

Using Business Analytics to Upgrade Sales Promotions



©shutterstock/whiteMocca

Lennart Baardman, Maxime Cohen, Kiran Panchamgam, and Georgia Perakis present a detailed case study of how business analytics, prediction models, and optimization methods can be used to improve promotion planning. They describe the entire process, from collecting data to computing promotion recommendations for retailers.

Lennart Baardman,
Ross School of Business, University
of Michigan

Maxime C. Cohen,
Desautels Faculty of Management,
McGill University

Kiran Panchamgam,
Oracle Retail Global Business Unit

Georgia Perakis,
Sloan School of Management, MIT

Sales promotions have become ubiquitous because customers expect them. When the American department store JCPenney changed from a pricing strategy based on promotions to one based on consistently low prices, their sales dropped substantially, largely because their customers were conditioned to look for promotions.¹ Given these expectations, it is important that retailers understand their customers' purchasing behavior and be able to determine the right promotion policy for any situation. Fortunately, the combination of large customer datasets and ever-increasing computational power create a unique opportunity for retailers to use advanced analytics to improve their decisions.

It is important that retailers understand their customers' purchasing behavior so they can ultimately determine the right promotion policy for any situation.

The Oracle Retail Global Business Unit (RGBU) provides clear evidence that retailers are increasingly interested in planning promotions efficiently. Oracle RGBU initiated this work after several retail grocery clients asked for software tools to support promotion planning. Retailers advertise promotions (temporary price reductions), through a variety of vehicles, such as product displays, flyers, and commercials.² Their goals include generating extra sales, increasing store traffic, introducing new products, building and maintaining brand loyalty, supporting price discrim-

ination, and retaliating against competing promotions. Oracle Retail's grocery clients found that making frequent use of promotions with only their experience and intuition to guide them was time-consuming and also made them worry that they would leave money on the table. The situation provided a great opportunity for us to develop efficient promotion planning software that would boost retailers' bottom lines.

In collaboration with Oracle RGBU, we developed a promotion planning tool rooted in business analytics. Most earlier promotion planning tools, by contrast, are based on simulating "what-if" scenarios to gradually arrive at the best plan.³ For obvious reasons, these techniques tend to be inaccurate and time-consuming. We thus created a systematic data-driven optimization model that maximizes profits by clearly determining which promotion is right for which product and at what time, all while conforming to business rules. Because consumer demand, and with it profit, is uncertain, we needed to accurately capture consumer behavior. We therefore built the demand functions of our model to be calibrated directly from data.

We created a systematic data-driven optimization model that maximizes profits by clearly determining which promotion is right for which product and at what time, all while conforming to business rules.

Our approach to promotion planning is divided into several stages, described here through

our work with the Oracle RGBU. For more information about the technical components behind this tool, please see our previous work.⁴ In it we showed that, by optimizing promotions, a grocery retailer could increase its profits by 3 to 9 percent. By here describing how we applied our approach to a hardline retailer, yielding a nearly 10 percent increase in profits, we demonstrate that our method is applicable to a broad range of retailers.

The top-tier clients of Oracle RGBU run weekly promotions for over 1,000 stores in roughly 200 categories. For each store and category, the retailer handles between fifty and 600 SKUs.

Business Problem

Promotion planning is an important challenge for retailers. Effective management of promotions can produce substantial benefits which can be absolutely vital to industries such as supermarkets, which are characterized by low profit margins. Nonetheless, planning promotions at a large scale is difficult. The top-tier clients of Oracle RGBU run weekly promotions for over 1,000 stores in roughly 200 categories. For each store and category, the retailer handles between fifty and 600 stock keeping units (SKUs). Effective promotion planning maximizes profits by scheduling price promotions (i.e., temporary price reductions) and promotion vehicles (such as commercials, flyers, and displays) for the right products during the right weeks. At the same time, any approach to promotions must satisfy various business rules set by the

retailer and its vendors. Data-driven analytics are ideal for solving the problems of promotion planning while significantly boosting profits.

For more than fifty years, our partner retailer for this study has

We here focus on these early stages of our process because it is during this period that Oracle RGBU could closely monitor the software’s performance before putting it into full effect. We began by examining

Year	Sales (Units)	Revenue (\$)
2012	1,243,897	12,521,783
2013	1,455,319	12,560,407
2014	1,495,591	12,598,077

TABLE 1: Yearly sales and revenue for the oil category during 2012-2014.




