toolsgroup

+ White Paper

Forecasting Intermittent Demand 101

+ Everything you need to know about long-tail demand and how to forecast slow-movers

Introduction – The Long Tail and Intermittent Demand

For most businesses, the "long tail" is growing. You don't have to be an internet retailer like Amazon to see it. Companies are dealing with more slow moving items with intermittent demand and unpredictable demand patterns.

The result is that inventory mixes are wrong. Some products are being over-served, locking up precious working capital, while others are being under-served, and causing loss of margin and market penetration. Companies have too much of the items they don't need and not enough of the items most in demand. This creates a significant opportunity to improve both their top and bottom line. Later we will discuss how, but first let's take a closer look at intermittent demand and what's causing long tails.

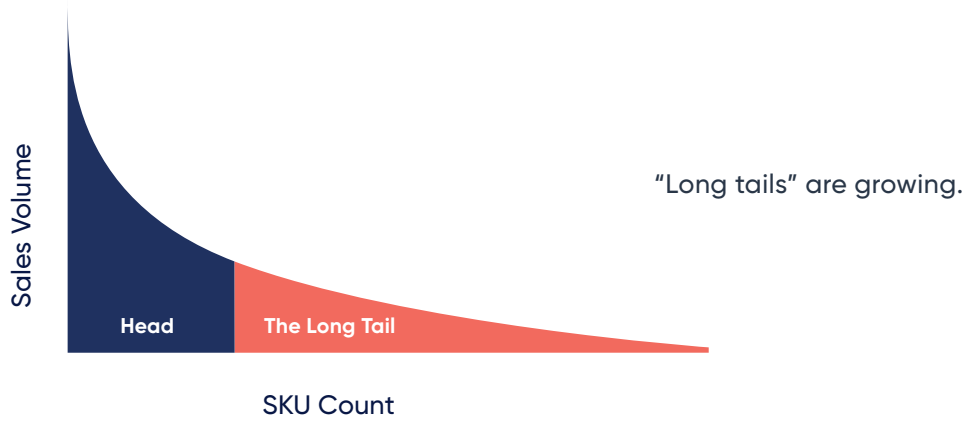

What's Causing Intermittent Demand to Grow?

From a supply chain perspective, long tail demand begins at the threshold between the "head" and the "tail", which is the line between fast and slow-moving products. This is the point where demand becomes "intermittent" and there are at least as many zero-demand as non-zero demand periods, or a 50% probability of having a zero-demand time bucket.

The long tail concept was originally developed to describe the fractured demand commonly seen at Internet retailers like Amazon, but even companies with more traditional supply chains also experience the long tail effect, along with intermittent, unpredictable demand patterns. Three business trends are causing intermittent demand to grow:

/ **Product proliferation:** Most companies have more products and product variants than 30 years ago. More products means splitting the demand between more buckets. When the percentage increase in the number of SKUs is larger than the percentage increase in sales, sales per SKU decrease. So product proliferation translates into more intermittent demand and higher variability at the individual SKU level.

/ **More frequent replenishment causing more granular forecasting:** While more frequent deliveries allow companies to react more quickly to changes in demand, they also mean that both replenishment and demand should be modeled and managed in shorter time buckets, typically daily. This shorter time bucket means higher demand variability. If demand behavior is observed in monthly buckets an SKU may look like a stable "fast mover". But if observed in weekly buckets it will look more like a lumpy "slow mover" and at the daily level will likely to be intermittent.

/ **Extended supply chain planning:** Increased collaboration between vendors, distributors and retailers and vendor-managed inventory (VMI) provide vendors with downstream demand visibility closer to the end-customer. It allows manufacturers who used to focus on planning their big regional distribution warehouses to obtain a more reliable demand signal, often at final customer facing demand. As the focus shifts from the "Ship-From" vendor DCs to "Ship-To" retailer DCs and even to retail shelves, demand is increasingly disaggregated into smaller demand streams. As a result, demand variability increases and slow-moving behavior becomes more common.

These three trends; product proliferation, more frequent replenishment, and demand streams closer to the end customer all segment demand into smaller and smaller buckets, causing the resulting demand to become more intermittent.

