What I'm looking for is whether we're considering the unique roles that products play within any given assortment or range within selling locations. For example, an item that's there to drive traffic—or in an extreme case to be a loss leader—is inherently different than an item whose role is one of a fringe assortment item. These are different from image or statement items, which differ again from the big margin drivers. And in between all of these are core assortment items of different velocities and lifespans. Each of these (and other) roles implies different goals.



Read more about strategic merchandising [here >>](#)

Ultimately the point is that each should be considered for its place among these unique roles, and deriving the correct level of inventory support should differ. In fact, if you have two items that appear the same based on unit velocity, but which play different roles and have different margins & retails, they may well require very different levels of inventory support.

Consider two items, each sells an average of three per week in a given location. One is a traffic driver with a low retail and a tight margin the other is a profit driver with a high retail and a large margin.

The traffic item is there to satisfy the customer. If you fall short of that item the customer may well leave and not purchase anything else. The cost of stocking this item a bit more aggressive and is offset by capturing the potentially missed opportunity on other items and is mitigated by the fact that the item is a lower retail.

The profit driver is how we make a disproportionate amount of our income, so we need to preserve the profit. The cost of stocking this item too aggressively can be very damaging, however, since there is a holding cost eroding the margin. This is amplified by the fact that the item has a high retail; it is consuming precious capital, especially if I'm over stocking across many locations. Here finding the sweet spot of profitability where I capture most demand, but may sacrifice some that is more volatile to preserve the capital and avoid the carrying costs, makes more sense.

The roles of products differ from retailer to retailer, but they exist in some form for all retailers. Each role does carry with it implications that can and should impact our stocking decisions. Beginning to consider this can quickly have an impact on making inventory more productive and efficient.

WHAT CAN YOU DO NOW?

Start having the discussion about what roles exist in the assortments you carry. Merchants intuitively understand many of these roles, but haven't necessarily defined them to the degree that they're applied to products for the rest of the organization to consider. If you can get to a point where you've defined a small set of roles and what each means to how they should be supported with inventory, you can then alter the way you're defining the parameters that drive inventory decisions. Doing this consistently can return significant inventory performance improvements.

WHAT SHOULD YOU CONSIDER IN THE FUTURE?

If you're investing in a new technology, beware of systems that aren't inherently able to consider and understand the roles of products. Set this key component into an environment that can interpret and execute against it and you will be on your way to better performance. The best systems internally translate a simple role into complex literal goals that can be evaluated and optimized against on a location-by-location basis for each replenishment decision.

5 LOCK AND LOAD //

Step #5: Lock in the target before loading the inventory

The [previous post](#) in this series discussed establishing your goal. What is it you're that trying to achieve? Odds of getting to the ideal inventory answer without knowing the goal are slim. That goal becomes the driver that helps us identify what inventory position we should stock to.



A major part of determining the goal is assigning the role of a product. Once the role of a product in the assortment is established (e.g. traffic driver, money item, fringe item, image item, etc.) it provides guidance that can be used to direct inventory targeting. If we use a traffic driver for example, the goal is to have a high in-stock because we don't want to disappoint customers who shop specifically for that item and are likely to buy other things while in the store. If we translate this into a metric value that you have available to drive your inventory targets, you'll be more aligned with the goals of each product. This perspective will allow you to more appropriately define targets than common approaches that focus mainly (or entirely) on sales volume.

Since most traditional replenishment systems use a service level target as the inventory driver, we can use that to accomplish the objective. In the case of the traffic driver, a very high service level is implied. Due to the nature of a traffic driver we can think of missing a sale as being more costly for these than other items because of the implication that other sales will be lost if we lose this item's sale. These items will, therefore, have some of the highest service targets of any products.