FIGURE 1: Flowchart describing the stages of the recommendation process.

sold outdoor equipment through more than 100 stores in the US Midwest. It has increased its product range over the years to include lawn, garden, farm, and ranch supplies, livestock feed, animal health, pet food and supplies, hardware, plumbing, electrical, automotive, toys, housewares, and work clothing. The company uses a mix of price promotions and vehicles, including temporary price reductions, coupons, buy-one-get-one-free offers, displays, flyers, commercials, and online advertising to continuously promote its products. By centrally planning all promotions several weeks in advance for all 120,000 SKUs in its stores, the retailer is able to tightly integrate its supply chain with its promotion management. Yet this complex planning is both costly and time-consuming. By using our proposed approach, the company can lower these costs by leveraging historical data, automating promotion planning, and significantly reducing the need for human labor.

the retailer’s data and selecting the products and stores we would use for pilot testing. Next, we estimated the demand function and created an optimal promotion plan. We used the large dataset of transactions between 2012 and 2014 to estimate and validate both models. Figure 1 illustrates the stages of the promotion planning process, arranged into three categories: (1) descriptive analytics for product and store selection, (2) predictive analytics for demand forecasting, and (3) prescriptive analytics for promotion optimization. We used different software tools for each stage: Oracle SQL for data collection, R for clustering and demand estimation, Python and Gurobi for optimization, and

Microsoft Excel to build the tool’s user interface.

Product and Store Selection

Our partner retailer provided us with sales data from 157 stores spanning 153 weeks from January 2012 to December 2014. We used the data from the first 104 weeks as a training set and that from the latter 49 weeks as a test set. These sets gave us data on sales, prices, and promotions for each possible combination of week, store, and product, according to month and year. It also contained data on holidays, the square footage of individual stores, and the brand and size of each product. Although several entries were incomplete, the dataset was large enough that we could discard those without significantly reducing its size.

In order to select a set of products and stores for the initial application of an optimized promotion plan, we went on to subsample this dataset. By so doing, we made it possible for both Oracle RGBU and the retailer to monitor the workings of the tool and the resulting sales performance. Selecting the most appropriate prod-




FIGURE 2: Prices and sales for one product from the oil category during 2014.



FIGURE 3: Example of features for a product.

ucts and stores allowed us to control for differences in products and stores if Oracle RGPU or the retailer wanted to assess the efficacy of our promotion planning tool.

Product Selection

Selecting a large group of frequently promoted products gave us many other products to compare with and allowed us to see how optimized promotion planning can increase profits. The greatest opportunity to improve promotion recommendations for this retailer was in the oil category. This category is large (in terms of both item count and sales volume) and the stores promote many of its products routinely (every few weeks). Table 1 shows the yearly sales and revenue of the oil category between 2012 and 2014. This product group contains 137 SKUs, only twenty-two of which have incomplete data. The remaining 115 SKUs form a clean, representative sales and revenue dataset composed of over 99 percent of the oil category.

Figure 2 provides more information on the frequent promotions of products in the oil category, showing the sales volume and

pricing of one product over the course of 2014. The retailer used four prices during this period: a regular price of 3.59 and three promotional prices of 1.99, 2.09, and 2.29. As expected, temporary price reductions immediately increased sales, but the extent of the increase was determined not only by the promotion but also by several other factors.

We applied our promotion recommendations to several products from the largest subcategory of engine oils, referring to them as treated products. We selected these treated products as forming a good representation of the engine oil subcategory. To this end, we examined specific features in engine oils, some of which are shown in Figure 3. Ultimately, we selected three treated products of the same brand but of different grades and oil types. We used all other products as control products (those sold according to the retailer's previous promotion policy).

Store Clustering

We wanted to find a cluster of similar stores whose various dimen-




FIGURE 4: Revenue in the engine oil category in different store clusters during 2012-2014.

sions differed as little as possible. We needed one group of stores to serve as a treatment group, in which our optimized promotion plan would be implemented, and another as a control group, which would continue to use the retailer's existing practice. We created clusters of stores based on traits such as revenue, promotional revenue, number of products sold, and square footage. For this clustering, we used the kernel k-means method with a Gaussian kernel to find stores with highly similar traits, for example, their revenues being close. To make it robust, we tested polynomial and sigmoid kernels, finding similar results with each, and we normalized our data to equalize the scale of our variables.