Let's Compare: Aftermarket Parts and CPG Tails

The trends described above are impacting demand and creating long tail intermittent demand in a wide range of industries. For example, the graph below on the left plots SKUs for a large automotive aftermarket parts business where 98% of the SKUs are intermittent or in the tail. This may not be surprising. After all, aftermarket parts businesses are known for lots of slow moving SKUs. Nonetheless, the sheer size of this long tail is impressive.

In the example on the right, what may be more surprising is that even for a globally branded "fast moving" consumer packaged goods company, the long tail consumes 86% of the SKUs. This is a lot of intermittent demand for a business that most people would not normally consider prone to this type of behavior. One important reason is that, like many consumer goods companies, this company now has a daily replenishment control frequency. The majority of our customers are now working with shorter time buckets, either because of the need for rapid demand sensing or in order to support high frequency replenishment processes.

The Anatomy of the Tail at an Automotive Aftermarket Parts Company

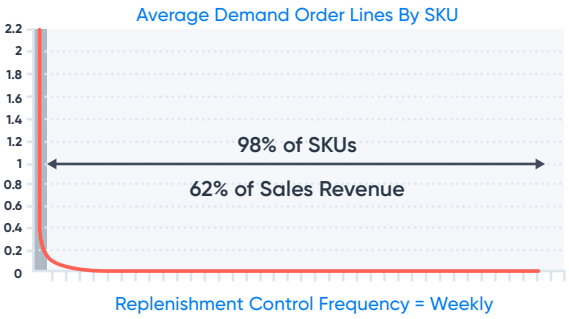

The Anatomy of the Tail at a Branded CPG Company

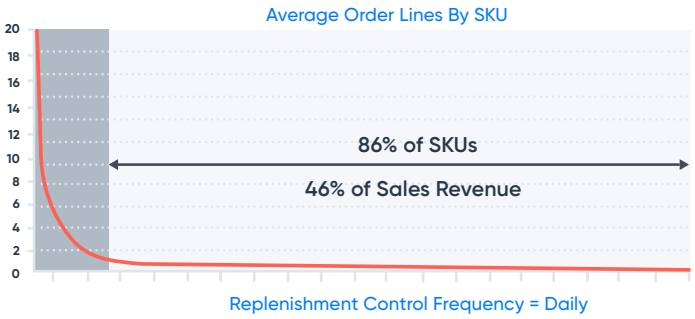

Problems Caused by Intermittent Demand

These days most every company seems to have a lot of intermittent demand. Let's look at the implications of this instability on the supply chain.

/ **Forecast Accuracy:** As a reminder, the critical issue related to dealing with the long tail is that the demand is intermittent: there are many zero-demand periods. This causes demand variability to be very high and the demand probability distributions to be highly skewed. The intrinsic variability caused by low order frequency has very little to do with forecast accuracy. For instance, if an SKU has average sales of one unit every 10 days, and the daily forecast is close to 0.1, the forecast is quite accurate overall. The problem is that it is very difficult to predict the next demand. So any attempt to improve "forecast accuracy" will likely be useless, since it will not lead to any relevant reduction in perceived demand variability.

/ **Inventory Modeling:** In an intermittent demand environment where demand variability is high and the demand distribution is skewed, classic demand and inventory models do not perform well. Intermittent demand behavior creates a significant mismatch between inventory levels required to meet the service-level objectives and the inventory-levels calculated by the most commonly used traditional supply chain solutions. Deterministic APS and ERP technologies are not designed to address the skewed distributions, particularly at high service levels (e.g., 98% and above).

/ **The Bullwhip Effect:** A final issue arises from the bullwhip effect caused by intermittent demand. Even if advanced, "integrated" replenishment logic is used, the presence of many slow-moving items produces significant gaps between what is planned and actual results. Planners typically focus on the head and manage the tail as exceptions. When inventories get out of line and stock-outs occur, extensive manual intervention tends to be the rule. Due to the difficulty of planning such inventories planners will tend to act directly on replenishment orders. The risk and the tendency is to overreact. Overstocks ensue. So inadequate technology for managing the slow-moving and intermittent items in the long tail causves shortages that lead to inappropriate manual intervention. This manually induced bullwhip effect creates overstock in the tail that accumulates over time.

