Alloy

The No-Nonsense Guide to Forecasting

"How can we improve forecast accuracy?"
"What does best-in-class look like?"
"What are the planning best practices for my industry?"

These are common question we get from consumer goods brands, not to mention topics that immediately catch the attention of anyone in sales operations, planning, or supply chain. Forecasts are never as accurate as you would like, and despite having more sophisticated technologies like AI that we can now throw at the problem, little has changed. A 30% error rate[1] when forecasting SKUs just one month ahead is average in the retail industry, based on IBF surveys.

Getting to single-digit error rates can feel like trying to chase the holy grail. You add more factors into your models, test increasingly sophisticated models based on the latest in data science, and build the business case[2] for a demand management investment to move beyond Excel spreadsheets[3].

Through it all, though, there's a tried-and-tested conclusion you should never forget: forecast accuracy is driven by true demand.

No matter what methodology you use, or what types of products you make, a best-in-class forecast is based on unconstrained demand—the amount that you would have sold if the product were always available on the shelf (physical or virtual[4]), when and where a consumer wanted to buy it. It's not that techniques like machine learning and taking into account causal factors like weather don't help; they're just not the biggest difference makers.

Key Components of the Demand Forecasting Process for Consumer Goods

We've written before about the three key principles for modern demand forecasting: (1) use an integrated approach, (2) keep methodology transparent, and (3) make results actionable. If you haven't already, we recommend reading the white paper[5], as these principles form a strong theoretical foundation for continuous forecasting to provide critical, real-time input for decisions across the organization.

Here, we'll walk through best practices for a best-in-class demand forecasting process in 3 key steps:

- Data preparation
- Baselining
- Forecast adjustments

Data Preparation

As we've said before[6], when it comes to data, it's "garbage in, garbage out." Cleaning and harmonizing data from disparate sources can be a time-consuming—not to mention mind-numbing—task, but is a fundamental requirement if you hope to have accurate forecasts. All data elements should be at the same level, either gross or net of returns, and as granular as possible, ideally down to the SKU/store/day level to allow you to be flexible when building your models.

The 5 Data Types Needed for a Best-in-Class Forecast

- Sell-through units
- Out-of-stocks
- Channel inventory
- Shipment units
- Weeks of Supply target

There are at least five types of data needed for an actionable forecast based on unconstrained demand:

- **Sell-through units.** This type of data can be hard to get, but is the most important, because getting to true demand[7] starts with knowing sell-through. For any brands who sell through retailers, using sell-in as a proxy for demand does not set the team up for success. Of course, there are tradeoffs: there may be time delays in receiving the data from partners, the data may be incomplete or inconsistent across retailers, and there may even be a cost associated with obtaining the data.

  However, getting to true consumer demand is worth the extra effort and investment given the resulting forecast improvements. You can start with the 80/20 rule: typically, a small number of retailers make up a large percentage of total sell-through, so start with the big ones and extrapolate. If the biggest retailers are generally representative of mix and seasonality, it's likely good enough.

  Sell-through data should be in units, at a product level. Having retailer-store level data is best for analysis and understanding geographic differences, but not a requirement for aggregate forecasting.

- **Out-of-stocks.** Along with sell-through, out-of-stocks are needed to understand unconstrained demand. If you had any out-of-stocks, they likely led to lost sales, or in a best case scenario, the consumer bought a similar product from your brand. Either way, the sell-through is not fully representative of true consumer demand, which will bias your predictions, causing you to underestimate future demand and risking additional stock-outs.

  Looking at out-of-stock data in the context of regular demand patterns (when products are fully in-stock), you can derive unconstrained demand. Similar to sell-through data, out-of-stocks should be in units, at a product and store level.

- **Channel inventory**. Knowing exactly how much inventory is in the channel, e.g., in retailer DCs and at retail stores, is critical to understanding future shipments. If, for example, sell-through softens and weeks of supply build up above your target, you can anticipate that shipments will be softer than seasonality would call for. Channel inventory should be in units, at a product level, by retailer. It's okay if it's aggregated, but the more granular, e.g. , by region or location, the more tactical you can get.

- **Shipment units.** Just like you need both sell-through and out-of-stocks, you should have both inventory and shipment data, too. The two are needed for a complete understanding of your total inventory, including on-hand and in-transit. Make sure you're accounting for 100% of shipments and gathering data at the same level as the sell-through, at a unit and product level, broken down to DCs/regions and by retailer.

- **Weeks of Supply target:** Both manufacturers and retailers should have a goal for channel weeks of supply that balances working capital needs, storage costs, out-of-stock costs, shelf life, and other factors. This goal may vary throughout the year, so it's not a one-time exercise. For example, during a new product launch, you might have a higher weeks of supply target, but during periods of transition or low seasonality, you might want to bring that target down to reduce the risk of markdowns.