The algorithm identified nine appropriate clusters of stores. Fig-




FIGURE 5: Map of the selected stores in Kansas.

Store	Sales (Units)	Revenue (\$)	Promotional Sales (Units)	Promotional Revenue (\$)
Store 1	16,936	250,419	3,058	78,176
Store 2	23,589	308,703	3,607	83,841
Store 3	26,446	347,434	3,664	88,965
Store 4	16,401	235,320	2,910	79,952
Store 5	21,989	237,626	3,040	66,138
Store 6	20,726	251,453	2,796	66,425
Store 7	41,846	274,781	4,704	78,036
Store 8	38,804	277,743	4,188	75,271
Store 9	42,443	338,921	6,173	108,605

TABLE 2: Yearly average monthly sales, revenue, promotional sales, and revenue in the engine oil category in the nine stores during 2012-2014 (above: treated stores, below: control stores).

ure 4 presents the average monthly revenue and promotional revenue of each cluster. The chart shows large differences in average monthly revenue between clusters, in part because some clusters were larger than others. To make our results

more robust, we wanted to include a sizable cluster with both large average revenues and large promotional revenues. Cluster 5, with its high revenues and twenty-one stores, satisfied our requirements.

Within this cluster, we wanted to ensure that all the treated stores were in close proximity so that state regulations, on taxes or price tags for example, would not complicate our comparison of different stores or our implementation. To mitigate this concern, we selected nine stores in the state of Kansas as our subset cluster, as shown in Figure 5. We then chose six Kansas stores to receive promotion recommendations, while another three Kansas stores continued to use the retailer's policy, providing a benchmark by which to assess the impact of our promotion planning method.

Table 2 reports the average monthly sales, revenue, promotional sales, and promotional revenue of the treated and control stores during 2012-2014. The revenue variation between stores is minimal. Although there are some significant differences in sales between stores, the variation within each group is small. Overall, this suggests that these stores are quite similar and allow for a robust comparison.

Demand Forecasting

Before we could formulate a promotion optimization model, we needed to create an estimated demand forecasting model. We therefore had to determine the main factors that drive demand, so that our model would yield an accurate forecast. When we were working with grocery products, the three most important factors were timing, products, and pricing. We found that these factors were also important in outdoor products. Indeed, the extensive marketing and economics literature has found that these three factors should be included in most demand forecasting models. However, while much of that literature focuses on causal inference and endogeneity, we focus on demand forecasting and generating accurate predictions.

Sales trends are generally small but steady. Figure 6 shows the average monthly engine oil sales of




FIGURE 6: Time series and trend line of sales in the engine oil category in the nine stores during 2012-2014.




FIGURE 7: Monthly sales in the engine oil category in the nine stores during 2012-2014.

the selected stores, revealing a slight upward trend over time. We built our model to account for this trend by including a variable which represents the focal week that corresponds to each data point.

Many products are also seasonal, meaning they tend to sell more or less during certain periods. Figure 7 shows the variation in the monthly engine oil sales of the selected stores over the three-year period. Specifically, the stores sold less oil than the average during the winter months, and considerably more during the spring and summer months. Spring sales may be partly attributed to the North American planting season, when the engine oil used in planting equipment needs to be refreshed, while summer sales may be influenced by the harvesting season, during which harvesting equipment is used. Another influence on this increase is the better weather which encourages people to travel and to use bicycles, motorcycles, watercraft, and other recreational machines more often. To control for seasonality effects, we include in our model several monthly variables for each observation.

Some product categories, of course, see increased sales during holidays. Figure 8 shows the average engine oil sales of the selected stores during the holiday weeks of the three target years. The largest demand spikes fall during the weeks of Fa-

ther's Day and Thanksgiving Day. The Father's Day spike may be explained in part by the holiday's falling towards the latter half of spring (June in the U.S.), when the weather is warm and people are drawn to outdoor activities like motorcycling. Thanksgiving seems to be a period in which cars are refreshed and agricultural equipment stored for the winter. By contrast, the demand at Christmas and New Year's is relatively low, largely because of store closures, winter weather, and Thanksgiving stockpiling. Our model includes several variables which account for these holiday effects.

