



Supply Chain


A Simpler Way to Modernize Your Supply Chain


by David Simchi-Levi and Kris Timmermans

Sponsored by:




This article is made available to you with compliments of Qlik. Further posting, copying, or distribution is not permitted.





PHOTOGRAPHER BERNHARD LANG



SUPPLY CHAIN

A Simpler Way to Modernize Your Supply Chain


How to spend less and accomplish more



AUTHORS

David Simchi-Levi
Professor, MIT

Kris Timmermans
Senior managing director, Accenture




SUPPLY CHAIN



ABOUT THE ART

Bernhard Lang's aerial shots of a beverage production yard explore the repetition and geometry found in everyday life.



Most

executives believe that digitizing a major corporation's supply chain costs tens of millions of dollars. The assumption is that it will be a massive three- to five-year transformation effort—requiring major investments in cloud technology, the installation of RFID tags and readers on every product container and in every facility, the deployment of 3D-printing and robotics technologies, and new instruments on machines on the shop floor to monitor their performance and condition. All that is necessary, the thinking goes, to break down the walls between functional areas and create an integrated supply chain that provides a competitive advantage.

But in our consulting work for a number of companies, we've discovered an alternative. The experiences of these

firms—which include a global fashion retailer, a large manufacturer of consumer packaged goods (CPG), a global appliance maker, and a high-tech company that produces PCs, tablets, and workstations—demonstrate that it's possible to reap substantial benefits by spending a few million dollars on a supply chain modernization that takes 12 to 24 months. In these more moderate efforts firms assemble readily available data; use advanced analytics to understand and predict customers' and suppliers' behavior; optimize inventory, production, and procurement decision-making to cut costs and improve responsiveness; and add some automation to revamp existing processes and introduce new ones.

The secret to the success of this approach lies in three initiatives: In the first firms replace consensus forecasts with a unified view of demand. In the second they move away from a one-size-fits-all supply chain strategy to a segmented strategy. In the third they create a single plan to continually balance supply and demand and identify and respond to deviations or disruptions.

Executed well, these initiatives lead to lower supply chain costs—and higher revenue because of fewer stock-outs and improved service levels (the proportion of orders delivered on time and in full). Equally important, they enable companies to increase customer retention. At the fashion retailer, they helped to boost market share by more than 28% and double operating profit in just three years. The operating and financial gains from the CPG company's initiative paid for its cost in just two years. The high-tech firm saw a 10% to 30% improvement in service levels. And the appliance maker realized a 20% increase in revenue, raised the proportion of

IDEA IN BRIEF

THE CONVENTIONAL WISDOM

Digitizing a company's system for managing its supply chain is a megatransformation project that takes three to five years and costs tens of millions of dollars.


THE REALITY


There is an alternative: Substantial benefits can be reaped from a modernization effort that takes 12 to 24 months and costs a few million dollars.

WHAT IT ENTAILS

Assembling readily available data; using advanced analytics to understand and predict customers' and suppliers' behavior and to optimize inventory, production, and procurement decision-making; and adding some automation to revamp existing processes and introduce new ones.

COPYRIGHT © 2021 HARVARD BUSINESS SCHOOL PUBLISHING CORPORATION. ALL RIGHTS RESERVED.





SUPPLY CHAIN



More-accurate order forecasts clearly translate into a more effective supply plan, which reduces lost sales.

customers to whom it could provide one-day delivery from 70% to 90%, and cut its operating costs by 3% to 4%.

In this article we'll focus on the CPG manufacturer's implementation of the approach. It's a particularly instructive case because of the extraordinary challenges the company faced in addressing the deficiencies of its existing system, which included multiple and time-consuming manual processes, excess inventory, and a large amount of expired and damaged products.

BUILDING A UNIFIED VIEW OF DEMAND

The journey starts by rethinking the demand-planning process. Traditional approaches employ consensus forecasting, in which each function—operations, finance, sales, and trade (which is responsible for marketing, promotions, discounts, and so on)—uses standard statistical techniques, historical sales data, and some external data to generate its own forecast. Then all the functions get together and hash out a compromise uniform forecast.