The process becomes a bit more complex when we start looking at other roles such as money items. Money items are all about profitability so the goal is to find the target that maximizes selling without being subject to overstocks, which erode profitability. In traditional environments, we are again constrained to using service level as a driver so an analysis and a few assumptions will be required to find the appropriate targets. Margin will have an impact on the answer as well. For example, an item with a high margin can lean toward a slightly higher service level since the value of having an extra unit that may capture a possible sale is less than the cost of a markdown that might result. If margins are tighter, we need to be more conservative since the cost of markdowns and/or holding can more quickly slash margin and profit.

Similar logic can be applied to any role you define for products in the assortment. It's also necessary to consider the remaining life of a given product. A product that is approaching the end of its life will have more margin exposure (due to inevitable markdowns ahead) than an item that has a long lifespan.

WHAT CAN YOU DO NOW?

First, define product roles and translate those into high-level definitions of what they mean as objectives relating to inventory. Since most traditional systems are driven by service level targets we must then begin to analyze the facts that reside in historical performance to determine what levels of service result in achieving the measures each role implies. How much service leaves you with a profitable ratio of inventory (money item)? How much service enables you to capture as many sales as possible without resulting in carrying an unacceptable level of inventory (traffic)? How much service enables you to minimize inventory investment and perhaps purposely be out of stock occasionally but not so much that the image that item represents isn't lost



(image)? How much service enables you to capture sales opportunistically but sell through quickly with little exposure to carrying costs or markdowns (fringe)? And so on and so forth.

WHAT SHOULD YOU CONSIDER IN THE FUTURE?

If you're investing in a new technology, this is an area where more modern solutions really begin to distance themselves from traditional approaches. Modern, sophisticated replenishment systems actually evaluate the literal profitability results that any given inventory decision results in and does this uniquely for each product-location combination. Some of these systems also consider impacts to profit including cost of holding merchandise and impact of likely future markdowns. Better systems actually work across multiple objectives to ensure more complex goals can be met. The best actually evaluate their past inventory recommendations and measure success, which are learned from and applied to future decisions so results continuously improve. If you're investing in new technology this area of inventory targeting is one dimension that can have a surprisingly large impact on the quality of results.

Learn more about next generation technology [here >>](#)

6 WHAT NOT TO DO: MIN/MAX //

Step #6: It's the exception, not the rule



Traditional replenishment processes have many constraints that cause us to come up with creative solutions. A few of those solutions have become standard and are accepted as foundations of the replenishment process. It is time to challenge these and look at what new approaches can do to improve results.

Many of these strategies result from the fact that the forecasting capabilities of traditional solutions have significant weaknesses when applied across the variety of product behaviors we see in retail. To make up for these weaknesses, we frequently supplement the forecast on the inventory side or ignore it all together. The most commonly applied example is the use of a min/max strategy.

Common min/max philosophy is that the forecast is too volatile or too low, so we ignore it and put an inventory target range in place. The use of “model stocks” is a variation on a similar theme: If we look at the level we are forecasting at for selling location level replenishment, min/max will often account for a large percentage of the assortments we carry. It is common to see retailers with 80%, 90% or 100% of their assortments on min/max. Why have a system capable of forecasting product demand if you are not going to use it? While min/max is a simple approach that is easy to understand and train users on, it is sub-optimal in virtually every case.

For starters, any fluctuations due to seasonality, lifecycle or other dimensions of behavior are generally overlooked using min/max. The [objective of the product](#) as we discussed earlier in this series is generally not considered either. This can have a profound effect on the results of slow movers – especially those with high retails and/or short lives.

For volatile products, short life products and product introductions, there are multiple ways to understand potential variability and to improve either the forecast or (more importantly) the inventory targeting process. For slow moving products, we need to understand if there are situations where we can sell multiple units within a replenishment cycle and how much is enough service to reach our goal (profit, image, core, etc.).

To further impede potential, min/max parameters are most commonly managed at store cluster level. This means that all stores in a cluster are treated the same. They are not the same. Even if they appear to be similar at one point in time, they will not be at other times.