Demand is also clearly affected by the product's characteristics. We therefore included many variables that indicate which product each observation indicates. Although we could include product features directly in the model, we have enough data to estimate product-specific parameters and their effects.

In order to optimize promotion planning, we must also consider the effects of pricing and promotions. Knowing that customers are more likely to buy at a reduced price, we used a current product price variable. We also needed to account for price interactions between products in which the promotion of a complementary or substitutable product may increase or decrease demand for the primary product. We have found, however, that cross-product price effects

are weak among products of different brands or sizes. We therefore included cross-product price variables only within each brand. That being said, one certainly could consider cross-price effects among products, brands, categories, and stores, depending on the context and available data. It is similarly important to determine what range of cross-time price effects to include. A recent promotion, for example, may have encouraged customers to stockpile so that they purchase less in the future. In building our datasets, we limited our use of cross-time price effects to the most recent sales, including product price variables only from the most recent weeks. Finally, we had to consider promotion vehicle effects in which displays make customers more aware of a product and thus more likely to buy. We used indicator variables to note whether a promotion vehicle was used, and which one.

With this large number of demand factors, we used a stepwise selection process to build our demand model. We began by estimating several linear regressions describing demand (and its non-linear transformations) as a function of the aforementioned variables, which we included with the guidance of managerial knowledge. We then used three methods, statistical significance, the Akaike Information Criterion (AIC), and the Bayesian Information Criterion (BIC), to iteratively remove some variables. From the models drawn from these three criteria, we selected the one which displayed the highest forecasting accuracy during validation. Because we are primarily interested in prediction, rather than causal inference, this stepwise selection process allows us to quickly generate good models. In estimating our model, we did not use all of our data for training purposes. To assess the model's accuracy, we deliberately kept a subset of our data separate for testing purposes. This allowed for an out-of-sample test which provided us a fair assessment of how our model




FIGURE 8: Average holiday weekly sales in the engine oil category in the nine stores during 2012-2014.

Variable	Product1	Product2	Product3
Intercept	5.2978	6.3014	6.9543
Price Elasticity	-5.7751	-5.9812	-5.6783
Past Price Elasticity	0.8310	1.1004	1.1966

TABLE 3: Estimated product-specific parameters using the training set (2012-2013).

would perform on previously unseen data. We note that other regularization methods for model selection can also be used in this context (e.g., Lasso, Ridge, Elastic net).

Sales in the winter months of December, January, and February were 36, 38, and 41 percent lower, respectively, than those of the spring months.

We have focused on linear regressions, which fit into the optimization framework and are easily interpreted by retail managers and estimated on a large scale, because we are primarily interested in the practical application of our approach. In the end, we used a log-log demand model, which interprets the estimated coefficients as elasticities, for each store and each product. The variables we chose to include were the following: product-specific effect, current price, last week’s price, and seasonal indicators both for the month and for holidays. Because promotion plans are determined at the chain level, we chose to jointly predict the demand for all treated stores.

To create our demand model, we divided the data into two parts: a training set, composed of the first 104 weeks, and a test set of the final 49 weeks. We applied the ordinary least squares regression to the training set to estimate our parameters.

Table 3 presents our parameter estimates for each of the three treated products (all estimates become statis-

tically significant at the 0.05 level). Base sales of product 1 are lower than those of products 2 and 3, but the intercept estimates are relatively close. And while all products show a similar estimated price elasticity (between -5.6783 and -5.9812), products 2 and 3 show a stronger past price effect.

