



Demand Forecasting with Advanced Analytics

Executives can apply these real-world lessons
to maximize their investments.

**By Cesar Brea, James Anderson, Robin Bartling,
Rodrigo Mayo and Florian Mueller**

Cesar Brea, James Anderson and Robin Bartling are partners with Bain & Company's Advanced Analytics practice. Cesar is based in Boston, James in Sydney and Robin in São Paulo. Florian Mueller is Bain's Advanced Analytics practice leader for Europe, the Middle East and Africa, and is based in Munich. Rodrigo Mayo is a partner with Bain's Retail and Strategy practices in Mexico City.

The authors wish to thank, for their meaningful contributions and lessons, Diane Berry, Sanjin Bicanic, Anton Bossenbroek, John Groves, Sumner Makin, Sriram Narayanamoorthy, Carla Nasr, Ben Rollins and David Hess.

Net Promoter Score® is a registered trademark of Bain & Company, Inc., Fred Reichheld and Satmetrix Systems, Inc.

Copyright © 2018 Bain & Company, Inc. All rights reserved.

At a Glance

- ▶ Companies can use advanced analytics to improve demand forecasting, but only if they manage the process well.
 - ▶ The most successful companies focus on results, treat forecasting as an operating process and build forecasting tools in-house only when strategic.
-

A major consumer products firm—selling perishables through grocery stores—chronically put 10% more on the shelves than it sold, a buffer that overcompensated for the risk of running out of stock. Over time, this added up to an expensive miscalculation. What kept going wrong?

After considering the problem, the company identified three causes. First, its demand forecast, while reasonably accurate for high-volume items, was significantly off for a large number of items that accounted for most of the volume sold. Second, the order management interface that the firm provided to its sales reps took too long to teach and learn. Finally, the various layers of sales management and planning support weren't always helpful, tending to simply adjust up the size of an order.

Each of these challenges suggested a solution: a better forecast, a better interface to turn that forecast into orders, and a change in the management and planning support structures for sales reps. The firm ultimately addressed all three. It improved forecast accuracy by more than 30% compared with its prior approach (*see Figure 1*). It built an ordering interface with a Net Promoter Score® near 60, compared with a score of negative 8 for similar applications. And it reduced the number of people involved in the planning process by two-thirds.


When tested in four major-market pilots representing the majority of the company's business, this new approach reduced overordering by one-third while actually improving growth. The company is now rolling out this comprehensive solution more broadly. Once it is in place across all markets, the company expects an annual nine-figure increase in operating profit, for a recurring threefold return on its investment.

This brief is about the forecasting part of the story. For leaders unhappy with their current demand forecasts, whether executives or analytics team managers, we highlight three useful lessons:

- Focus on results, not sophistication.
- Treat forecasting as an operating process, not a modeling exercise.
- Build when forecasting is strategic; buy when it isn't.

Figure 1: With better forecasting, a consumer products company reduced expensive overordering

Mean absolute percentage error, weighted average by week



Notes: Since launch; adjusted for corrected data
Source: Bain & Company

Lesson 1: Focus on results, not sophistication

No one would disagree that the point of doing advanced analytics is to get results. And yet, many boards instead pressure CEOs to develop a broadly thought-provoking AI strategy. Some data scientists can seem more focused on using neural nets for their own personal development than for achieving business results. In truth, progress toward an objective, quantifiable business goal is the coin of the realm. Earn some and you get to keep playing.

In the push for progress, one approach won't fit all product scenarios. A typical firm's product mix will include perennial products, seasonal products, new products, products with lots of volume and products with less volume. Shelf life, frequency and depth of promotion, and maybe many other factors will vary. Different scenarios may benefit from tailored algorithmic approaches.

For example, when creating a forecast for a new product, you might take an autoregressive integrated moving average (ARIMA) model of the sales history of a different product with similar attributes, and then blend it with a moving average or exponential smoothing approach on the more limited history of the new product. Over time, you might try a neural network approach to discern likely new product adoption patterns from multiple past examples.

