



# Which Products Should You Stock?

A new technique to help retailers improve assortment planning by Marshall Fisher and Ramnath Vaidyanathan

**G**etting product assortment right isn't easy, yet it's absolutely critical to retail success. Unlike inventory management and pricing, where retailers have lots of data and analytical tools to guide decision making, assortment optimization is still much more art than science. And making the wrong call can be disastrous. Consider these examples:

- Following a survey in which customers said they'd like less cluttered stores, Walmart introduced Project Impact, in **2008**, removing **15%** of the SKUs it carried. Sales declined significantly, and it was forced to roll back most of the changes.

ILLUSTRATION: JENNY BOWERS



- Super Fresh, owned by the grocery retailer A&P, stopped carrying many of its low-selling dry grocery items to allow for an expansion of fresh offerings. But the eliminated products turned out to be essential to customers; when they couldn't find them, they took their business elsewhere, and the retailer entered bankruptcy.

- A retailer of home goods used demographic data to localize its assortments to better cater to customers' tastes. It started with fashion bedding and was thrilled to see an 18% revenue lift. But when it applied the data to the fashion bath category, revenues didn't budge. Discouraged, the retailer abandoned the effort.

- When the new CEO of a tire retailer shifted its assortment from low-priced tires to more-expensive ones, he learned the hard way that price mattered to his customers. The CEO was replaced after two years, and his successor restored most of the products that had been eliminated.

Like so many assortment-strategy shifts, these moves were largely acts of faith. It's easy to spot the dogs in your assortment, of course—sales data will tell you that—but it's far from obvious what slow sellers should be replaced with. And there is always the nagging concern that a slow seller you delete might be an important product to some of your best customers, prompting them to defect to competitors. As all retailers know, picking the best assortment is a balancing act; any one change can have ripple effects.

Plenty of software tools claim to support assortment planning, by helping retailers decide which combination of products will maximize sales. But with very few exceptions, they lack the ability to forecast demand for new products or to estimate how much demand would transfer to other products if a slow seller were dropped. The tools do little more than facilitate a manual planning process that relies on the judgment of managers for key inputs. They do nothing to reduce the risk inherent in every product-assortment decision.

To address this deficiency, we've developed a technique that makes assortment planning vastly more scientific. It is rooted in our observation that most of the time customers don't buy products; they buy a bundle of attributes. Think about the last time you bought a TV. Did you say, "I want TV X"? Or did you think about screen size, resolution, price, LCD versus plasma, and brand? Our approach uses sales of existing products to estimate the demand for their various attributes and then uses those es-

timates to forecast the demand for potential new products. Armed with these data, retailers can test their hunches more scientifically.

Our approach is especially useful for retailers in the hard-goods and grocery segments; it's less applicable in the fashion-sensitive apparel segment, where products change fast. Grocery retailers currently use the abundance of available market data to identify potential additions—SKUs they don't carry that sell well at other retailers. But research we conducted shows that our attribute-based approach has a lower margin of error.

It also helps retailers gain insight into the following questions:

- Can we improve our assortment by replacing slow-selling products with new ones? What is the likely demand for the potential new items?

- If customers don't find their ideal product, what is the likelihood that they will substitute another?

- How will sales change if we increase or decrease the number of products we carry?

- Does localization—customizing assortments by store or store cluster—make sense? If so, for which categories? If we decide to create clusters of stores with distinct assortments, how many should we create, and what criterion should we use in creating them?

By focusing on the attributes of products, retailers can maximize the number of customers who say either "That's exactly what I want" or "This product may not be what I'd ideally like, but it's close enough, and I'll buy it." Let's now look at the process, using two examples in auto parts retailing: the tire department of one chain (a research project we conducted) and the appearance-chemicals division of another (a consulting job). While the process is described here step-by-step, it is in fact multidimensional and highly iterative; much of the analysis is handled by a computer model, which produces the final recommendations.