Table 4 shows the parameter estimates for the effects of trend, season, and holidays (the variables again become statistically significant at 0.05). The table also includes the demand factor (the exponentiated estimate), which shows the extent to which demand increases or decreases during a given month or holiday. The first row indicates the small positive estimate of the demand trend. The demand factor indicates a sales increase of approximately 0.17 percent every week, equivalent to a yearly increase of 9.24 percent. The second part of the table shows the parameter estimates and demand factors for sales by month. We have omitted April, May, and June because the estimation method determined their parameter estimates to be statistically insignificant (i.e., indif-

ferent from 0). The estimated sales for the other nine months were lower than those of these three base months. Sales in the winter months of December, January, and February were 36, 38, and 41 percent lower, respectively, than those of the spring months. These numbers confirm our expectation that the winter would have lower sales, while spring would see the highest sales. The final part of the table reports the

estimates for holidays. Having corrected for the demand trend and for the monthly base demand, only New Year’s, Martin Luther King Day, and Christmas have a significant impact on sales. These three holidays lead to lower sales than other holidays, such as Father’s Day and Thanksgiving. During the New Year’s week, demand drops by 32 percent; during the Christmas week it decreases by 29 percent. We attribute the large Christmas decrease to the holiday closure of most stores.

Having estimated the demand model, we could now test its accuracy in forecasting demand. We used our log-log demand model to estimate future demand and then applied the exponential function to calculate the actual predicted demand. By applying our estimated model to the test set we could compute out-of-sample forecasting metrics that assess the model’s accuracy with data it has not seen before. We focused on three metrics: R^2 (coefficient of determination), MAPE (mean absolute percentage error), and MdAPE (median absolute percentage error). Our aim was for the R^2 to be

Variable	Estimate	Factor
Trend	0.0017	1.0017
January	-0.4751	0.6218
February	-0.5275	0.5901
March	-0.1753	0.8392
July	-0.1319	0.8764
August	-0.1560	0.8556
September	-0.1445	0.8655
October	-0.1736	0.8406
November	-0.3305	0.7186
December	-0.4434	0.6419
New Year’s Day	-0.3834	0.6815
MLK Day	-0.0973	0.9073
Christmas Day	-0.3408	0.7112

TABLE 4: Estimated time-specific parameters using the training set (2012-2013).

Forecasting Metric	Oil Category	Treatment Brand	Product1	Product2	Product3
R^2	0.89	0.90	0.82	0.86	0.93
MAPE	0.6357	0.3728	0.3838	0.4359	0.2987
MdAPE	0.2994	0.2305	0.2777	0.2018	0.2544

TABLE 5: Forecasting metrics of the estimated demand model on the test set (2014).

close to 1, while we preferred the MAPE and MdAPE to be close to 0.

Table 5 presents our out-of-sample forecasting metrics. Our results showed very good prediction accuracy for the retail industry, with an in-sample R^2 for the oil category of 0.92 and an out-of-sample R^2 of 0.89.⁵ The proximity of the in-sample and out-of-sample R^2 also indicates that the model generalizes well. We observed similar results at the brand level, at the individual product level, and using the MAPE and MdAPE.

Overall, our demand model produces highly accurate predictions.

In Figure 9, we present a comparison of the actual and predicted sales of one of the treated products during the testing period. The predictions follow the same pattern as the actual sales, and often with a similar magnitude. Only in some of the highest selling periods does our model under-predict, and this difference is relatively small with a MAPE of 22.76 percent and an MdAPE of 19.68 percent throughout all the promotion periods. Overall, our demand model produces highly accurate predictions.

Promotion Optimization

Having built and tested the demand forecasting model, we were

ready to design the mathematical optimization model that would prescribe promotions. Our goal was to maximize profits during the upcoming planning window by determining which products to promote, to what extent, and when. A typical retailer must plan for around 250 products (our representative category), choosing from about twenty prices (several of them ending in 99 cents), and roughly thirteen time periods (for example a quarter consisting of thirteen weeks), for a total of 65,000 binary decision variables which

difference between the unit price and the cost multiplied by the expected demand.

Our goal was to maximize profits during the upcoming planning window by determining which products to promote, to what extent, and when.

We also had to include business rules set by retailers through constraints in our mathematical optimization model. For each product, many retailers only offer prices from a pre-determined price ladder. We




FIGURE 9: Actual and predicted sales of one treated product during 2014.

determine whether each product should be offered at one of the prices during each period. We have not considered the choice of promotion vehicles here because our partner retailer was interested primarily in price decisions. We have considered methods to schedule promotion vehicles in previous studies.⁶ We used the total expected profits, equal to the sum of the expected profit of each product in each period, as the objective in our mathematical optimization model. Here, the expected profit of a given product equals the

built this constraint into our model to ensure that the prices it recommended corresponded to this price ladder. In order to preserve their image, brands and stores often limit the number of promotions for each product. We therefore included an optional constraint on the number of promotions for each product. Similarly, most stores try to avoid running two promotions back-to-back, so we added a constraint to ensure a minimum separating period between promotions for each product. Our model makes it easy to include

additional business rules, according to the requirements of the retailer.

