Fundamentals of Modern Demand Forecasting

3 Key Principles For a Rapidly Changing World

Author: Aaron Hoffer, PhD

Alloy

Contents

Alloy

Historically, retail forecasting has been a static input into a manual process for sales operations and supply chain management. Consumer goods companies and their retail partners created forecasts for predicted sales (or, in many cases, sell-in), updated them on a quarterly or, at most, monthly basis, and then made decisions accordingly. Due to how time-consuming and slow the process was, forecasts were not revisited until it was time to evaluate results at the end of the quarter—too late to make any meaningful corrections.

In today's fast-moving retail environment, this old mode of forecasting is no longer sufficient for brands that want to succeed and thrive. Instead of an isolated, one-time exercise, forecasting now needs to be a process that happens continually, providing critical, real-time input for decisions across the organization.

Mission-critical forecasting

Shifting from old practices to new ones isn't just about pulling in more data or using new methodologies; it's about creating a new approach and mindset for forecasting. To make this transition, forecasting must become a core component of how organizations operate, critical to the success of all the departments that depend on its results. To achieve this goal, leading companies follow three key principles when implementing forecasts: use an **integrated** approach, keep the methodology **transparent**, and make results **actionable**.

Use an integrated approach

Using an integrated approach when forecasting means using multiple sources of input data and incorporating sound data science methodologies when working with that data. Brands must be able to seamlessly connect the various data sources a company has access to across its supply chain, including point-of-sale, inventory, bill of materials, logistics data, and more. Then, brands need to apply good data science practices to that data, including normalizing it and evaluating multiple models to assess best fit.

Keep methodology transparent

With the increasing complexity of modeling tools and data inputs/outputs, it can be hard to see how all the components of a forecast fit together. Having clear explanations for the way each component impacts the overall forecast makes it easier to fix problems, explore multiple scenarios, and have confidence in the end result.

Make results actionable

This is where the ROI of forecasting comes in. Once the hard work of creating a forecast is done, how is it used to improve the business? For course-correcting actions based on clear triggers, such as exceeding a pre-set out-of-stock threshold, allow the response actions to be activated automatically by an open platform. When response actions are more ambiguous, have collaboration and ticketing workflows that allow users to track progress as they work together to solve problems.

Through the rest of this white paper, we'll review the industry and how it has evolved over time. Then, we'll dive deeper into these three principles, adding more detailed discussion of how to put them into practice so you can follow a similar framework to maximize the value you get from forecasting.

A quick note: while there are many metrics that can be forecasted, we focus on sell-through (demand) forecasting because we believe this is the most critical forecast; we'll explain more below.

Evolving industry practices

With the critical role forecasting plays today, the shortcuts taken to simplify the process in the past will no longer provide the needed rigor. The main culprits are:

- Monthly/quarterly forecasting: these timescales are too long, because brands today need more rapid feedback
- Only forecasting top products or regions: inventory costs associated with the long tail can add up and significantly impact profitability
- Forecasting at only one node of the supply chain: the complexity of a multi-tiered supply chain network makes it necessary to forecast along each step, from manufacturing to end-consumer sales

To help illustrate what needs to change, it's helpful to first take a look at where the industry has come from and what new opportunities are now available to take advantage of.

Evolution of data

It's now the world of big data, where the volume and velocity of available data has dramatically increased from what was readily available in the recent past. As a result, the possible inputs for forecasting models—and even the models themselves—have increased in complexity.

Old practices	New practices	Benefits
Aggregate sales	Store and SKU- level data	Avoid incorrect trends that may reverse when disaggregated[1]
Average metrics	Derived store/SKU-level metrics, like out-of-stock %	Understand true cost of out-of-stock occurrences
Forecasting from sell-in data	Forecasting from sell-through data	Reflection of true demand
Top SKUs/Locations	Automated forecast of long-tail for approximation of demand	Capturing long tail of growing number of SKUs reduces inventory costs[2]
Spreadsheets, EDI, separate systems	Unified data repositories	Data normalization, end-to-end analysis

Table 1: Summary of changing data practices and benefits of more granular data and automated processes

While these new practices, see Table 1, have created new opportunities, they also present new challenges in data management and maintaining data quality, which is why integrating advanced data science techniques into forecasting processes is now even more important.