That process has two drawbacks. First, it takes a long time—typically four to five weeks—to generate the various forecasts and reach a consensus that satisfies all business requirements. By that time the sales data used is old. Second, rather than agreeing on the data and having the analytics produce a single forecast, the people involved typically focus on finding a balance between conflicting forecasts and rely on gut feelings about what drives sales, revenue, and margins.

A much better way to generate a unified view of demand is to start with the sets of data that all participants agree will yield the most accurate picture. The CPG manufacturer, for example, chose four kinds:

- internal data on shipments to retailers, prices, discounts, promotions, and various product characteristics
- data on consumer demand, which can be accessed through retailers' point-of-sale technology or provided by companies such as IRI and Nielsen
- macroeconomic information—including quarterly GDP, the Purchasing Managers' Index, the Consumer Price Index, and unemployment and inflation rates—that helps explain consumer behavior, seasonality, and trends
- external data on other factors that can indicate or affect demand, such as web searches, social media mentions of

products, average temperature, precipitation, holidays, and competitors' prices

Using such data and advanced analytics, firms can set up an automated five-step circular process that generates supply, financial, and trade plans for the next 50 to 80 weeks—the planning horizon for most companies. (See the exhibit "Five Steps to a Better, Cheaper Supply Chain.") Here's what that process looks like at the CPG manufacturer:

First, trade-planning information—about future promotions, discounts, and marketing investments—is combined with consumer, macroeconomic, and external data to generate a market demand forecast by SKU and retailer for each week of the entire horizon. From what we've observed, most CPG companies have never tried to predict demand at such a granular level.

Second, the demand forecast for each retailer is combined with historical data on the company's shipments to that retailer to generate a weekly forecast of the retailer's orders of each SKU for the horizon.

Third, the company aggregates all the order forecasts and converts them into a feasible supply plan. The plan considers available resources, including inventories of raw materials and finished goods; manufacturing capacity constraints; and market targets (say, for increased sales of a product category at a given retailer-region combination). It also aims to achieve certain performance goals. The CPG firm focused on minimizing total supply chain costs, but the chosen objective will vary from firm to firm. At some companies, for instance, it may be to maximize revenue or the amount of supplies produced.


The fourth step is to use the weekly SKU supply plan for all retailers to generate revenue and gross margin forecasts at the brand level for every month of the planning horizon.

The fifth step is to compare that financial forecast with the firm's business objectives. A gap between the two may trigger a change in the trade plan—for example, the addition of more-aggressive discounts or increased investments in marketing to stimulate sales.

When they were considering the adoption of this new process, the CPG firm's managers raised a number of questions—which are representative of the kinds of concerns most executives express about our approach. Let's examine them one by one.

Five Steps to a Better, Cheaper Supply Chain

Using advanced analytics, firms can set up an automated five-step circular process that generates supply, financial, and trade plans for the next 50 to 80 weeks.



STEP 1

The market demand forecast

Trade-planning information—about promotions, discounts, and marketing—is combined with consumer and macroeconomic data to forecast weekly market demand for each SKU and retailer.

STEP 2

The retailer order forecast

The demand forecast for each retailer is combined with data on past shipments to it to generate a weekly forecast of the retailer's orders of each SKU.

STEP 3

The supply plan

The order forecasts are combined with information on available resources (including inventories of raw materials and finished goods), manufacturing capacity, and market targets to create a supply plan.

STEP 4

The financial plan

The weekly supply plan is converted into monthly revenue and gross margin forecasts at the brand level.

STEP 5

The alignment of plans and goals

The financial plan is compared with the firm's business objectives. If gaps are identified, the firm may adjust the trade plan—by, for example, adding more discounts or increased investments in marketing—and start the cycle again.

What degree of forecast accuracy can the process achieve? Research has proved that variability in customer demand is significantly lower than variability in retail orders—a reality that underlies the well-known bullwhip effect in supply chains. This implies that predicting consumption should be easier than predicting retail orders, and indeed, the accuracy of the CPG firm's forecast for market demand is quite high. At any moment the demand forecasts at the SKU, week, and retailer level for five to eight weeks out have proved to be 85% accurate.

Combining the more exact consumption forecast with historical retail orders allowed the CPG company to improve its forecast of retailers' future orders. The accuracy of the weekly order forecasts has been 15 to 20 percentage points higher than that of the standard, consensus-based forecasts the company previously used. And more-accurate order, or shipment, forecasts clearly translate into a more effective supply plan, which reduces lost sales—therefore boosting revenue—and improves service levels and the customer experience.

