A new algorithm helps retailers make better inventory decisions

13 December 2017

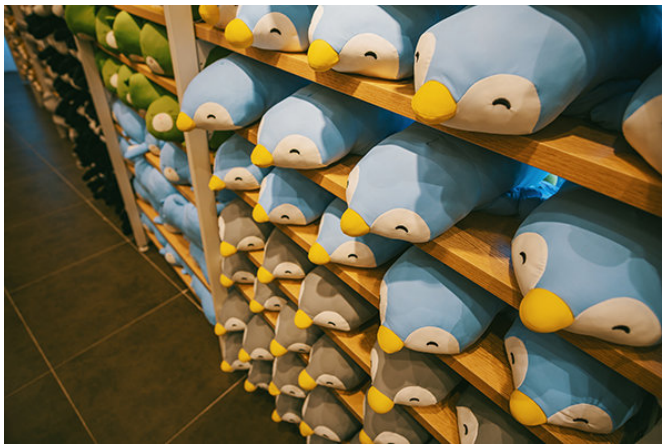

Offering multiple, similar items can complicate inventory decisions. Credit: MIT Sloan School of Management

Stocking too much of a product, or not enough, costs retailers hundreds of billions of dollars annually. If they stock too little and run out, the customer will likely take their businesses elsewhere, costing the retailer money. If they stock too much, though, the retailer ends up with excess inventory.

A forthcoming paper in *Operations Research*, co-authored by MIT Sloan visiting professor Amr Farahat, PhD '04, and Joonkyum Lee, an assistant professor at Sogang Business School in South Korea, presents a new way to tackle this problem. Their data-driven approach can help retailers make more informed decisions and increase profits.

While having retailers simply restock an item when they run out of it may seem reasonable, that approach doesn't work for many. "Shoppers are not going to wait until the retailer replenishes to make a purchase," Farahat said.

Instead, retailers have to predict in advance how much they are going to sell of specific items. That's not easy—especially for products whose replenishment lead-times are long compared to the duration of their peak selling season.

The effect of substitution

To help, Farahat and Lee have developed what they call the approximate similarity transformation.

"This algorithm recognizes that there is a relationship between how much retailers stock and their profits. This relationship is complicated, so we replaced it with a simpler one that provides an upper bound on sales, but it is a tight upper bound. Dealing with that simpler, yet approximate, sales function leads eventually to better decisions," Farahat said.

Retailers typically base their stocking decisions on past sales while taking into account the time of year, how the economy is doing, what is in fashion, and what new products have come out that are expected to sell, among other things.

According to Farahat, if a retailer stocks a single item, determining the optimal amount of stock to carry is simple. When a vendor sells many items, this becomes more complicated due to customers substituting one item for another. "If I want to buy a striped blue shirt for a gift, and I don't find the size I need in stock at Macy's, instead of deciding to buy the gift at a different store, I may look at a solid blue shirt, a striped purple shirt, or a different brand. Complex substitution effects are going on—it is the nature of consumer choice," Farahat said.

Since consumers could make any number of decisions based on the inventory that is available, it is practically impossible to determine optimal inventory levels. "Mathematically, this is among the most challenging problems in computer science," Farahat said.

A data-driven approach

Determining precisely optimal inventory needs is unattainable, but the approximate similarity transformation produces recommendations based on "demonstrably good approximations." Farahat and Lee's research indicates that by following these recommendations, some retailers can expect profit increases of 2-3 percent.

It achieves this by using the data that retailers have already collected on their customers, such as traffic expectations based on the season and how their consumers make choices. It then approximates a sales forecast that can help with inventory decisions.

"We are trying to make recommendations for retailers that they can use as a starting point for their final decisions. For us to be able to do that we need to build on the predictive analytics capabilities that a lot of retailers are already developing," Farahat said.

As companies hone their data gathering capabilities, the quality of that data will improve—and Farahat and Lee's algorithm will become more useful. "As these models and forecasts become more accurate, this prescriptive decision-making piece becomes more relevant," Farahat said.

Sharing the research

Farahat and Lee have run thousands of numerical experiments based on retailer experiences to test their algorithm. Those tests have all indicated that the algorithm works as well or better than previous methods to help retailers plan their inventory since it provides stricter upper bounds, or more precise profit expectations, in more than 99 percent of the tests.

The researchers would still like to test the approximate similarity transformation with retailers. For now, though, they have made it available on GitHub for two reasons: They want colleagues who could build on their work to be able to do so without having to start from scratch; and "if a company has some in-house capabilities that they can take this and test it they are welcome to do so."

More information: Amr Farahat et al. The Multiproduct Newsvendor Problem with Customer Choice, *Operations Research* (2017). DOI: 10.1287/opre.2017.1654

Provided by MIT Sloan School of Management