

Upgrading Promotions Using Business Analytics

Lennart Baardman

Ross School of Business, University of Michigan, Ann Arbor, MI 48109, baardman@umich.edu

Maxime C. Cohen

Desautels Faculty of Management, McGill University, maxime.cohen@mcgill.ca

Kiran Panchamgam

Oracle RGPU, Burlington, MA 01803, kiran.panchamgam@oracle.com

Georgia Perakis

Sloan School of Management, MIT, Cambridge, MA 02139, georgiap@mit.edu

Sales promotions are ubiquitous to the extent that customers are expecting retailers to offer promotions. When the American department store JCPenney changed from a pricing strategy based on promotions to one based on low prices, their sales dropped substantially mainly due to customers being conditioned to promotions.¹ Given that customers may expect promotions, it is important for retailers to understand their customer's purchasing behavior and to determine the right promotion policy for each context. Fortunately, large customer datasets in conjunction with increasing computational power create a unique opportunity for retailers to upgrade their decision-making process by using advanced analytics. In this article, we focus on how business analytics can improve promotion planning.

Retailers are increasingly interested in planning promotions efficiently, as witnessed by the Oracle Retail Global Business Unit (RGPU). This work was initiated by the ask of several grocery retail clients for software tools supporting promotion planning. Promotions (i.e., temporary price reductions) are frequently used by retailers via several vehicles such as product displays, flyers, and commercials (Blattberg and Neslin 1990, Anderson and Fox 2019). Promotions are used with the goal of generating extra sales, increasing store traffic, introducing new products, creating and maintaining brand loyalty, aiding price discrimination, and retaliating competitor promotions. Due to their frequent use and their current management based on experience and intuition, the aforementioned grocery clients of Oracle Retail found the planning of promotions to be time-consuming and they were worried about leaving money on the table. Altogether, this presented us with a great opportunity to impact retailers' bottom lines by developing efficient promotion planning software.

¹ <https://www.priceintelligently.com/blog/bid/152018/lessons-from-the-failure-of-j-c-penney-s-new-pricing-strategy>

In collaboration with Oracle RGBU, our intent is to develop a promotion planning tool founded upon business analytics. Most of the promotion planning tools currently used in the industry are not based on predictive and prescriptive analytics, but rather on simulating different “what-if” scenarios.² In contrast to these approximate techniques, we develop an optimization model that explicitly determines the right promotion for the right product at the right time so as to maximize profits, while accounting for business rules. As consumer demand, and hence profit, is uncertain, it is important to capture consumer behavior accurately. For this reason, our model incorporates demand functions that are directly calibrated from data.

The rest of this article presents the different stages of implementing our promotion planning approach at a large client of the Oracle RGBU. The technical components behind this promotion planning tool were described in previous work (see Cohen et al. 2017, 2020b, Baardman et al. 2019). Specifically, it was shown that optimizing promotions at a grocery retailer could yield a profit increase of 3 to 9%. In this article, we show that these findings generalize by presenting the implementation of our approach at an outdoor retailer with a profit increase close to 10%. Our hope is that this article can act as a case study for various applications, not only grocery and outdoor retail, but also for other verticals.

1. Business Problem

Promotion planning is an important challenge for retailers. The importance is evident from the substantial benefits that can result from managing promotions effectively. Certain industries, such as supermarkets, are characterized by low profit margins, and thus, can benefit from promotions that efficiently manage these margins. The challenge comes from the difficulty of planning promotions at a large scale. Oracle RGBU’s top-tier clients run weekly promotions for over 1,000 stores with roughly 200 categories each containing 50 to 600 stock keeping units (SKUs). An effective promotion planning approach 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. While scheduling promotions, this approach also needs to satisfy various business rules set by the retailer and vendor funds. Our goal is to show that data-driven analytics can solve the promotion planning problem while significantly boosting profits.

In this article, we focus on the initial stages of implementing such a data-driven model to a large retailer. The partner retailer has been supplying outdoor equipment to farms and ranches in over 100 stores in the Midwest of the U.S. for over 50 years. The product offerings have increased over the years to include lawn, garden, farm and ranch supplies, livestock feed, animal health,

² <https://www.sdcexec.com/software-technology/press-release/21138328/blue-yonder-only-15-of-global-retailers-supply-chains-are-prescriptive-or-autonomous>

pet food and supplies, hardware, plumbing, electrical, automotive, toys, housewares, and work clothing. These products are continuously promoted through a mix of price promotions and vehicles: temporary price reductions, coupons, buy-one-get-one-free offers, displays, flyers, commercials, and online advertising. The retailer plans these promotions centrally for all 120,000 SKUs in all stores, several weeks in advance. This allows the retailer to tightly integrate supply chain and promotion management. However, it also creates a costly time-consuming process.

In the context of this retailer, we will present the recommendation process prior to the implementation of our promotion planning tool. We focus on these early stages as this is where Oracle RGBU can keep a close eye on the software’s performance before it is implemented. First, we describe the retailer’s data and how we selected the products and stores for initial pilot testing. We next estimate the demand function as well as optimize the promotion plan. Both models are estimated and validated using a large transaction dataset from 2012 to 2014. As illustrated in Figure 1, the stages of the promotion planning process can be categorized as follows: (1) product and store selection, (2) demand forecasting, and (3) promotion optimization. In each of these stages, we used different software tools: Oracle SQL for data collection, R for clustering and demand estimation, Python and Gurobi for optimization, and Microsoft Excel for building the user interface of our promotion planning tool.




Figure 1 Flowchart describing the stages of the recommendation process.