The formulation we built around these requirements is a non-linear integer optimization problem, known to be difficult to solve⁷. By drawing upon the methods developed in previous studies, however, we were able to generate approximate optimized promotion plans.⁸ The solution relies on using a linear approximation of the original non-linear problem. Interestingly, this approximation often produces promotion plans with a profit very close to optimal. It also runs significantly faster than the optimal method, allowing us to solve the problem in seconds on an ordinary computer.

“Without altering our business processes, just with optimizing the price-point for a promotion, the team of researchers showed us that we can improve our profit margins by as high as 10 percent for some of our products. This is a very significant improvement, considering that our margins are thin.”

We first formulated the promotion optimization problem by combining it with our demand forecasting model and design parameters. Next, we solved the optimization problem for 2014, producing a schedule of recommended promotions. We were interested to note that these promotions all give the same substantial discount and are spaced out evenly over the year. Managers will find this insight, corroborated by the results of previous real-world studies, very useful.

We went on to compare our optimized promotion plan to the actual promotion policy which our partner retailer used in 2014 to determine how much optimized promotion planning could have improved that year’s profits. Table 6 reports the store’s potential improvement in sales, revenue, and profit. Column (a) shows the retailer’s actual sales, revenue, and profit. Column (b) applies our demand forecasting model to historical prices to compute the sales, revenue, and profit as if our demand forecasts were reality. The differences between these two columns thus indicate the aggregate error in our demand prediction model. Our model’s 12 percent difference between yearly actual and forecasted sales is notably smaller than the 24 percent of the MdAPE prediction, showing that our model works well for the treated products and treated stores in 2014. Column (c) reports the sales, revenue, and profit which our optimization model would produce. To ensure that our comparison is fair, we compare Col-

umns (b) and (c), both of which use the same demand forecasting model. The final column shows that our optimized promotion policy would likely have increased our partner retailer’s profits by 10 percent and its revenue by 1 percent, all at a similar sales level. In short, by using our tool to optimize promotion planning, a retailer can significantly improve its profits, while maintaining its revenues and sales. In looking over the recommended promotion policy, we noted that our model is able to keep prices high during periods of high demand by spacing promotions farther apart than the retailer had been. Additionally, as our partner’s chief information officer put it: “Without altering our business processes, just with optimizing the price-point for a promotion, the team of researchers showed us that we can improve our profit margins by as high as 10 percent for some of our products. This is a very significant improvement, considering that our margins are thin.”

KPI for 2014	Actual Promotions		Optimized Promotions	
	Actual Demand	Forecasted Demand	Forecasted Demand	Improvement
	(a)	(b)	(c)	$\frac{(c) - (b)}{(b)}$
Sales (Units)	690,414	612,265	612,138	-0.02%
Revenue (\$)	\$1,478,905.89	\$1,319,001.47	\$1,334,419.66	+1.17%
Profit (\$)	\$169,190.53	\$157,535.49	\$173,193.78	+ 9.94%

TABLE 6: Key performance indicators (KPI) for the treated products in all stores during 2014.




FIGURE 10: Comparison of the revenues and profits of current practice and two optimization scenarios.

As well as optimizing our partner retailer's promotion policy for 2014, we also ran several hypothetical scenarios. Because our models run very fast (processing hundreds of products within milliseconds), we could rerun the model with various parameter settings. Our promotion policy, described by Table 6, included the business rule that yearly revenues should never decrease relative to the previous year. Figure 10 compares the profits and revenues without this business rule (Scenario 1) to those with it (Scenario 2) as well as to the revenue and profit of the actual promotion plan (Current).

Author Bios



Lennart Baardman is an Assistant Professor of Technology and Operations at the Ross School of Business at the University of Michigan. His work focuses on using analytics to solve

operational challenges in revenue management, pricing, supply chain management, and logistics. This research is inspired by his collaborations with Adobe, CMPC, Johnson & Johnson, and Oracle, among others.