Evolution of methodologies

With today's increased availability of high-quality, granular data and the increase in computing power and storage, machine learning and artificial intelligence methods have been added to the forecasting toolkit. But while these newly popular methodologies have been making a big splash in forecasting, that doesn't necessarily mean older methods have to be discarded entirely. In many cases, since each of these methods evaluate different aspects of the time series, models built using various methodologies can be combined into an ensemble method that produces an improved forecast by taking many factors into account. This is another reason why the integration of data science is important for today's forecasting.

Table 2 below highlights the major classes of forecasting models being used today:

Model class	Example methods	Pros	Cons
Historical Average	Simple Moving Averages, Holt-Winters Exponential Smoothing	Simple, can easily be implemented everywhere	Without regressors, the prediction can be slow to react to changes in demand or be overly responsive to outliers
Time Series with Added Regressors	Seasonal ARIMAX, Generalized Additive Model	Well understood methods with consistent variation. Can consider multiple/complex time variation along with other drivers	Expects particular seasonal or time-dependent structures, limited history requires too many assumptions
Machine Learning and Artificial Intelligence	Gradient Boosted Machines, Neural Networks (LSTM-RNN, CNN) Support Vector Machine	Non-linear and complex relationships can be discovered without a need to pre-select the model type or make assumptions regarding external factors	Requires a lot of data and initial investment in setup, may require parameter tuning, and results may not be easily interpretable (this is why having a forecasting partner is valuable)

Table 2: Summary of forecasting model types that are applicable to different types of data due to different tradeoffs

Different models can be combined to create a more complete picture of demand

The forecast lifecycle

Forecasting can be valuable to nearly every part of the organization, but we'll focus on supply/demand planning and marketing/sales. As forecasts have become more integral to different processes, the forecast lifecycle is extending to include more partners and becoming a more collaborative decision making process.

After forecasts are generated, they're inputted into planning systems to help calculate supply and demand planning. However, it is an iterative process, whereby the effects from different variables can be considered in isolation and drive decisions made by the sales, marketing, and operations teams, which can in turn impact calculations of supply and demand planning.

As shown in Figure 1 with transparent forecast models that clearly indicate the effects of adjusting different variables, the marketing team could explore how increasing or decreasing marketing spend in certain geographies might affect sales. Similarly, with continually-updated forecasts, the operations team can make key allocation decisions in the event of a manufacturing disruption.

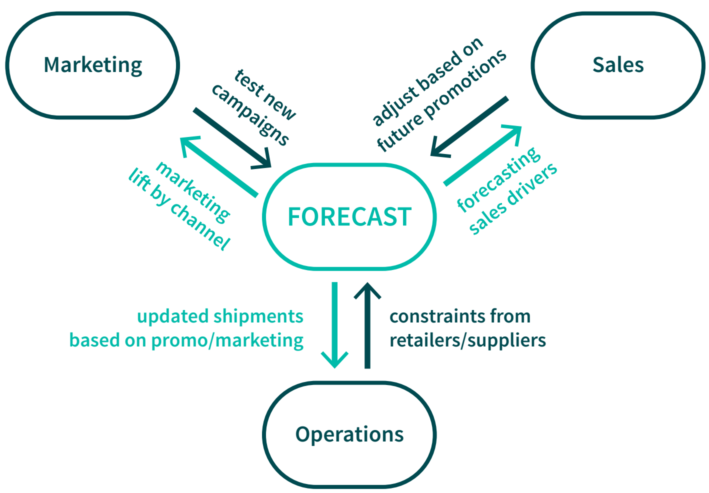

Figure 1: Example iterative process between different parts of the organization and their interactions with forecasting results. Forecast gives an estimate of lift from marketing campaign and marketing can test impact of new campaigns on demand. Sales sees demand drivers, e.g., promotions, on forecast and can test impact of new promotions on demand. Operations receives order and shipment estimates, and can set constraints around lead time and supply for different suppliers and retailers

In addition to the demand forecasts created in-house, retailers and other trading partners may create their own forecasts for sell-through demand, sell-in shipments, or inventory. These may not be as accurate or informed, but they can still help shape business strategy.

For example, a retailer's sell-through demand forecast indicates how they plan on ordering and how they will stock their individual stores. If they also provide a shipment or sell-in forecast, it can be compared to the sell-through forecast to estimate how they're looking to stock up before promotions or the holiday season. This information can help brands better plan for inventory and timing, reducing the number of forced markdowns or clearance items at the end of a season or product lifecycle.