For heavily promoted items, you could begin by forecasting base demand and then layer the effects of price promotions on top of that. You may be able to forecast demand for products with a longer shelf life at a higher level of aggregation—say, by product category rather than by sales per store—if you have stocking flexibility. For shorter-lived items where accuracy is more important, however, you'll need to get closer to predicting demand by SKU and by store.

In short, don't expect just one model. There will be many. Think instead about a model platform.

Work top-down. It's tempting to start your forecasting at the most granular level—for instance, at the individual store-SKU combination. But it's more computationally efficient, and often as effective, to work your way down the product or market hierarchies (starting, say, with overall sales, then chain-category sales, then chain-region-subcategory and so on), reserving your efforts at the store-SKU level for “hero” high-volume combinations.

Look for the point of diminishing returns, when increased granularity stops resulting in significant improvements in forecast. And for really low-volume combinations—where, for example, one item might or might not be sold each period—maybe substitute a standing order for the one jar of jam that mostly sells, but sometimes doesn't.

Think accuracy and consistency. Because improvements in accuracy can carry asymmetric value—underordering is more expensive than overordering—being able to consistently walk up to the edge of the store inventory cliff, but not over, is critical. Don't just think in terms of mean absolute percentage error (MAPE), a common measure of forecast accuracy. Also consider the standard deviation of MAPE over time—in other words, how much your forecast accuracy bounces around.

Emphasize relative improvement. There's an expression: “You don't have to outrun the bear, just the guy next to you.” Your job is not just to improve forecast performance, but to do better than the alternatives—whether “naive” out-of-the-box benchmarks from software providers or the results of a Kaggle competition. If you'd like to dig deeper into relative improvement, see Rob Hyndman and Anne Koehler's excellent paper on mean absolute scaled error (MASE) as a better measure to use.

Lesson 2: Treat forecasting as an operating process, not a modeling exercise

Forecasting starts with picking the most important things to forecast; proceeds through getting and improving data, thinking about how you'll scale and maintain your forecasting models, and trusting the output enough to use it; and finally concludes with evaluation of model and business results in order to continuously improve the whole process.

Everyone should agree on the mix of specific target metrics used to assess the forecast. Are you measuring year-on-year forecast improvement, or are you making a concurrent comparison with other forecasting approaches? Are you forecasting unit volume or dollar volume? Are you focusing on reducing overordering or on avoiding out-of-stocks? In many cases, the answer is to do some of both,

balancing these considerations and including others as well. The metrics you emphasize most may evolve over time. The main thing is to be clear and aligned with other stakeholders about where your cannons are aiming. It seems obvious in principle. In practice, it's often not.

Better data is probably the single most important investment you can make. A consumer products company with operations in more than a dozen countries recently discovered this when it began training its forecasting model. The company split its data, training the model on one subset of that data and then comparing those results to the data not used in the training. Initial results were promising: The new modeling approach produced a 30%–40% improvement in MAPE. But, when real-world tests began, the improvement was just half as much as it had been using estimated price and promotion data. The problem was that the initial data wasn't accurate. Once the team went back and fed better data into the model, they started to realize the potential they'd first seen. They designed the user interface so that reps using the forecast to place an order could enter up-to-date pricing and promo information, ensuring the best possible data for future modeling. Good forecasting operations are often paired with initiatives to improve data quality.

The company further improved the accuracy of the model platform by including quickly available scan-based trading reports and cross-elasticity variables, such as whether it helps or hurts product Y when product X goes on sale.

The main thing is to be clear and aligned with other stakeholders about where your cannons are aiming. It seems obvious in principle. In practice, it's often not.