Maxime Cohen is a Professor of Retail and Operations Management, codirector of the Retail Innovation Lab, and a Bensadoun Faculty Scholar at McGill University. His core

expertise lies at the intersection of data science and operations. He has collaborated with Google, Waze, Oracle Retail, IBM Research, Via, Spotify, Aldo Group, Couche-Tard/Circle K, and Staples as well as several retailers and startups.

We see this work as one important step in improving retail operations through data analytics.

The steps described here are only the early stages of implementing our promotion planning approach at a

large retailer. This collaboration between industry and academy allowed us to develop a systematic data-driven approach to optimizing promotion planning. This method can be applied in many retail settings, works for general demand models, can incorporate a wide range of business rules, and is calibrated using real transaction data.

We see this work as one important step in improving retail operations through data analytics. As our partner retailer said: "We have worked with the team of researchers for a little more than a year and it is truly amazing to see the growth and value that this work has brought, starting with

Retail Promotion & Markdown Optimization, and Oracle Retail Offer Optimization.



Georgia Perakis is William F. Pounds Professor of Operations Management, Operations Research, and Statistics at MIT Sloan. She codirects MIT's Operations Research Center and is Faculty Director of its Executive MBA Program. She is editor in chief of MSOM and an INFORMS Fellow in recognition of her lifetime contributions. Perakis' analytics work has been applied to operations, retail, supply chain and healthcare. Her collaborations include Adobe, IBM Research, Lahey Clinic, Oracle Retail, Wayfair, Zara, and many more.



Kiran Panchamgam is a Senior Principal Data Scientist at the Oracle Retail Global Business Unit. His work spans optimization, prediction, and retail analytics. He identifies retail challenges,

devises solution methods, develops enterprise software systems, and implements them for a range of retailers. He has three granted and nine pending patents with USPTO and is the product manager for Oracle Retail Assortment & Space Optimization, Oracle

Endnotes

1. <https://www.priceintelligently.com/blog/bid/152018/lessons-from-the-failure-of-j-c-penney-s-new-pricing-strategy>
2. Blattberg RC, Neslin SA (1990) "Sales promotion: Concepts, methods and strategies" (Prentice Hall).
3. Anderson ET, Fox EJ (2019) How price promotions work: A review of practice and theory. Dubé JP, Rossi PE, eds, Handbook of the Economics of Marketing, Volume 1, volume 1 of Handbook of the Economics of Marketing, chapter 9, 497-552 (North-Holland).
4. <https://www.sdcexec.com/software-technology/press-release/21138328/blue-yonder-only-15-of-global-retailers-supply-chains-are-prescriptive-or-autonomous>
5. Cohen MC, Leung NHZ, Panchamgam K, Perakis G, Smith A (2017) The impact of linear optimization on promotion planning. Operations Research 65(2):446-468.
6. Cohen MC, Kalas J, Perakis G (2021) Promotion optimization for multiple items in supermarkets. Management Science 67(4): 2340-2364.
7. Baardman L, Cohen MC, Panchamgam K, Perakis G, Segev D (2019) Scheduling promotion vehicles to boost profits. Management Science 65(1):50-70.
8. Ali OG, Sayin S, Van Woensel T, Fransoo J (2009) Sku demand forecasting in the presence of promotions. Expert Systems with Applications 36:12340-12348.
9. Ferreira KJ, Lee BHA, Simchi-Levi D (2016) Analytics for an online retailer: Demand forecasting and price optimization. Manufacturing & Service Operations Management 18(1):69-88.
10. Cohen MC, Kalas J, Perakis G (2020) An efficient algorithm for dynamic pricing using a graphical representation. Production and Operations Management 29(10):2326-2349.
11. Cohen MC, Leung NHZ, Panchamgam K, Perakis G, Smith A (2017) The impact of linear optimization on promotion planning. Operations Research 65(2):446-468.
12. Baardman L, Cohen MC, Panchamgam K, Perakis G, Segev D (2019) Scheduling promotion vehicles to boost profits. Management Science 65(1):50-70.
13. Cohen MC, Leung NHZ, Panchamgam K, Perakis G, Smith A (2017) The impact of linear optimization on Promotion planning. Operations Research 65(2):446-468.