Learn how to cluster stores [here >>](#)

While there can be times where use of min/max is appropriate, it should be viewed as an exception, not the rule. If you find yourself with a large proportion of your assortment, or worse a majority of it, being managed by min/max you are sure to have a lot of opportunity for improvement.

WHAT CAN YOU DO NOW?

If you are constrained by a traditional system's limitations, it does not mean you are stuck. The road to getting better results starts with a little more effort in deciding on the approach you are using and setting parameters while considering a bit more information and/or detail. First, have you chosen to use a min/max strategy because it is better or because it is simpler? If your answer is because it is simpler, then consider putting in the effort to dial-in the forecast and service parameters to get the capabilities of your forecasting solution to provide some returns. Doing this will allow you to consider the impacts of seasonality and other forward looking activities rather than being based on an average of historical boundaries that are common to defining min/max parameters.

If your forecasting capabilities cannot support the product and your alternative choices are limited as most traditional replenishment systems are, then we turn to setting the min and max parameters more effectively. This generally comes down to more detail. Here are some things to consider:

- How often do you review the min and max? If it is annually or seasonally, you will want to validate that the seasonal fluctuations do not warrant establishing more frequent revisions to the ranges.
- Have you considered the product's goal? If it is centered on profitability, being stocked out for a few days may be more profitable than having a min parameter that is ensuring you always have a unit on hand. If it is an item that is conveying an image, have you really lost the image entirely if you are out of stock occasionally?
- Are you driving the inventory decision by space or by demand? Lean more toward demand as often as possible. Sure some items require a strong presence, but not all items do.
- For slow movers with a none-or-one min/max, do you have situations where you can sell multiple units within the replenishment cycle? If so, is the potential additional sale enough to warrant carrying an additional unit?
- Have you revisited your replenishment cycle? If you are replenishing every fourteen days, would it be possible to change that to every seven days? Doing so may enable you to reduce the level of inventory

you commit to individual locations that may have more volatility, while still being able to react quickly enough to avoid missing opportunities.

- If you are setting your min/max at cluster level, can you get down to the store level? If that is not practical, can you make your clusters more representative of how stores sell that specific product? More detail on location or product levels gets you closer to the unique behavior that happens at item/store level.

WHAT SHOULD YOU CONSIDER IN THE FUTURE?

If you are investing in a new technology, you should be looking for systems that have leveraged current thinking and advancements in technology. If you evaluate a new system and min/max is one of the featured replenishment methods – especially if it appears like it will be used on many of your items – it is time for some healthy skepticism around whether that solution has matured to finding a better way.

Learn more about forecasting and replenishment solutions [here >>](#)

Forecasts still are not (and not likely to ever be) perfect, but forecasting capabilities continue to improve. Some modern replenishment solutions are capable of understanding the weaknesses in their forecasts and automate their understanding of how to compensate for those weaknesses in the process of deriving inventory targets. The best ones constantly monitor their recommendations, learn when they do well and determine how to improve when they do poorly.

7 WHAT NOT TO DO: SAFETY STOCK //

Step #7: Please mind the gap

While ideally replenishment is driven by the perfect forecast, as we discussed [earlier in the series](#) there is no such thing as the perfect forecast. This is particularly true in localized replenishment where we must execute at very low levels of detail. Among the most difficult products to forecast are those that have high amounts of variability or volatility. These conditions are inconsistent and challenge even the most sophisticated forecasting methods.



Understanding this situation, replenishment systems have endeavored to compensate for it using something commonly referred to as safety stock. This is basically an acknowledgement that there are gaps between the forecast and actual sales. Safety stock is commonly calculated based on some measure of demand volatility, the implicit assumption being that volatile demand implies difficult-to-measure demand, which in turn implies high forecast error. The problem is that demand volatility is itself constantly in flux, making it difficult to reliably calculate the ideal safety stock.

WHAT CAN YOU DO NOW?