## **Understand Which Attributes Matter Most to Customers**

Using our method still requires some judgment about which attributes are important to consumers and how those preferences might influence their purchase decisions if they don't find their first choice. The steps below can help retailers tackle those questions.

**Identify which attributes are important to customers.** Most retailers already think about their products in terms of attributes and can readily iden-

## Idea in Brief

When figuring out which existing products to drop and which new ones to add in their stores, retailers still largely rely on judgment. The result is often disappointing and sometimes disastrous.

A new technique makes assortment planning much more scientific. It uses sales of existing products to estimate the demand for their various attributes—for example, TVs' screen size, resolution, price, and brand. It then uses those estimates to forecast the demand for existing products and for others that could be added to the assortment.

This technique helps retailers do a better job of replacing slow sellers with new offerings, understanding whether customers are likely to settle for another choice if they don't find their ideal product, and customizing assortments for individual stores or clusters of stores.

Because retailers have inherent biases of what's their assortment, they might not be choosing the best assortment. This is where the science comes in.

tify those that matter in their category. They might include price, brand, size, flavor, and color.

When we began our research project on tires, the retailer's category manager told us that the important attributes for tires were brand, the mileage warranty, and size. The retailer offered several nationally advertised brands that the manager believed customers regarded as interchangeable. We grouped those together as National Brands. The retailer also offered three house brands of varying quality and price, which we'll call House 1 (the premium brand), House 2 (mid-level), and House 3 (low-end). A number of mileage warranties were offered, but the retailer believed that consumers considered many to be equivalent. Therefore, we grouped the mileage warranties into three levels: Low (15,000 to 40,000 miles), Medium (40,001 to 60,000 miles), and High (greater than 60,000 miles).

The four brands and three mileage-warranty levels implied 12 brand-warranty combinations that the retailer theoretically could offer, but some made little sense, such as a high-mileage warranty on a low-priced brand. Only six combinations were actually offered (in decreasing order of quality): National High, National Medium, House 1 High, House 2 High, House 2 Medium, and House 3 Low.

The third key attribute for tires, size, includes type (for instance, radial), and whether it is for a passenger car or some other kind of vehicle. Tires come in 64 sizes, which means that there were 384 possible SKUs the retailer could have carried (64 sizes x 6 brand-warranty combinations). But it carried only 105 in most stores. The count varied across the chain, mostly according to the size of the store. The assortments varied as well, but most stores carried the most popular SKUs.

**Account for what customers will do if you don't offer their preferred product.** Customers' willingness to buy another product if they don't find

## It's easy to spot the dogs in your assortment, but it's far from obvious what slow sellers should be replaced with.

their first choice is a crucial input when a retailer considers dropping items. Their willingness depends greatly on the attribute. Customers probably won't substitute one dress size for another, but they might buy the blue one if red is not in the assortment. Similarly, people are not going to buy a 14-inch tire for a 15-inch wheel, but they might substitute one brand and mileage-warranty combination for another. So in building an assortment, retailers need to account for the fact that if customers don't find their ideal item, some of them will buy the next-best option and some won't. In our tire example, we were interested in the percentage of customers that would shift up by one quality level if their first choice were unavailable and the percentage that would shift down.

### Analyze Current and Potential Sales By Attribute

Now we'll figure out how well items you don't currently carry would sell and how adding them to your product mix would affect overall sales. This is where the science comes in.

**Assemble sales data for a recent period.** Start with what you know: unit sales of the SKUs you currently carry and each brand-warranty combination's share of your total sales. This is the foundation of the model. We typically look at six months' to a year's worth of recent data.

In the tire project, we assembled sales data by SKU for every store over a recent six-month period.

**Figure 1: Analyze Sales Data**

Customers don't buy products; they buy attributes. A tire retailer we worked with analyzed its sales data according to what mattered most to its customers: size, brand, and mileage warranty.