Finally, because the inputs into it are more accurate, so is the financial plan. In multiple implementations of this approach at several CPG companies, the accuracy of the financial forecast made at the beginning of a given month for the next month rose to 95% to 97%.

Will we be able to understand what drives the behavioral and other changes the plans predict? This question is probably the most critical. Indeed, in our experience, virtually all executives are reluctant to blindly follow the recommendations of a black box developed by data scientists. They rightly want to be able to interpret and explain the output of the demand-forecasting process.


For instance, is an increase or decrease due to competitors' behavior, cannibalization across products, promotions

and discounts, or merely a special event or holiday? The good news is that the analytic technology today is mature enough to allow a single SKU weekly forecast to be decomposed into its basic components. This is done by explicitly modeling the data as a combination of key variables (competitors' behavior and so on) and estimating the contribution of each one to the forecast.

Executives also want to know the reasons why, say, the forecast generated last week is different from the one generated this week. This, too, is information that today's analytic technology can provide, by comparing the input data used to generate each of the forecasts.

Last, executives want to understand why forecasts and actual sales sometimes deviate. At the CPG firm the answer is that sales are affected by the way pricing, promotion, discounts, and inventory decisions are executed by retailers—a dimension that the manufacturer's planning team can't see. For instance, the forecast might be off when a retailer experiences operational challenges in moving inventory onto the shelf or in implementing promotions or discounts according to plan. Information about the retailer's inventory and prices paid by consumers at the cash register can reveal these problems, but in our experience most retailers don't provide it to their CPG suppliers. Thus, at the CPG firm any significant gap between the forecast and actual sales triggers an investigation of the reason for the difference.

How can we ensure that all the functions follow the new approach? The answer is to establish a forecast center of excellence that brings together people from the various functions, information technologists, and data scientists. Their role will be to agree on the data to be used and let the analytics generate the forecasts and the supply plan according to the five-step process.



How frequently should we run this process? Here, the answer depends on the market cycles of the various businesses and brands. For most businesses the demand forecast, retailer order forecast, and supply plan should be updated weekly or biweekly, while the financial forecast and the comparison with the firm's objectives should be done monthly. But there are clear exceptions. Some of the CPG manufacturer's products have short life cycles of only six or seven weeks. In such cases companies need to update the demand forecast, retailer order forecast, and supply plan twice a week. (The same is true for makers of fashion products, whose selling seasons last no more than 10 or 11 weeks.)


REDEFINING THE SUPPLY CHAIN STRATEGY

Traditional supply chain strategies have often focused on either operational efficiency or responsiveness. When operational efficiency is the priority, a firm strives to squeeze as much cost out of the supply chain as possible, and that goal drives supplier selection, manufacturing strategies, product

design and distribution, and logistics. Typically, production and distribution decisions are based on long-term forecasts, inventories of finished goods are located close to customer demand, and components are often sourced from low-cost countries.

The objective of a responsive strategy is to compete on time to market, satisfy demand quickly, and eliminate stock-outs. Manufacturing or product assembly is based on actual orders rather than forecasts; products may be customized; inventories of components are maximized but inventories of finished goods are minimized; and speed is prioritized over cost in decisions about sourcing and transportation.

Although seasoned operations and supply chain executives understand the difference between efficiency and responsiveness, many are nonetheless confused about when to apply each strategy. That's because different products have different characteristics, with some requiring a strategy focused on efficiency, some a strategy focused on responsiveness, and some a hybrid approach. Until recently, executives




didn't have the tools to segment products and decide which strategy was appropriate for a particular segment. But that has changed, thanks to digitization and analytics.

The CPG manufacturer began by exploring variations in sales data, focusing on products' sales volatility, volume, and profit margin, because each is directly related to risks associated with stock-outs, service levels, inventory, and transportation. The higher sales volatility is, the lower the forecast accuracy, and the riskier the product. That in turn translates into frequent stock-outs and lower service levels. Similarly, the higher a product's profit margin is, the higher the risk is, since missing an order will have a bigger impact on the bottom line. Volume, in contrast, is inversely proportional to risk—that is, the higher the volume, the lower the impact of missing an order, and the lower the risk. These relationships are consistent with those we've seen at other CPG and retail companies, though sometimes other companies focus on price or product cost rather than product margin, depending on which one is more stable and as a result easier to apply.