If you have no measure of volatility in your current process, you can start to make progress by looking at rate-of-sale. Larger locations and/or products that have a higher rate-of-sale tend to be easier to forecast and therefore require lower *relative* levels of safety stock. As an example, suppose you are currently holding 4 weeks of supply in all locations for a given product. You can likely do better holding only 3 weeks of supply in the higher-volume locations and 5 weeks in the

lower-volume locations given that the higher-volume locations are likely to have more consistent week-to-week sales.

If you already have a safety stock measure, you then need to look for ways to make it more accurate. Reanalyzing for different seasonal periods or timeframes in product life can begin to accommodate this. Any opportunity to calculate safety stock at a more detailed level can help as well. Different location levels or lower level of product (rather than using averages) can enable you to get better results. The effort can be tedious, but – especially for high margin products – it can be quite fruitful.

WHAT SHOULD YOU CONSIDER IN THE FUTURE?

Some solution providers have invested the efforts of their science teams to generate processes such as “safety stock optimization” to try to reduce the gap left by volatile and variable behavior. While better than manual

processes, these are generally a signal that traditional approaches are still the foundation of the replenishment logic and that newer thinking has yet to be applied to the solution space being covered.

More forward thinking approaches use processing capabilities of today's technology in addition to new scientific approaches which understand the facts of what different recommendations resulted in to measure their productivity rather than trying to estimate forecast error as a foundation. Better systems measure historical facts against the goals that we've been discussing throughout this series. If an item is a money item, did the system recommended position result in achieving the goal of highest possible profitability? If not, what level would have? How, when, and why is that? The best systems ask and answer these questions, learning from the results and constantly improving the answers and subsequent performance.

Learn more about systems that constantly improve and learn [here >>](#)

8 WHAT NOT TO DO: PRESENTATION AND COUNTER STOCK//

Step #8: Step outside of the comfort zone

One of the most common ways replenishment systems are abused is by using capabilities for things that they were never intended. The area I have observed as the most common of these is the abuse of presentation levels, minimums, and/or counter stocks. In a properly executed environment these values are intended to be precisely what they describe: the minimum inventory level a product should ever be at within its active life.



Traditional replenishment systems are quite complex, especially for new users. As such, it is difficult to keep an appropriately trained group of people in place to manage them so they perform to their full potential consistently. Efforts to simplify this frequently result in simple shortcuts. One common shortcut is to compensate for erratic results by using presentation as a buffer. Doing this is in effect a shortcut to addressing the challenge of last week's 'what not to do' - [safety stock](#).

Another example is setting presentation levels equal to a pack quantity. The thinking is often that this will ensure we keep shelves in stock. The reality is that, in a properly configured and running replenishment environment, this will lead directly to costly overstocks.

If you find yourself using presentation or its equivalent for anything other than being the minimum you need to make the inventory presentable, you more than likely have some real opportunity available. As an example, an item that lives for five months is considered to expend a holding cost per week equivalent to five percent of its underlying cost; these costs are higher for shorter-living items.

WHAT CAN YOU DO NOW?

Before we can turn off the "overrides" that an inflated presentation value represents, we need to be careful to understand why it is being done. More often than not it is just a safety or comfort level reaction. However, if there is an underlying weakness when presentation is not inflated, that needs to be addressed. This can generally be controlled by understanding and setting other replenishment criteria more accurately. If you have invested time in the areas discussed in prior posts to this series, you should have better inventory/service level targets, which are considering [product roles](#), and you are compensating for forecast weaknesses with an appropriate safety stock.

If those areas are defined appropriately, presentation measures should be open to being set properly, which will quite likely reduce over-inventory situations.

WHAT SHOULD YOU CONSIDER IN THE FUTURE?

If you are investing in the future, this topic has already been covered in “future” sections of the posts in in this series. Having the capabilities described in those will free you entirely from needing to make presentation represent anything but the pure value it is intended for – aesthetics!