SIZES	BRAND-MILEAGE WARRANTY COMBINATION						TOTALS
	National HIGH	National MEDIUM	House 1 HIGH	House 2 HIGH	House 2 MEDIUM	House 3 LOW	Total units sold
A	100			29	28	190	347
B	282	21		30	203		536
C				11	12	86	109
D				53	50	284	387
E	72	64	20	172	570		898
F	59		97	285	763		1,204
G	10		16	14	76		116
H		7	33	157	377		574
I		10		183	524		717
J		39		225	568		832
K			8	10	73		91
L			8	47	223		278
M				43	298		341
N				72	221		293
O	8					200	208
Total units sold	531	141	182	1,331	3,986	760	6,931
Share of total sales	7.7%	2.0%	2.6%	19.2%	57.5%	11.0%	

Figure 1, showing one store's data for 15 of the 64 tire sizes, represents the raw data for our analysis. (The data have been changed to protect proprietary information.)

**Forecast demand for all potential SKUs.** The fact that some SKUs had only single-digit sales suggested that replacing them could increase revenue,

but the challenge was to figure out which new SKUs would sell better. The first step is to use sales data to forecast total demand for each tire size if all brand-warranty combinations were offered.

To illustrate, let's look at size F. (See Figure 2.) Notice that the retailer currently carries size F in four of the six brand-warranty combinations. We start by

**Figure 2: Estimate Total Demand**

The retailer next determined what total demand would look like if all brand-warranty combinations were offered in each size.

**AS AN EXAMPLE, LOOK AT SIZE F:**

$1,204 \div 87\% = 1,384$   
total sales for size F      share captured      total demand for size F

SIZES	BRAND-MILEAGE WARRANTY COMBINATION						TOTALS		
	National HIGH	National MEDIUM	House 1 HIGH	House 2 HIGH	House 2 MEDIUM	House 3 LOW	Total units sold	Share of demand captured	Total demand
A	100			29	28	190	347	95.3%	364
B	282	21		30	203		536	86.4	620
C				11	12	86	109	87.7	124
D				53	50	284	387	87.7	441
E	72	64	20	172	570		898	89.0	1,009
F	59		97	285	763		1,204	87.0	1,384
G	10		16	14	76		116	87.0	133
H		7	33	157	377		574	81.4	705
I		10		183	524		717	78.8	910
J		39		225	568		832	78.8	1,056
K			8	10	73		91	79.4	115
L			8	47	223		278	79.4	350
M				43	298		341	76.7	444
N				72	221		293	76.7	382
O	8					200	208	18.6	1,119
Total units sold	531	141	182	1,331	3,986	760	6,931		
Share of total sales	7.7%	2.0%	2.6%	19.2%	57.5%	11.0%			

Figure 3: Refine the Forecast

SIZES	BRAND-MILEAGE WARRANTY COMBINATION						TOTALS		
	National HIGH	National MEDIUM	House 1 HIGH	House 2 HIGH	House 2 MEDIUM	House 3 LOW	Total units sold	Share of demand captured	Total demand
A	100			29	28	190	347	97.3%	357
B	282	21		30	203		536	28.8	1,862
C				11	12	86	109	94.9	115
D				53	50	284	387	94.9	408
E	72	64	20	172	570		898	30.3	2,960
F	59		97	285	763		1,204	29.2	4,123
G	10		16	14	76		116	29.2	397
H		7	33	157	377		574	27.9	2,054
I		10		183	524		717	26.4	2,717
J		39		225	568		832	26.4	3,153
K			8	10	73		91	26.8	339
L			8	47	223		278	26.8	1,037
M				43	298		341	25.3	1,350
N				72	221		293	25.3	1,160
O	8					200	208	72.0	289
Total units sold	531	141	182	1,331	3,986	760	6,931		
Share of total sales	7.7%	2.0%	2.6%	19.2%	57.5%	11.0%			
Best-fit demand shares	2.4%	1.1%	1.5%	6.7%	18.6%	69.6%			

Because estimates have inherent margins of error, the retailer refined the forecasts to produce "best fit" demand shares.