The analysis revealed that the CPG company had four product segments, although other companies may have more segments given their products' characteristics. Each segment required a different supply chain strategy. (See the exhibit "The CPG Firm's Segmented Strategy for Supplying Retailers.") The first segment comprises products characterized by high volatility. Because their stock-out, service-level, and inventory risks are high, they require a responsive supply chain strategy. Finished-goods inventories for them should be located in central distribution centers. Each center will be responsible for many retail outlets, which allows a company to aggregate demand, improve forecast accuracy, and reduce the inventories needed to supply the retailers collectively while maintaining high service levels. Because fast delivery is critical, these products are often shipped through cross-dock regional facilities—at which items from incoming large trucks are reloaded onto outbound smaller trucks with no storage in between.

The second segment comprises products with high volume and low volatility, which require an efficiency strategy.



SUPPLY CHAIN

In their case, forecasts are reliable, and managing transportation costs is important. Because of this, the products are stored in regional warehouses close to customers, and inventory is replenished on a fixed schedule. That allows a company to fully load trucks taking products from manufacturing facilities to regional warehouses, which keeps transportation expenses down.

The remaining two segments are both characterized by conflicting drivers: low demand volatility (which suggests that an efficiency strategy would be best) and low product volume (which alone would call for a responsive strategy). What distinguishes these two segments are product margins.

Let's look at the high-margin ones first. Because these products are riskier, many of them are stored at both centralized locations and regional warehouses and are replenished on the basis of actual store sales. That strategy allows a firm to strike a balance between efficiency and responsiveness, though it leans toward responsiveness.

Low-volatility, low-volume, low-margin products, in contrast, call for a hybrid strategy that leans toward efficiency. Indeed, because the risks and cost of holding inventory are low while demand is predictable, a firm can ship these products on fully loaded trucks to regional warehouses close to its customers, supply them from those locations, and minimize transportation costs.

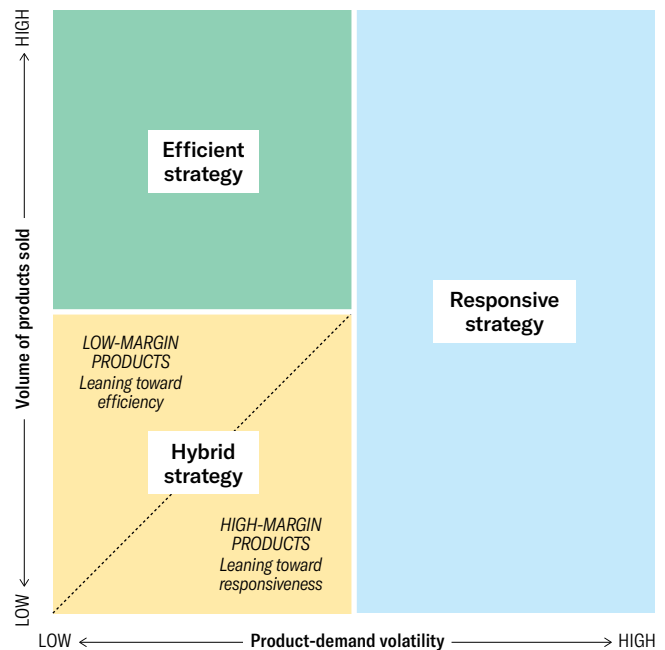
Once a company has done the segmentation, it needs to develop detailed sourcing, manufacturing, and logistics strategies. One objective should be to identify synergies across the segments that will allow the firm to benefit from economies of scale. They can be achieved by leveraging volume across segments to reduce procurement costs; sharing capacity and infrastructure in manufacturing and logistics; and consolidating demand and supply information for better planning and execution. We'll now look in more detail at that last activity.

BALANCING SUPPLY AND DEMAND

An important supply chain management process that has been applied since the mid-1980s is sales and operations planning (S&OP). It continually balances supply and demand, and historically it has called for managers launching new products and leaders from manufacturing and

The CPG Firm's Segmented Strategy for Supplying Retailers

An analysis of data on sales volatility, volume, and profit margins revealed that the consumer products manufacturer had four segments of products. Each required a different supply chain strategy: a responsive one prioritizing speed; an efficient one prioritizing low costs; or a hybrid one leaning toward either efficiency or responsiveness.



distribution to come together and agree to a single plan. Typically, it involves analysis at the business unit level or the product family level, not the individual product level.