The Bottom Line: Misaligned Inventory

To recap, intermittent demand has an intrinsic variability that traditional enterprise systems cannot handle, creating a highly "nervous" and unstable operating environment. The inventory mix--and consequently the service levels—across the network get imbalanced and out of line with business objectives. Planners try to compensate for this lack of performance by manually adding inventory, which typically accumulates in the tail. Meanwhile, faster movers in the head may be underserved.

Without good statistical demand and inventory models, there is neither the right inventory nor the desired service levels. The bottom line business impact is simple: the tail accumulates working capital, without delivering desired performance.

The graphs below show a plot of optimized stock levels (blue dots) against traditionally planned stock levels (violet dots). The x axis is the SKU count sorted in descending order of demand-frequency.
For the automotive aftermarket company represented in the graph on the left, you can see that inventories are largely misaligned, with a general tendency to overstock across the entire product range.

The well-known branded CPG vendor represented in the graph on the right illustrates another common problem: not only are many items clearly overstocked, but many others are also understocked. They are not delivering an acceptable service level, causing a high percentage of lost sales and margin. In other words, the business is under-serving many high margin items that should be receiving more inventory.

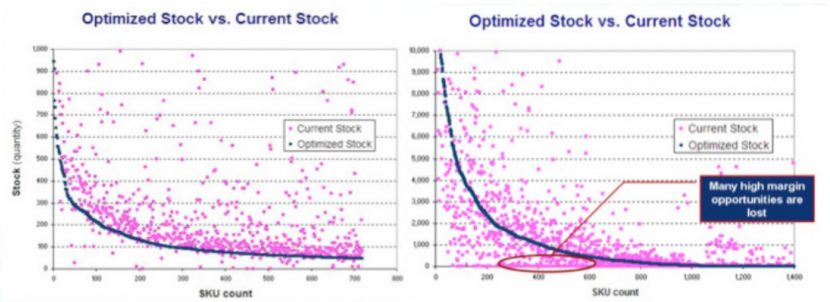

Automotive Aftermarket Company Branded CPG Vendor

The Key Shortfall is Demand and Inventory Modeling

Successfully forecasting the long tail and managing intermittent demand requires accurate and robust demand and inventory modeling that generates reliable service levels in the long tail. Secondarily, but also important, it requires disciplined processes that eliminate manual intervention and bullwhip behavior.

The problem with traditional solutions is that they oversimplify the problem. While such simplifications might be acceptable for fast movers, they don't work well in the tail. With inadequate models of demand and inventory behavior, no amount of optimization can solve the problem.

For example, a typical demand modeling oversimplification assumes that the demand probability distribution corresponds to the forecast error probability distribution. The resulting distribution, which is typically assumed to be "normal", is then used for inventory modeling and safety stock calculation. The example of intermittent demand in the figure below shows that the actual demand distribution is highly skewed and far from a normal distribution.