USING BEST-FIT SHARES, DEMAND FOR SIZE F ALMOST TRIPLES:

$$1,204 \text{ total demand for size F} + 29.2\% \text{ share captured} = 4,123 \text{ total demand for size F}$$

ESTIMATES MAY NOT BE PERFECTLY ACCURATE:

$$4,123 \text{ total demand for size F} \times 18.6\% \text{ share captured for House 2 Medium size F} = 767 \text{ demand estimate for House 2 Medium size F}$$

$$4,123 \text{ total demand for size F} \times 6.7\% \text{ share captured for House 2 High size F} = 276 \text{ demand estimate for House 2 High size F}$$

adding up the shares of total sales enjoyed by each of the combinations offered in size F (7.7% + 2.6% + 19.2% + 57.5%). That tells us the share of total demand for size F that the retailer is currently capturing (87%). In other words, the retailer is theoretically forfeiting the shares of sales associated with the two combinations it does not offer: National Medium, 2%, and House 3 Low, 1%.

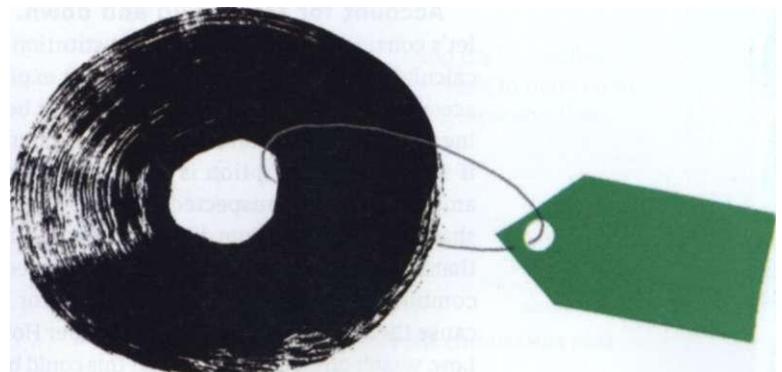
To calculate total demand for size F, we simply divide total sales for size F by the share of demand captured:  $1,204 \div 87\% = 1,384$ . Once we know the total demand for size F, we can estimate demand for any SKU in that size, by multiplying the total demand for the size by the share of sales enjoyed by the brand-warranty combination. For example, House 3 Low has an overall share of 11%; applying that percentage to 1,384 gives us a forecast of 152 units for House 3 Low in size F.

**Refine the forecast.** This calculation takes us only partway to an accurate forecast, as you'll see if you forecast the sales of an item you actually carry. We determined above that if all combinations for size F were offered, the total sales would be 1,384 units. Thus, the demand for House 2 Medium, the retailer's top seller, is estimated at 796 ( $1,384 \times 57.5\%$ ). Actual sales, however, were quite a bit lower: 763 units.

One reason for the discrepancy is that sales shares are influenced by the assortment offered. House 2 High and Medium were offered in almost every size, which raised their total sales shares relative

to those of brand-warranty combinations offered in fewer sizes.

To correct for such discrepancies, we need to tweak the brand-warranty sales shares to minimize the average difference—what statisticians call the mean absolute deviation—between estimates and actual sales. This highly iterative process is done using an optimization tool like Excel Solver. Essentially, the tool plugs trial values for the brand-warranty share numbers into the demand estimate calculations for all current SKUs and sees how close the resultant forecasts are to actual sales. Then it adjusts the share values to make forecasts closer, and repeats until it arrives at the share values that minimize the sum of all discrepancies over the SKUs offered. It's exactly the way you get a prescription for eyeglasses: Start with a trial lens, try a different lens-