Traditionally, S&OP is simply an extension of the consensus forecast, and because of that it suffers from similar limitations: It doesn't start with a unified view of demand, doesn't create a plan at the SKU level, doesn't distinguish between supply chain segments, and is driven mostly by common sense, experience, and intuition, not data and analytics. Because it's a manual process, it generally takes a month.



Companies need key performance predictors: metrics that indicate what the state of the supply chain will be in the next three to six weeks. These are central to smart execution.

A better approach to S&OP replaces the manual process with an automated one that can be performed weekly, and ensures that the engineering, finance, sales, supply chain, manufacturing, sourcing, and trade functions are all working to achieve the same business goals. The new process begins when an analytics-driven optimization system generates the SKU-by-SKU supply plan we described earlier. This plan will inform everything from master production schedules to materials planning to logistics, including inventory and transportation decisions.

While not every company or business unit needs to produce a plan weekly, such frequency is critical for products whose demand is highly volatile and whose marketing and promotion strategies often change.

The new S&OP process also calls for monitoring activities. Firms should collect information throughout the supply chain about key performance indicators (KPIs) such as supply lead times, raw-material and finished-goods inventories, and service levels, looking for any problem or deviation that could undermine the sales and operations plan. Firms can then work to address those issues and, if they turn out to be significant, adjust the plan itself.

Companies also need to keep an eye on data and events that portend what may happen in the near future. For example, while inventory and service levels may suggest that everything is going smoothly, shipment-tracking data may indicate that lead times are likely to increase and that as a result service levels could go down in the next few weeks, signaling a need to build inventories or expedite shipments. Similarly, if a disaster causes the shutdown of a supplier's manufacturing facility in Asia, it could affect available supply down the road—perhaps forcing a firm's manufacturing and assembly plants on the U.S. West Coast to lower or stop operations in five weeks. But traditional KPIs alone might not provide any warning.

For this reason, companies need key performance predictors (KPPs): metrics that indicate what the state of the supply chain will be in the next three to six weeks. KPPs are central to what we call *smart execution*, a new business process that complements smart S&OP. While S&OP focuses on the next 50 to 80 weeks, smart execution homes in on the short term (no more than six weeks) and tries to identify and quickly respond to disruptions and deviations from the plan.

Smart execution involves three automated capabilities: (1) the real-time capture of internal and external data that reveals potential deviations from the plan, supply disruptions, or changes in demand; (2) artificial intelligence that identifies the potential impact of those developments on supply chain performance; and (3) analytics-driven optimization that determines the best response, considering various trade-offs and objectives.

Here's an illustration. By gathering financial information on suppliers that are public companies and internal data on supplier performance (for instance, on lead times, service levels, or product quality), firms may be able to identify distressed suppliers. An AI system can then project the likelihood and impact of a supplier default on future commitments to on-time delivery and product quality. Finally, the automated optimization system can identify an alternative supplier for sourcing the material.

FOR MOST OF its history, the CPG manufacturing company had used a one-size-fits-all strategy. Its forecasts were achieved by consensus, S&OP was a monthlong process, the supply chain strategy didn't distinguish between different products, and deviations from the plan and supply disruptions were managed ad hoc. The company had excelled at operational efficiency by embracing continuous improvement in its production, packaging, distribution, and order fulfillment processes, but it hadn't fundamentally changed any of them. Our approach to supply chain digitization allowed the firm to transform the organization in less than half the time and at less than a quarter of the expense that such efforts take most corporations.

Other firms can do the same. A comprehensive, automated approach can allow them to redefine their supply chain strategies and respond quickly to deviations from the plan. And because it's driven by AI, it will free up executives to devote more time to value-added activities, such as identifying the best opportunities for growing the business. ☺

HBR Reprint R2105K



DAVID SIMCHI-LEVI is a professor of engineering systems at Massachusetts Institute of Technology and the head of the MIT Data Science Lab. **KRIS TIMMERMANS** is a senior managing director at Accenture and the head of its supply chain and operations practice.