Gathering and preparing high-quality, timely data for input is also important. Working with the most up-to-date data available ensures you are taking into account the latest sales and inventory to produce the most relevant forecast. Long lag times between receiving data and producing a forecast are often a source of frustration, both for the team developing the forecast and the teams using it. A best-in-class process minimizes this frustration by pulling data as soon as it becomes available and dynamically updating baseline forecasts accordingly.

Once you have the right data, it's time to develop a baseline demand forecast and size the effect of adjustments on top of that. A forecast can be decomposed into parts, generally divided into the baseline forecast and adjustments to it. The baseline typically includes seasonality and trend, which tell you how sales in one month typically compare to sales in another month. The adjustments on top include events and changepoints, which reflect your future plans.

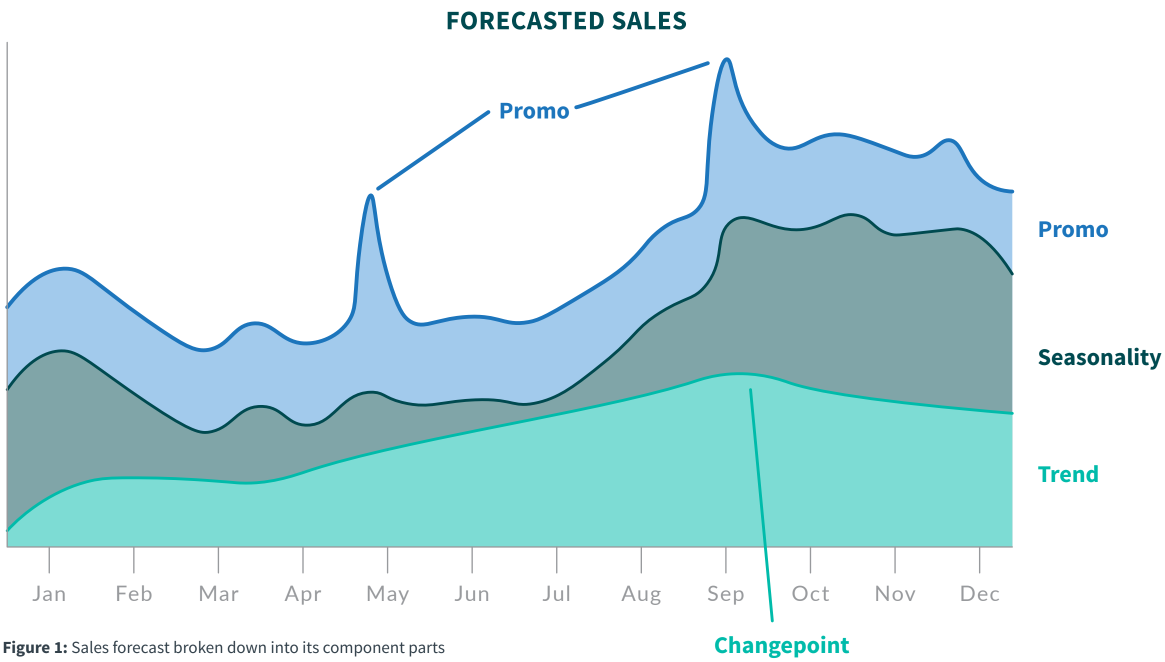

Figure 1: Sales forecast broken down into its component parts

These high-level steps are essentially what different forecasting models[8] mimic. From simple historical averages to machine learning, the goal is to determine historical demand and seasonality, taking into account different causal factors using different statistical methods, to arrive at the best baseline and impacts of different factors.

> "***Consumer behavior is already unpredictable****, and when we add in events like new product launches, promotions, and inventory shortages, everything gets really messy.*"

Building the Baseline

The first step is to create a “clean” version of history, adjusting for out-of-stocks and event outliers, to understand the typical seasonality for your products.

As mentioned above, you want this historical data to represent unconstrained demand, not just historical sales. By interpolating over out-of-stock occurrences, the historical baseline will more accurately reflect true consumer demand.

Consumer behavior is already unpredictable, and when we add in events like new product launches, promotions, and inventory shortages, everything gets really messy. To develop an accurate forecast, you’ll need to clean up that mess by handling these outliers as best as you can. Add and subtract the impact of all the past events you can think of that would affect consumer demand for your products, from an influencer unexpectedly promoting your brand to major holidays. This process will be a combination of art and science; make sure to track your assumptions and use as much data as possible to determine the right sizing of these events and justify your decisions.