Mind how you cook your data ingredients. At scale, you will produce many forecasts, very frequently, in very short time frames. If something goes wrong, it will be expensive and messy to fix. You can anticipate this. Know what data you expect, and make sure it's coming in. Have guesses you can apply when the data is missing or looks wrong. The fancy term for this is “interpolation,” an approach similar to the nautical triangulation sailors once used, before GPS, to guess where they were. Make sure you produce the number of forecasts people are expecting. Automate checks to make the model platform work well at scale. When one thing looks incorrect, trace it to its root, and then double-check the other things that look like it. Create a VIP lane with extra care and exploratory data analysis for high-volume and other strategic SKUs, such as key new products, and related stores—items for which you want to be sure nothing automated goes off the rails. Staff accordingly!

On the path to more predictive power for each forecast, two or more models are often better than one. As you go from simple models, like moving averages, to exponential smoothing techniques, to

ARIMA and onward, make room for an *ensemble* of these (essentially, try several and average their answers) before moving on to more complex individual algorithms. And even once you have deployed sophisticated approaches, keep them honest by benchmarking them against ensembles of simpler techniques. An advantage of using multiple approaches is that if one of them occasionally provides a weird recommendation (which you should investigate, of course), the others will reel it in and together provide a plausible forecast.

Test in the real world. Once you've established a model platform's accuracy and consistency using a test sample (including backtesting on prior periods), the next challenge is evaluating its performance at scale in real-world situations. Model platforms that don't pass this test—which may be great for a specific data set but don't generalize well (“overfitted” in industry parlance)—may need to be simplified or rebuilt frequently. Reflect on what any degradation in forecast performance may be telling you about new developments in old data. It could be revealing shifting customer habits or the unexpected effect of a new competitor, for example. Also, measure the full dollar impact of your process changes, not just narrow gains in forecast accuracy. That may lead you to spend your next dollar on something besides the model, like a better interface or better coaching for users.

About those pesky users: They just won't use a model they don't understand. That's another reason to start simple. It's also important to track reps' usage of the ordering interface, not just the forecast, and the outcome of the order. Surprisingly, people don't always do what you tell them. Often, it's not enough to present a forecast for reps to use; you need to contextualize the forecast with data on recent sales, or from a similar period, for example. When they adjust something, capture why they did it, and show them the impact of their adjustments, good or bad. Here's an especially radical suggestion: Involve them on your model platform development team from the very start of the program. Form an advisory board to support ongoing improvement, and include them on it, too. They will have ideas for features to add to your data set well beyond what you can imagine (e.g., “This store is near the site of the county fair; when the fair is on, everything shuts down”).

Often, it's not enough to present a forecast for reps to use; you need to contextualize the forecast with data on recent sales, or from a similar period. When they adjust something, capture why they did it.

Vary how you evaluate your model over time. Obviously, performance should always be the dominant measure, but overweight transparency early on. That way, if there is a wacky result, you can look at variables and coefficients of the model that produced it—perhaps spotting a weird set of selected variables or a strange coefficient—and then track down what's behind the issue. And afterward, when you've begun to run out of performance improvement ideas for your model, your next

dollar of ROI might well come from making the whole model platform run more efficiently, to save on your server bills.

While you might start with a multiple regression approach like ARIMA, as described earlier, some companies have found performance and efficiency improvements using tree-based algorithmic approaches designed more for classification. However, greater sophistication doesn't always pay off. One company that pushed into deep learning realized that the incremental performance boost wasn't worth the diminished understanding of the predictions.

Lesson 3: Build when forecasting is strategic; buy when it isn't

If you have an internal data science team, they will have a bias to build, because, well, it's what they do. Conversely, most managers don't want to build, because it's scary, and it's easier to have a vendor to choke than an internal team to "manage constructively." Beware of both in the extreme.

Rule one when deciding whether to buy or build is "know thy requirements." If a better forecast is strategic to your business—for instance, if you make your money on thin margins and fast turns—then you need continuous improvement of forecast error and a lot of control over your forecast platform. Build it. Also know that 80% of your effort will go into building the data pipeline that feeds the models and that checks and exports the results. When you buy a commercial modeling tool, you still have to do this work, and no amount of algorithmic sophistication will bail you out of a bad data pipeline (see the pricing/promo advice above). Be sure to focus on the full cost of your solution, and budget accordingly.

