Safety Stock

We describe how predictive analytics helped a large retailer improve its Buy Online & Pick Up in Store program by making smarter inventory reservation decisions.

Business Problem

In virtually all industries, including retail, supply chain management is a mission-critical function. Inventory shortages and overages caused by inaccurate demand forecasting or imperfect replenishment directly impact both the customer experience and the operating costs related to warehousing and transportation.

Supply chain management is a challenging problem, and its complexity stems partly from the large number of factors that influence demand. Accurate forecasting of demand requires consolidating multiple datasets and analyzing dependencies between various signals. Machine learning can greatly ease this problem, both improving the accuracy of demand forecasting and enhancing replenishment performance so that it closely follows demand. Some studies show that AI technologies can reduce sales lost due to product unavailability by up to 65% and reduce overall inventory by 20 to 50%.[5]

Grid Dynamics was engaged by a large department store chain to improve inventory management using machine learning. The priority business case involved merchandise reservations for its Buy Online & Pick Up in Store (BOPUS) program, which was experiencing two major problems. The first problem was stockouts, when orders could not be fulfilled because the ordered items were sold out, and the second was overbookings, when online orders could not be placed because the desired items had been reserved for in-store selling. Both problems were caused by a nonoptimal merchandise reservation strategy, and machine learning could significantly improve reservation accuracy.

Analysis: Environment and Controls

The BOPUS program allowed customers to place orders online and then pick up their merchandise from a brick-and-mortar store. The order was assembled in the store, with a store employee taking the ordered items directly from store shelves. Sometimes an item was available in a store when the order was placed but had been taken by a customer before the employee came to pick it up.

To avoid stockouts, the retailer had a special order placement protocol for BOPUS orders, as shown in Figure 12. All in-store transactions were tracked, and the inventory-on-hand numbers were updated in near real time. The online system checked this inventory number in the selected store before confirming the order, accepting the order only if the inventory level was above a specific threshold (reservation level). Otherwise, it suggested that the customer choose another store. An out-of-stock exception could still occur at the time of order processing, however, even if the order had been accepted by the online system, because the items were served directly from shelves.

The reservation level check was supposed to decrease the probability of out-of-stock exceptions, but the thresholds were configured manually (and thus suboptimally). The reservation levels were too high for some products, effectively making those products unavailable for online ordering. On the other hand, products for which reservation levels were too low caused out-of-stock exceptions.

Analysis: Objectives

The optimization of reservation levels is driven by two conflicting objectives:

Figure 12. Buy Online & Pick Up in Store environment.

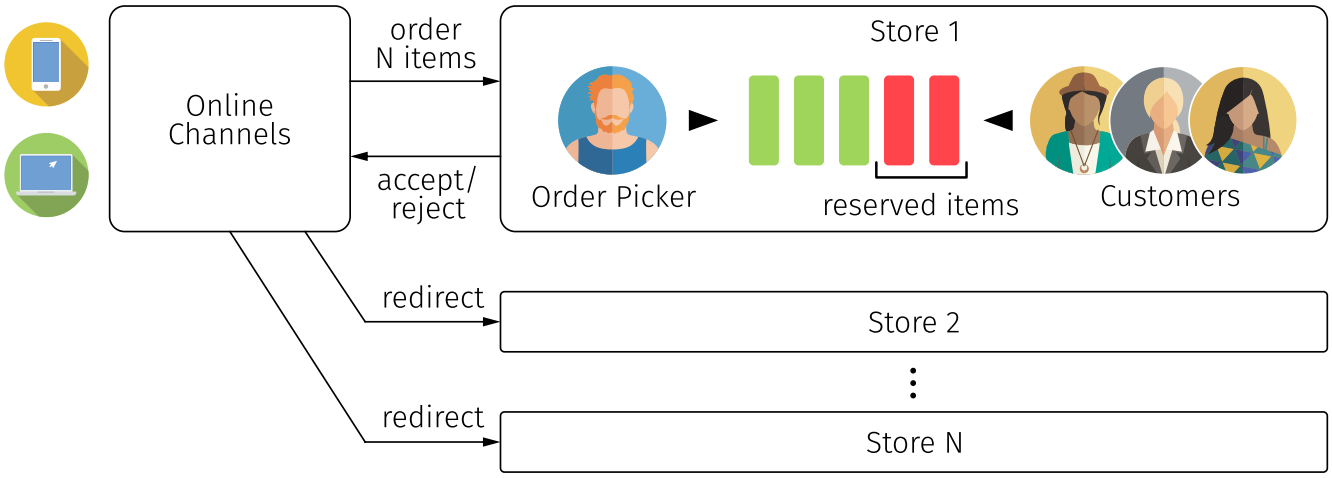

- Maximize product availability. Low reservation levels increase the percentage of products available for BOPUS ordering.
- Maximize order fulfillment rate. High reservation levels increase the order fulfillment rate (or, alternatively, minimize the number of out-of-stock exceptions).

The system should take both of these objectives into account when searching for optimal reservation levels.

Solution: Architecture and AI Usage

One way to address the optimization problem for reservation levels is by posing the following question: How many units of a product should be reserved for in-store customers so that this quantity will not be sold out when an employee needs to process an order? It is essentially a demand prediction problem, and thus the ability to optimize the reservation level for a given product depends on the ability to accurately predict the demand for the product.

The solution included a data collection subsystem to consolidate transactions, product catalog data, prices, promotions, and other datasets useful for demand modeling. Our data science team trained a demand prediction model based on this data, and the model was evaluated on a regular basis to forecast the demand for individual products. The reservation thresholds were automatically adjusted based on the forecasted demand and updated in a transactional system that made decisions on accepting or rejecting orders.

Results

- The percentage of products available for BOPUS was increased from 41% to 50%. This improvement was achieved mainly by lowering the reservation thresholds for slow-moving goods.
- The number of order processing exceptions due to stockouts was maintained at the same or lower levels.

Related Business Cases

The demand prediction capability developed in this solution can be adopted in other inventory-management cases, such as replenishment and warehouse management. Demand forecasting is also a crucial component of price optimization systems, in which the price-demand function needs to be accurately estimated.