Once you have a historical baseline, the seasonality will become more apparent. Look back as many years as you can to identify consistent seasonality patterns from year to year. For newer products or product categories, you may need to extract seasonality patterns from similar products by leveraging syndicated market data. In most cases, best-in-class forecasting leverages some level of triangulation between multiple data sources to build a seasonality curve.

Then, you can develop future baseline forecasts by determining an anchor month (usually the most recent month’s actuals) and applying the seasonality and trend to it for future months. For example, let’s say that your July actualized at 40,000 net units and we know based on the seasonality that August is usually 90% of July and September is usually 130% of August. Your baseline forecast is a simple calculation using this information: your August forecast is 36,000 units (90% of 40,000) and your September forecast is 46,800 units (130% of 36,000).

	Seasonality and Trend vs. Prior Month	Actual/Baseline Forecast
July	n/a	40,000
August	90%	36,000
September	130%	46,800

Figure 2: Sample baseline forecast calculations based on seasonality and trend

However, even the best baseline is just that—a basis for your future forecast. At this point, it is like a car with no engine in it; it looks like a car from the outside, but won’t get you very far.

Forecast Adjustments

On top of the smooth variation in the baseline forecast from trend and seasonality, there will be adjustments caused by special events or changes in the dynamics of the product. This part of the process is where the most variability can occur and the most scrutiny and review should happen.

You'll need to add a couple of things to the future baseline and reflect their impact on demand appropriately.

1. **Changepoints.** Changepoints in your forecasts are points in time following which the baseline activity shifts in some way.
   - **Growth.** Seasonality extrapolates growth based on a current number, but if you plan to make a significant change, like doubling the number of stores where your product is sold, the impact will need to be forecasted. Conversely, there may be a softening of the industry overall or new regulations that could hurt growth.
   - **End of Life (EOL) or New Production Introduction (NPI).** Are products going end of life? Are you launching new products? These are hands-down the most challenging events to forecast for and require significant forecast adjustments to account for winding down or winding up sales and operations for the products in question.

2. **Events.** Events are more discreet than changepoints and have a short, fixed period when the events occur:
   - **Promotions.** Will there be a promotional price, rebate, gift with purchase, or some other promo activity that should drive demand over the typical seasonality? We stripped promotions out to create a "clean" version history and develop a baseline, but now need to know how to add them back in and account for the expected increase in demand.
   - **Brand marketing activity.** Although extremely hard to determine impact to sell-through, it is important to have some agreed-upon methodology for accounting for above-the-line (ATL) or brand marketing campaigns, in addition to below-the-line (BTL) promotions. Especially if the level of brand marketing activity will be increasing or decreasing compared to the previous year, it will have an impact on demand.

Once all of these factors are added in, you'll have a forecast of future demand. To get the sell-in forecast, it is just a math exercise of looking at the channel inventory, shipment units, and weeks of supply goal to determine the orders customers should be placing to meet the forecasted demand.

Your forecasts are now ready for your demand planning and sales and operations planning (S&OP) processes. You have a best-in-class forecast, based on unconstrained demand, that you can confidently share with cross-functional teams and leadership to develop and refine short- and long-term plans.

There will, of course, be many company-specific nuances to your forecasting needs, but overall, this process should lead to an accurate and highly-defensible forecast because it puts consumer demand data at the core. Alloy's forecasting solution automatically integrates up-to-date demand, inventory, and shipment data at the store-SKU level to deliver dynamic baseline forecasts and forecast adjustments, so you can focus on planning.

To learn more, please visit www.alloy.ai/forecasting.

Footnotes

[1] http://demand-planning.com/2018/06/22/what-are-the-benchmarks-in-retail-forecasting-accuracy/

[2] https://www.gartner.com/smarterwithgartner/how-to-build-a-business-case-for-demand-management-investment/

[3] https://blog.alloy.ai/the-most-popular-planning-software

[4] https://ref.alloy.ai/nola-infographic

[5] https://ref.alloy.ai/modern-demand-forecasting

[6] https://blog.alloy.ai/how-executives-can-better-understand-forecasts-to-inform-decision-making

[7] https://blog.alloy.ai/not-all-demand-is-true-demand

[8] https://blog.alloy.ai/choosing-the-right-demand-forecasting-model

About Alloy

Founded in 2017, Alloy is the modern analytics and planning software solution for consumer-driven brands. The purpose-built platform enables leading and fast-growing consumer goods companies to evaluate, predict, and respond to true demand with agility and efficiency. With Alloy, brands can analyze sales and inventory at both macro and micro levels, down to the individual store and SKU, to accelerate insights, drive growth, and optimize the supply chain.