Forecasting principles in practice

Integrated: Develop a tight feedback loop for data collection and modeling

In forecasting, “integration” applies to two key areas: data collection and data modeling. For data collection, you must collect all available forecasting data sources, including from 3PLs and suppliers, POS, sell-in, and more. Even if the model ultimately used for forecasting doesn’t incorporate every available data source, it’s still important to have all options available.

Once data is collected, you need to incorporate sound data science processes[3] (see Figure 2) in order to glean the most accurate and valuable insights possible. For instance, after raw data is obtained, you’ll need to fill in missing values, normalize attributes, estimate unconstrained demand, and flag outliers (like promotions). In addition to cleaning and normalizing the raw data, you’ll also want to compute potential drivers (e.g., % discount) that can be used as features in your model. While this can be time-consuming, it’s critical setup work needed to build high-quality forecasts[4].

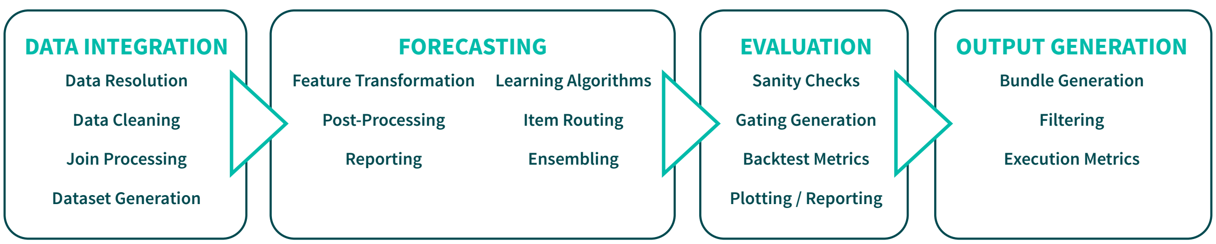

Figure 2: End-to-end process for developing forecasts using sound data science processes, from data preparation to model implementation[3]

You also need a constant feedback loop on the best methodologies’ fit quality. Once you’ve finished your models, you’ll want to apply all the processes outlined above to them. In most cases, a single method or model doesn’t result in the best forecast; instead, using an ensemble gives optimal results[5]. Copying from Excel doesn’t work (large error rates are seen in spreadsheets[6]); you’ll want robust automatic model evaluation and quality detection to ensure consistent forecasting.

Transparent: Institute data democracy and model transparency

The teams that are the end consumers of the models (e.g., planning, marketing, finance) may not be the same as the group that’s building the models. Therefore, it’s critical for data scientists and analysts to be transparent about what they’re building. For every model that might be used, the effects of each variable need to be explained, the performance of the model (and how that may change over time) should be made clear, and any assumptions used when building the model should be listed.

When building sales forecasting models, there are many choices that need to be made to drive results. There are also many different methodologies to choose from, each with its own pros and cons. Furthermore, different models work best[7] for different industries and product types. Explaining how and why methodologies were chosen will help give all business users context.

While you will want to determine and then implement the best model, it's a good idea to still document other methodologies, as they can yield helpful insights. For example, different models might showcase the impact that certain variables have more clearly than others. Understanding these factors helps everyone understand the levers they may be able to manipulate to shape demand. For example, Figure 3, shows how different forecasts can be compared alongside sales for easy comparison.

Another component of transparency is continuing to collaborate with cross-functional users. For example, operations people on the ground have particular insight into the products and your customers, and can help decide what data you can source and use to improve the models. To that end, it's good to make the data integration processes described previously as transparent as possible. Data democracy[8], the ability for people across the organization to read and understand critical data, helps bring more adoption and confidence in the end results generated.

With all the potential areas that forecasting can impact, it's very important to make it clear to stakeholders across the business how different models work and why certain choices were made. Not only does this improve confidence when relying on models, but it also may lead to new insights driven by area-expertise that ultimately improve the models.