9 SCARCITY AND ABUNDANCE//

Step #9: You can't always get what you want



Retail is a volatile environment. Change is the norm. One area where this is often apparent is in the receipt of orders. Whether due to demand changes, supplier error, poor initial allocation decisions, shipping damage or other constantly evolving activity, it is generally more common to have too much or too little product than it is to

have precisely what you need. Since this will have an impact on how you are fulfilling demand, it is critical to have a strategy for how to deal with it.

Simple methods will just scale by spreading the overage or underage across the contribution base of the target inventory. This has proved to be unreliable because it frequently fulfills high volume locations and completely eliminates low volume locations. As this became known as an undependable practice, a number of more sophisticated strategies evolved such as using need percent (which evaluates the importance of need relative to total need). The results are often complicated by packs or prepacks, which traditionally rely on the use of somewhat arbitrary rounding rules to get to a result. This again will typically lean toward prioritizing high volume locations, which will come closer to rounding cutoffs.

To make matters more difficult, when we are dealing with 100's or possibly 1000's of locations, it is common that many of them will have exactly the same targets or remaining need. How to break the tie? It is eye opening to see how many traditional environments end up giving the next unit to the location with the lowest store number in these situations. It would be tough to get less optimal than that!

WHAT CAN YOU DO NOW?

First off, if you have a situation where a tie goes to the lowest numbered location, find a way to make a change. Any change, even random selection, would be an improvement.

If you are using a simple spread in areas of shortage, try to get a more representative base in place. This can be done a number of ways depending on how your targets are being derived. For example, if you are using model stocks, consider reworking models to reflect values that are more aligned with your constrained inventory position.

If you have control over how the shortage/overage itself is being divvied out then consider the use of percent needed rather than straight unit needed. For example:

Location A needs 4 units to fulfill its target of 16
Location B needs 3 units to fulfill its target of 8

Simple logic will give the next available unit to Location A. Location B, however, has a higher percent needed since the three units it needs represent 38% of its target whereas Location A's need is only 25%. Therefore, the importance of the next unit is greater for Location B than it is for Location A.



While not a perfect solution it can be a simple way to move in the right direction.

WHAT SHOULD YOU CONSIDER IN THE FUTURE?

As we look at solutions that produce a more ideal result, we need to carry forward some of the concepts discussed earlier in the series to solving this problem. If we know what your **objective** is (profit, revenue, etc.), we can utilize that to assist in making the shortage or overage determination by location.

First, we can evaluate the importance of each unit of need. Rather than just relying on the target units, or even the slightly better percent need, it becomes important to understand the variables that lead to that target. Two locations with a need of six units may have gotten there for very different reasons. One location may be very volatile selling anywhere from one to 10 units in the period, while another may sell five, six or seven very consistently. These differences represent very different prioritization of each needed unit. Better, modern, functionally integrated replenishment solutions can look beyond the unit need target into components of the forecast including volatility, seasonality, lifecycle, etc. and use those to properly weigh the importance of each unit of need.

Learn more about establishing targets [here >>](#)

The best solutions then take that understanding and convert it onto the value of each unit of need as it relates to your defined objective. For example, what is the likelihood that sending an additional unit to one location will result in greater profitability vs. sending it to another location?

Carrying the logic down to individual units can net out an even better answer when dealing with packs. This allows the system to determine not only the value of the needed units, but the cost of unneeded units – resulting in, for instance, sending packs to the location that can sell more of the units contained in the pack profitably than another location.

Images courtesy of Evgeni Dinec and creationzs/FreeDigitalPhotos.net

10 NOW, HOW DO WE MAKE MONEY WITH ALL OF THIS? //

Step #10: Become profit aware

Retailers have many different definitions of success and talk about many things as being ‘drivers’. The metrics we use – revenue, margin, turns, sell-through, etc. – to monitor and measure success can have a different focus at any given time. Overall, they all translate directly or indirectly into a retailer’s ability to generate profit. Regardless of the ultimate motivation, we all need to generate profit to remain viable.