**Figure 4: Estimate How Well New SKUs Would Sell**

		BRAND-MILEAGE WARRANTY COMBINATION					
		National HIGH	National MEDIUM	House 1 HIGH	House 2 HIGH	House 2 MEDIUM	House 3 LOW
SIZES	A		4	6			
	B			29			1,297
	C	3	1	2			
	D	10	5	6			
	E						2,060
	F		46				2,862
	G		4				276
	H	49					1,430
	I	65		43			1,892
	J	75		50			2,196
	K	8	4				236
	L	25	12				721
	M	32	15	21			940
	N	28	13	18			808
	O		3	5	19	54	

Using the best-fit shares, the retailer estimated demand for all SKUs it could add to the product mix.

better or worse?—adjust, and repeat until there's no further improvement.

This process resulted in "best-fit" demand shares for the six brand-warranty combinations: **2.4%**, **1.1%**, **1.5%**, **6.7%**, **18.6%**, and **69.6%**. (See Figure 3.) Compare the demand shares with actual sales shares, and a dramatically different picture of optimal assortment begins to emerge.

Note that the forecasts, while close to actual sales, are not perfectly accurate. Two factors contribute to forecast errors: First, there are random fluctuations in sales. And second, our assumption that demand for a SKU equals the demand for the size multiplied by the brand-warranty share is imperfect, because the shares of brand-warranty combinations can differ by size. (For example, demand was higher for low-end tires in sizes that fit older, less expensive cars than in other sizes.)

Having determined the best-fit share values, demand can be estimated for all potential SKUs. (See Figure 4.)

**Account for trading up and down.** Now let's consider another wrinkle: substitution. The calculations we described above do not explicitly account for the fact that customers might be willing to buy a different brand-warranty combination if their preferred option is not offered. For example, the retailer suspected that the **57.5%** sales share of House 2 Medium did not necessarily mean that more than half of its customers preferred this combination; they might have settled for it because their preferred option, the cheaper House 3 Low, wasn't offered. One clue that this could be the case was that when House 3 Low was offered in a

given size, it outsold House 2 Medium by about six to one.

Making matters even more complicated, the degree to which customers trade up and down may not be the same for all quality levels. If you think that's the case for SKUs that are especially important to your business, you need to account for this in your calculations. For the tire project, we assumed that the fractions of customers who would trade up or down were equal for all brand-warranty combinations with the exception of customers shifting from House 3 Low to House 2 Medium. (Those two brand-warranty combinations accounted for more than two-thirds of sales.)

So our model now requires nine parameters: the six brand-warranty shares and three substitution parameters, which include the fraction of customers who trade up one quality level, who trade down one quality level, and who shift from House 3 Low to House 2 Medium. As before, we use a tool such as Excel Solver that plugs in trial values for the shares and the fractions, calculates demand estimates and sees how close the resultant forecasts are to actual sales. It adjusts the shares and fractions to make forecasts closer, and repeats until there is no further improvement. The final results: **35%** of customers who couldn't find House 3 Low in the assortment in their size would trade up and buy House 2 Medium. For other quality levels, **2%** would trade up and **1%** would trade down if they couldn't find what they were looking for.

Once you know the fractions of customers trading up or down, you can account for substitution in your demand estimates. Consider House 2 Medium

in size F. Because House 3 Low is not offered in this size, take the demand estimate for House 2 Medium and add to it that of House 3 Low multiplied by the fraction of customers who would trade up. For size A, in contrast, both House 2 Medium and House 3 Low are in the assortment, so no value for substitution demand is figured into House 2 Medium's estimate.

**Look for self-fulfilling prophecies.** Now consider this familiar scenario: A retailer thinks its customers don't want to buy a certain product type (or the retailer doesn't want to carry it). So the company offers a limited amount and thus doesn't sell much of it, seemingly confirming the original assumption that customers don't want it. But in the end, the assortment reflects the products the retailer wants to carry rather than those its customers want to buy—a risky proposition. One benefit of our technique is that it allows retailers to spot such situations.