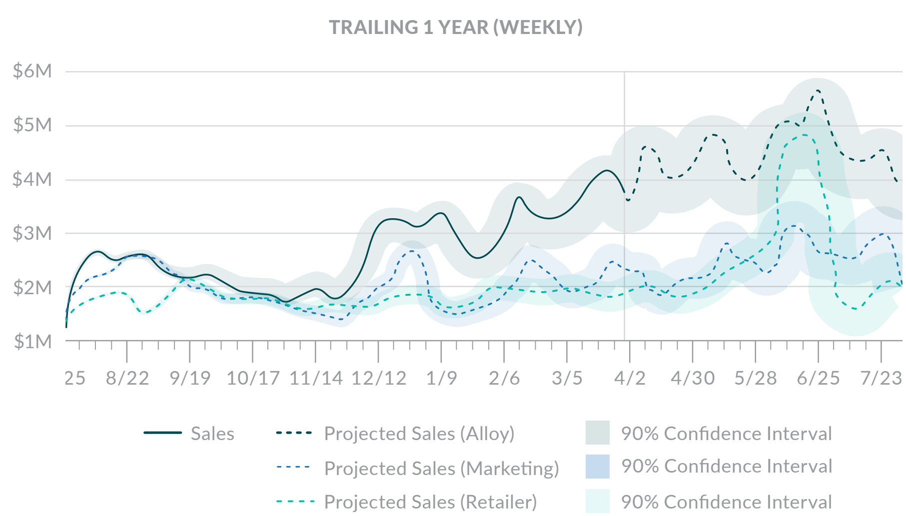

Figure 3: Example output comparison between different model types and actual sales

Actionable: Implement different forecast components in different outputs

To help make forecasts actionable, there are a few key points to keep in mind. First, don't overwhelm your team with numbers; you may be able to produce thousands of granular forecasts for each SKU, but you only want to show things that are important (i.e., manage by exception). While everyone should have access to the raw data and forecasts themselves, no one should have to parse through an overwhelming amount of information in order to draw a conclusion. Once the issues have been identified, the raw data can be readily available to dig deeper, but is not the primary focus.

Second, while innovating and adopting new best practices is valuable, it's also important that forecasts can be compared against one another over time, so make sure that new methodologies and variables are documented and accounted for properly.

Lastly, don't make your models more complicated than they need to be—in today's environment, continual forecasting is key, and adding variables that are hard to predict and don't add much value (for instance, weather data) will only slow you down. Your forecast should be easy to integrate into other systems and produce results that enable timely decision making. Figure 4 demonstrates a hierarchy of needs for forecast usage:

Not all applications need the same level of detail, but detail should be made available if needed

FORECAST IMPLICATIONS

- Highlighting products at risk for out-of-stocks given shipping delays
- Optimizing inventory distribution across channels when supply is limited

FORECAST DRIVERS

- Price sensitivity to promotion type when the sales team is designing new promotions
- A/B testing different consumer marketing messages and how they change sales lift

FORECAST VALUES

- Forecasted sales vs. actual when account managers are in CPFR conversations with their buyers
- Deriving estimated purchase orders from a retailer for planning and operations

Figure 4: Hierarchy of forecasting output detail, from the highest level (forecast implications) to the most detailed level (forecast values)

Make forecasting a core competency

Forecasts can affect many areas of a brand's business, so it's important to invest the time and effort into doing them well. Following the three principles outlined above will help put you on the right track to making timely, useful, and accurate predictions. In some cases, the steps needed to achieve these principles may be a departure from the way things have traditionally operated, but it is well-worth taking the time to make these changes. And if the data science feels more complex and time-consuming than you have resources for, there are experts and tools that can help.

Footnotes

[1] http://www.statisticshowto.com/what-is-simpsons-paradox/

[2] https://ibf.org/knowledge/journals/forecasting-performance-for-north-american-consumer-products-1002

[3] Böse et al. 2017 http://www.vldb.org/pvldb/vol10/p1694-schelter.pdf

[4] https://www.nytimes.com/2014/08/18/technology/for-big-data-scientists-hurdle-to-insights-is-janitor-work.html

[5] https://ieeexplore.ieee.org/document/8073492/

[6] http://panko.shidler.hawaii.edu/SSR/Mypapers/whatknow.htm

[7] https://insidebigdata.com/2018/09/03/demand-forecasting-mistakes-retail-industry/

[8] https://www.forbes.com/sites/bernardmarr/2017/07/24/what-is-data-democratization-a-super-simple-explanation-and-the-key-pros-and-cons/#121248ca6013