The way we generate profit is by procuring and placing inventory in a way that allows us to sell it profitably. Given this reality, shouldn’t we all be measuring effectiveness and driving execution with a literal evaluation of the profit impact of every decision we make?



If we apply this thinking to the space of localized replenishment, we will quickly see that profit is rarely considered literally. There are a small number of criteria that can be manipulated and set in hopes of finding a “sweet spot” of more profitable results that are rarely utilized. That sweet spot is where we have enough inventory to capture demand given all its volatility and variability without holding so much inventory that we erode profit through markdowns or by tying up capital that could be better spent on more efficient inventory. In this day and age of technology, *not* evaluating any potential inventory decision on its literal impact to profit begs one simple question:

Why not?

WHAT CAN YOU DO NOW?

If you are executing in a traditional replenishment environment, your best chances of driving profit are by taking the recommendations of the [previous posts](#) in this series and applying them to the best level that your resources are capable of. Looking at what level of service resulted in the most profitable outcome historically and learning from that is a big step. It is unfortunately unrealistic to expect that you can re-define an aging solution to be literally profit aware let alone profit optimized. Systems based on evaluating only unit movement and driven by static service targets are severely limiting.

WHAT SHOULD YOU CONSIDER IN THE FUTURE?

Today’s technology has been applied to solutions that evaluate a variety of inventory positions and determine all of the criteria that impact profitability before recommending an inventory target. There are options that will support you by providing “optimized” results based on evaluating some components of profitability. These are generally provided as a service whose outcome is a list of better-informed static parameters for your traditional system. It will improve results, but the improvement will begin to deteriorate almost immediately as behavior changes and will continue to degrade until the process is repeated.

A better alternative is investing in a system that automates this process and dynamically considers the profit impact of a variety of possible answers for each product/location on an ongoing basis. This capability is a

significant departure from traditional solutions and is one that can obviously provide impressive and continuing returns. If your goal is to be a leader in your retail vertical you can expect to be competing with others that have begun to adopt this approach. The result is higher service to customers with less inventory investment leading to more revenue and ultimately more profit.

Read more about the art of localizing inventory [here >>](#)

WHERE DOES ALL THIS GET US?

PUTTING IT ALL TOGETHER //

So where does all this get us? Well, regardless of what your current selling location replenishment capabilities are, there is likely opportunity for improvement. Most solutions in use today have their roots in technology and processes that were derived decades ago. The sheer number of decisions these systems are making means that even small improvements when applied consistently can return huge rewards.

WHICH COMES FIRST?

Start with a better base of data - keep the data fresh and accurate. Make sure to keep your eye on the goal of better inventory decisions and, more specifically, individual product goals. Stay away from distractions with less return, such as minimal forecast accuracy improvements. Reduce your dependence on sub-optimal parameters like min/max and presentation as drivers of inventory decisions; this will free up the capabilities of whatever solution you have. And be able to react to the inevitable variations in supply so that you stay on task.

Doing all of these things with a traditional solution is unlikely. To get near-term benefit, focus on one or two of these areas and apply some effort into improving just those. This can often lead to revenue and profit gains of 1-2 percent making an unchanged supply of inventory more productive.

YOUR MOST VALUABLE INVESTMENT

If you have opportunity to prioritize investing capital in new capabilities, be sure to include the potential of better inventory management. Many are fooled into thinking that what they have is “good enough”. There are many higher profile areas of activity that compete for prioritization. E-commerce/omni-channel capabilities are often cited as the “fastest growing” areas of retail, but as bricks-and-mortar retailers grow those businesses, it’s even more critical that inventory is managed within and across all channels. The three biggest investments most retailers make are: people, real estate and inventory. The area within those having most opportunity is invariably inventory. It’s not uncommon to see migration from traditional replenishment solutions to a modern solution that provides everything discussed in this series to return anywhere from three percent to as much as double digit improvement in revenue and profitability depending on current baseline.

Thank you for following this series. If you have any questions or comments, please feel free to contact me at greg.wilson@quantumretail.com.