For instance, in comparing estimated demand and actual sales at the auto retailer, we found one surprising result: House 3 Low had an 11% sales share, but our estimates pegged demand at a whopping 69.6% of total sales. The low sales share occurred because the retailer offered this very-low-priced tire in only a few sizes and therefore didn't sell many. But as the data show, when customers had a choice between House 3 Low and House 2 Medium, they strongly preferred the former. This pattern persisted

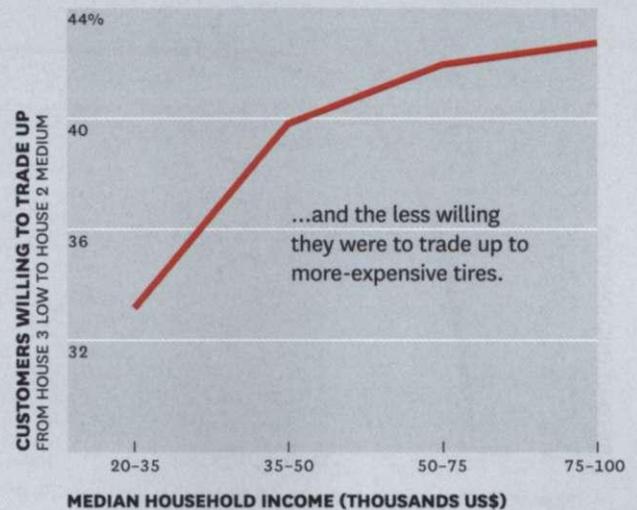
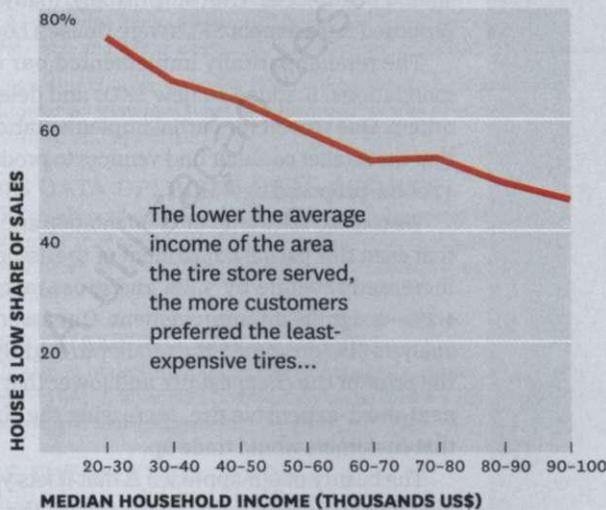
## This is not as simple as identifying the top 100 revenue-generating SKUs and calling it an assortment.

chainwide. There were nine sizes in which House 3 Low and House 2 Medium were both offered, and in every case, House 3 Low outsold House 2 Medium in total chain sales—by more than seven to one.

The retailer offered a limited selection of the cheapest tire because its managers thought they could trade customers up to the somewhat-higher-priced House 2 Medium. But they were successful in upselling only 35% of the time. Indeed, our model shows that by ignoring the estimated 69.6% share of House 3 Low, the company was losing 45% of its potential sales (the 65% of the 69.6% of customers who want House 3 Low and don't trade up).

To gain more insight into this finding, we tabulated average income in the area served by each store and used it to create Figure 5, which shows that the share of the cheapest House 3 Low and unwillingness to trade up were inversely correlated with income. In other words, the lower the average income

Figure 5: How Income Affects Demand and Willingness to Trade Up



# Our approach is superior to the conventional one, which requires retailers to guess how attributes influence demand.

of the area that a store serves, the more its customers preferred the least-expensive tires and the less willing they were to trade up to more-expensive ones.

## Optimize the Assortment

We will now describe how to use our model to decide which existing and new SKUs would constitute an optimal assortment.