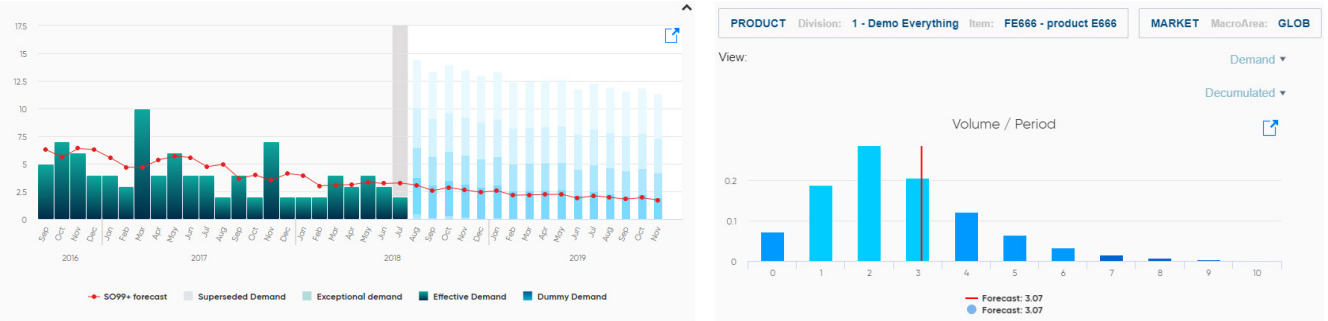

The statistical forecast (red curve) closely follows the "average" demand, but the infrequency of demand causes random deviations that are difficult to predict and are asymmetrically distributed or skewed with respect to the average and the statistical forecast. This can be clearly seen in the shape of the probability distribution in the image on the right. Unless this asymmetric behavior is recognized, the resulting forecast will be extremely unstable and very inaccurate. Most important of all, inventory behavior will be modeled in a way that does not provide reliable safety stock settings.

The far more accurate way to model the demand is by determining the demand distribution illustrated by the blue bars on the right side of the image whose different shadings represent the different probability ranges expected for future demand occurrences. These model the variations in the right-hand tail of the demand distribution.

The bottom line is that items in the long tail cannot be treated using common simplifications or "shortcuts". For instance, assuming that demand conforms to a normal distribution creates a very large difference between the safety stock required under a skewed distribution and that required if forecast errors are normally distributed.

What About Croston and the Two-Model Approach?

Some have realized the serious inadequacy of using a normal demand distribution for modeling slow-moving item behavior. The most common remedy has been to create a second model explicitly dedicated to "slow movers" based on one of the many derivations of Croston's method. However, this method does not offer a reliable way to set safety stocks or to target a specific service level, particularly for slow moving items.

Moreover, the two model approach requires an a-priori classification of the SKUs, which separates the normal ones from the intermittent ones. Products are split into two

separate categories, making use of two entirely different modeling techniques for the two categories. However, real world inventory behavior presents an infinite variety of different types of behavior, not just two extremes. The modeling technologies used must take this into account.

Succeeding in an Intermittent Demand Environment

Customer demand has become more geographically scattered, occurs in smaller quantities and is calling for more specialized products. In this environment, to achieve high customer service levels, you have to master the shape of the demand distribution and, in particular, its right-hand tail. You must avoid shortcut simplifications and approximations in the modeling that will prevent you from getting close to your goal. Achieving a high service level in an intermittent demand world requires technology that can master demand behavior and probability distributions across a very wide variety of demand patterns. This technology should have the following capabilities:

- Reliable demand modeling that automatically and seamlessly adjusts all relevant statistical parameters across a wide range of SKU behavior types; that understands all elements of demand uncertainty and considers not only forecast error, but also the "intrinsic demand variability" caused by order-line frequency and order-line size distribution.

- Advanced inventory modeling that provides highly reliable descriptions of statistical inventory behavior by eliminating the gross approximations of traditional inventory management theory.

- Logic for demand signal propagation that combines the superior capabilities of both demand modeling and inventory modeling to model the impact of replenishment policies and constraints at each echelon of the supply chain. A solution with these characteristics assures good adherence between plans and actual performance, reduces the need for emergency intervention, and enables disciplined inventory and service planning processes.

Wrapping it Up

We have seen the sequence of problems that make intermittent demand and the long tail environment challenging. First, intermittent demand has intrinsically high variability and probability distributions that are highly skewed. Classic demand and inventory models do not perform well in this environment. Finally, this shortcoming in business systems causes shortages that inevitably lead to inefficient manual intervention and create some degree of bullwhip effect. Inventory becomes seriously misaligned, and company business goals cannot be achieved.

The alternative is to have systems and processes in place that effectively forecast and manage the long tail. The key to mastering intermittent demand is fairly basic and yet very critical. You must accurately and reliably model demand both in the “head” and “tail”. It must be done without taking convenient shortcuts which may simplify the problem but which lead to inaccurate models and output erroneous decisions. This “no-shortcuts” approach will also encourage system confidence with users and planners. It will reduce manual overrides to a minimum and eliminate a key source of bullwhip amplification.

By mastering the full probability distributions of both demand and inventory across a wide range of possible types of behavior, unprecedented efficiency and service level excellence are achievable, even in an increasingly challenging world full of intermittent demand.