Also know that 80% of your effort will go into building the data pipeline that feeds the models and that checks and exports the results.

You may be thinking about buying a forecasting tool in the context of a broader platform upgrade, such as an enterprise resource planning (ERP) system. The main events in an ERP are how it models or represents your business, its data model and the associated business rules. Often the forecasting tools offered alongside those rules are relatively inflexible black boxes with fairly basic forecasting models inside. If customizing an ERP is strategic for you, much of your budget will go to modify the data pipeline that feeds the model, and fewer resources will be available to swap out and evolve the forecast engine you bought.

Open-source options are often the right starting point, but there's no need to be dogmatic about sticking to them. Commercial solutions can be the right choice, too. The important thing to remember is that tech strategy follows business strategy.

It's normal to wonder if your organization is capable of owning and evolving a sophisticated forecast platform. Whether or not you choose to outsource your forecasting operations to a modeling shop, you should retain some internal data science and engineering capability. That way, you can better understand and guide what you're getting, bridge the outside shop more effectively to your evolving strategy and organization, and make sure that you retain control over and visibility into the predictions. You should insist on controlling the data pipelines that feed and consume a shop's models; this might extend to all data-quality checks. It's fine if the vendor uses proprietary transformations of raw data and synthetic features as part of its offering, but you should know what those are even if you don't get the source code, so that you can better evaluate what the model is telling you.

Building models is not just an analytic exercise; it's software development. That means you must involve IT, because they govern and support the infrastructure and policies that will direct model development. Even if you're off doing your thing in the cloud, you'll need data they provide. You'll likely be producing outputs they consume. And you'll be complying (or you better be) with security, privacy and usability requirements. These conversations will lead you to broader IT adoption considerations, like training and support. So don't go rogue!



As the poet might have said, had he been a data scientist, "Bliss is it in this day to forecast, but to manage forecasting smartly is very heaven!" If you're a senior executive charged with a key business process that relies on forecasting, we hope these lessons help you ask better questions. And if you're an analytics leader, we hope these suggestions provide you with better answers. In any event, we look forward to hearing about your results.

Demand Forecasting with Advanced Analytics

Shared Ambition, True Results

Bain & Company is the management consulting firm that the world's business leaders come to when they want results.

Bain advises clients on strategy, operations, technology, organization, private equity and mergers and acquisitions. We develop practical, customized insights that clients act on and transfer skills that make change stick. Founded in 1973, Bain has 57 offices in 36 countries, and our deep expertise and client roster cross every industry and economic sector. Our clients have outperformed the stock market 4 to 1.

What sets us apart

We believe a consulting firm should be more than an adviser. So we put ourselves in our clients' shoes, selling outcomes, not projects. We align our incentives with our clients' by linking our fees to their results and collaborate to unlock the full potential of their business. Our Results Delivery® process builds our clients' capabilities, and our True North values mean we do the right thing for our clients, people and communities—always.



For more information, visit www.bain.com

AMSTERDAM • ATLANTA • BANGKOK • BEIJING • BENGALURU • BERLIN • BOSTON • BRUSSELS • BUENOS AIRES • CHICAGO • COPENHAGEN • DALLAS • DOHA • DUBAI
DÜSSELDORF • FRANKFURT • HELSINKI • HONG KONG • HOUSTON • ISTANBUL • JAKARTA • JOHANNESBURG • KUALA LUMPUR • KYIV • LAGOS • LONDON • LOS ANGELES
MADRID • MELBOURNE • MEXICO CITY • MILAN • MOSCOW • MUMBAI • MUNICH • NEW DELHI • NEW YORK • OSLO • PALO ALTO • PARIS • PERTH • RIO DE JANEIRO
RIYADH • ROME • SAN FRANCISCO • SANTIAGO • SÃO PAULO • SEATTLE • SEOUL • SHANGHAI • SINGAPORE • STOCKHOLM • SYDNEY • TOKYO • TORONTO
WARSAW • WASHINGTON, DC • ZÜRICH