**1. Decide whether to maximize revenues or profits.** The most natural profit measure in a retail context is total gross margin—typically, revenues minus cost of goods sold. Business schools and economists preach profit maximization, but retailers also care about revenues, in part because Wall Street watches that metric closely. In both the tire example and in the appearance-chemicals case, which we'll look at later, the goal was to maximize revenue.

**2. Decide on pricing for potential SKUs.** To optimize an assortment, you need to know how

much revenue (or margin) each SKU would generate. Prices are a key input in this calculation. The prices of existing SKUs are known, of course. If prices for the new SKUs are unavailable, come up with estimates by comparing the attributes of current SKUs with those of potential new ones.

In the tire example, we observed that the prices of a given size decreased consistently from the highest-priced brand-warranty combination (National High) to the lowest-priced (House 3 Low). We applied those decreases to estimate the prices of SKUs not carried.

**3. Decide on the final assortments.** Next, calculate the potential revenue of each SKU by multiplying its forecasted unit sales by its retail price.

Now you have the data you need to begin building your assortment. Start with the SKU that would generate the greatest revenue or profit for the store or the chain. Then add the SKU that would yield the second-greatest increase in revenue. Continue to add SKUs until you hit the maximum number of SKUs you want to carry, say, **100** out of a possible universe of **400**.

Make no mistake: This is not as simple as identifying the top **100** revenue-generating SKUs and calling it an assortment. Because of demand substitution, each time you add a SKU to the assortment, you have to adjust your figures to account for how that new SKU affects demand for the ones you've already added. The process, obviously, is highly iterative.

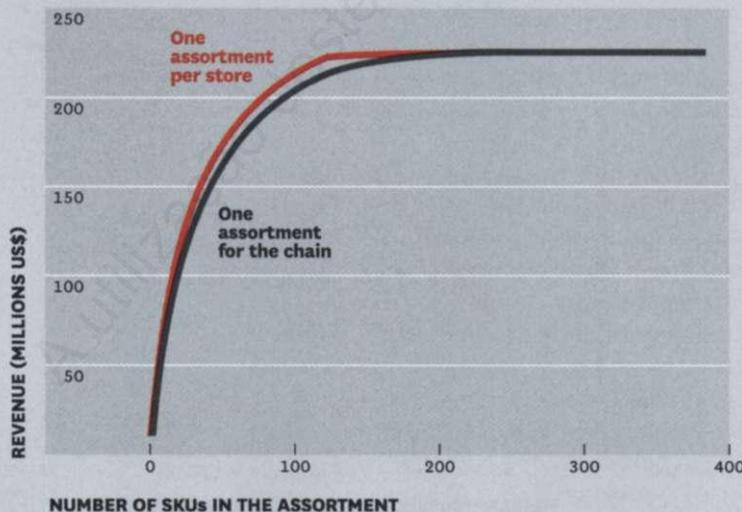
When we applied this process to create an optimal chainwide assortment of tire SKUs for the auto parts retailer, we found that **47** of the **105** SKUs should be replaced. (Not surprisingly, many of the proposed replacement SKUs were House 3 Low.)

The retailer partially implemented our recommendations: It added **10** new SKUs and deleted **10** others. One reason for partial implementation was that the retailer couldn't find vendors to produce all **47** of the proposed new SKUs.

We tracked sales after implementation and found that even this partial adjustment of the assortment increased revenue by **5.8%** and gross margin by **4.2%**—a significant improvement. Our assortment analysis also prompted the retailer to slightly raise the price of the cheapest tire and lower that of the next-most-expensive tire, increasing the chances that customers would trade up.

The beauty of our approach is that it lets you see how revenue varies with the breadth of the assortment. Figure 6 shows how the tire retailer's revenues

Figure 6: How Many SKUs Should You Carry?



are influenced by the number of SKUs in the assortment. The upper line shows the revenues if each store had its optimal assortment, while the lower line gives the revenues for a single optimal assortment for the chain. Graphs like this can be used to change the shelf space allotted to categories in order to increase sales. They also can help retailers avoid the mistakes made by Walmart and Super Fresh, which reduced their assortment breadth with disastrous results.

## Localizing the Assortment

Creating localized assortments is complicated. A retailer needs to understand how demand differs between stores and then create assortments that cater to store-specific tastes. Most retailers find it much too complicated to carry a unique assortment for each store; instead, they create clusters of stores that use the same assortment. In such cases, they need to decide how many clusters to create, on what basis clusters should be formed (for instance, income or weather), and what assortment to use in each cluster.

Our attribute-based technique is an excellent way to answer these questions, as we can see in a study of the appearance-chemicals category that we performed for an auto parts retailer that had hundreds of stores. Appearance chemicals comprises an array of liquids and pastes used to wash, wax, shine, polish, and protect cars. In the study, we identified six attributes of products in the category: the car surface to be treated, what is to be done to it, the application mode, the package size, the brand, and the quality level (good, better, or best). The retailer was eager to understand how demand patterns differed across stores and then use that information to localize its assortments. It considered at most five clusters, believing that it would not be operationally feasible to have more than five assortments.

We applied the method described above to estimate demand shares for the various attribute levels, used those estimates to forecast demand for potential new SKUs, and generated a revenue-maximizing assortment for each store.

Using the assortments we had generated, we then created store clusters. We began the process by assuming that each store represented a cluster. We then identified the two stores that would suffer the smallest reduction in revenues if they were forced to share the same assortment and combined them to create a two-store cluster. We repeated the process, identifying the next two stores that could best share an assortment or by adding a third store to the two we had

already combined—whichever would result in the smallest reduction in revenue. We kept going, each time reducing the number of store clusters by one, until we were left with a single cluster of all stores. This gave us revenue numbers for all levels of localization, ranging from a single assortment for the entire chain up to an individual assortment for each store.

Figure 7 shows the revenues for five localization options, ranging from one to five store clusters, adjusted to make the revenues of a single assortment **100**. As you can see, the returns from adding store clusters diminish, which led the retailer to implement the two-cluster solution.

The data also revealed that one of the two clusters, accounting for about a third of the chain's stores, sold much higher levels of tire-related products. That cluster had a distinctive ethnicity, which the retailer called urban/bilingual. So these stores' assortments featured more tire-related products, and the retailer created signage that called attention to them.

After tracking sales for six months, we found that the chain's revenue in the appearance-chemicals category was up 3.5%. This gain resulted both from localization and from improving the base assortment. Moreover, the new assortment and new signage helped the retailer in another way: It had been losing sales to the competition in the urban/bilingual stores, but after the changes, it began to show increases in comparable-store sales for the category. We believe this demand-based approach to clustering is superior to the conventional approach, which requires retailers to guess how store attributes influence demand.

**ANALYTICS HAVE** not been heavily applied to assortment planning, especially at the operational level of deciding which SKUs to carry. Our method uses analytics to glean insight into the product attributes your customers prefer at each store and then create localized assortments on the basis of those insights. Assortment planning can add significantly to same-store sales; but done wrong, it can cripple a retailer for years. Our method can help retailers do it well. 

**Figure 7: How Many Clusters Do You Need?**

Number of store clusters	Revenue
1	100.0
2	102.1
3	103.5
4	104.6
5	105.7

REVENUE NUMBERS HAVE BEEN INDEXED

The returns from adding store clusters diminishes, which led the retailer to implement a two-cluster solution.

 **Marshall Fisher** is the UPS Professor of Operations and Information Management at the University of Pennsylvania's Wharton School and a coauthor, with Ananth Raman, of *The New Science of Retailing: How Analytics Are Transforming the Supply Chain and Improving Performance* (Harvard Business Press, 2010). **Ramnath Vaidyanathan** is an assistant professor of operations management at McGill University's Desautels Faculty of Management.