2. Product and Store Selection

The partner retailer provided us with data from 157 stores. The dataset spans a period of 153 weeks between January 2012 and December 2014. We split the data into a training set composed of the first 104 weeks, and a test set of the final 49 weeks. For each week-store-product combination, we have data on sales, prices, and promotions of the products in the stores. The dataset also contains data on the month, year, and relevant holidays, the store square footage, as well as the brand and

size of each product. Several data entries are incomplete, but the dataset is large enough for us to be able to discard these observations without significantly reducing the size of the dataset.

We further subsample this dataset, as we want to select a set of products and stores that would be the initial recipients of an optimized promotion plan. This treatment allows Oracle RGBU to check the workings of the tool and its resulting performance. By selecting the most appropriate products and stores, we can control for structural product and store differences if Oracle RGBU or the retailer want to assess the potential impact of our promotion planning tool.

2.1. Product Selection

Naturally, we want to select a group of products such that there are a large number of frequently promoted products. This selection provides us with many other products to compare against and allows us to see how optimized promotion planning can improve profits. In this case, the retailer saw the most potential for improved promotion recommendations within the oil category. The oil category is large and many of the products are often promoted.

Table 1 presents the yearly sales and revenue in the oil category between 2012 and 2014. The category contains 137 SKUs, out of which only 22 SKUs have incomplete data. The remaining 115 SKUs form a representative and clean dataset that constitutes over 99% of the entire oil category in terms of yearly sales and revenue.

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 entire oil category during 2012–2014.

Specifically, products in the oil category are promoted frequently as seen in Figure 2, which presents the volume of sales and the corresponding price levels over time for one of the products in 2014. During this period, the retailer used four price levels: a regular price of 3.59 and three promotional prices of 1.99, 2.09, and 2.29. As expected, a temporary price reduction immediately increases sales, yet by how much depends on the promotion as well as on other factors.

Within the oil category, we select several “treated” products (i.e., products that receive our promotion recommendations) from the largest subcategory of engine oils. These treated products are chosen to yield a good representation of the engine oil subcategory. To this end, we look at specific features of engine oils, some of which are shown in Figure 3. Ultimately, we select 3 treated products from the same brand but with different grades and oil types. All other products can be used as control units.




Figure 2 Time series of prices and sales for one product from the oil category in 2014.



Figure 3 Example of features for a product.

2.2. Store Clustering

We aim to find a cluster of similar stores. This will allow us to compare the results between stores as well as to assess the actual implementation of an optimized promotion plan. To cluster stores, we use the Kernel K-means method with a Gaussian kernel. This method creates clusters of stores based on the similarity of their features such as revenue, promotional revenue, number of products sold, and square footage. For robustness, we normalized our data and tested polynomial and sigmoid kernels and found similar results.

The algorithm identified 9 clusters of stores. Figure 4 presents the average monthly revenue and promotional revenue in each cluster. The chart shows a large difference in average monthly revenue between the different clusters, but the average monthly promotional revenue is quite stable. Part of this difference can be explained by the fact that certain clusters are larger than others. To increase the robustness of our results, we want to select a sizable cluster with both large average revenues and promotional revenues. Satisfying all these requirements is cluster 5, which has sizable revenues and consists of 21 stores.

Within this cluster, we want to ensure that all the treated stores are in close vicinity. Otherwise, state regulations, for example on taxes or price tags, could affect the ease of comparing different stores as well as our actual implementation. To mitigate this concern, we specify our cluster to

Store	Sales (Units)	Revenue (\$)	Promotional Sales (Units)	Promotional Revenue (\$)
Store1	16,936	250,419	3,058	78,176
Store2	23,589	308,703	3,607	83,841
Store3	26,446	347,434	3,664	88,965
Store4	16,401	235,320	2,910	79,952
Store5	21,989	237,626	3,040	66,138
Store6	20,726	251,453	2,796	66,425
Store7	41,846	274,781	4,704	78,036
Store8	38,804	277,743	4,188	75,271
Store9	42,443	338,921	6,173	108,605

Table 2 Average monthly sales, revenue, promotional sales, and revenue in the engine oil category in the 9 stores during 2012–2014 (above: treated stores, below: control stores).

literature on demand models. In contrast to the focus on causal inference and endogeneity in this literature, we focus on demand forecasting and generating accurate predictions.

Generally, we see small but steady trends in the sales. Figure 6 shows a time series of monthly engine oil sales averaged over the selected stores. This plot supports the existence of a slight upward trend over time in monthly sales. To account for this trend, our model will include a variable representing the focal week that corresponds to each data point.




Figure 6 Time series and trend-line of sales in the engine oil category in the 9 stores during 2012–2014.

Additionally, many products exhibit seasonality, meaning that certain time periods experience significantly more or less sales than usual. Figure 7 presents the monthly engine oil sales averaged over the selected stores and a period of 3 years. We observe a considerable variation in the sales. Specifically, the winter months sell less on average, while the spring and summer months sell well beyond average. A partial explanation for the first period could be that it corresponds to the North American planting season when the engine oil used for planting machines needs to be refreshed, whereas the second season corresponds to the harvesting season during which harvesting machines are used. Other explanations include the fact that the weather improves, so that people use bicycles, motorcycles, and other recreational machines more often. To control for seasonality effects, we include several variables that capture the month associated with each observation.




Figure 7 Monthly sales in the engine oil category in the 9 stores during 2012–2014.

Furthermore, certain product categories see increased sales during holidays. Figure 8 shows engine oil sales during holiday weeks averaged over the selected stores and years. During the weeks of Father’s day and Thanksgiving day, we notice the largest demand spikes. The spike on Father’s day could be explained by the fact that this holiday falls towards the latter half of the spring (June in the U.S.), when the weather has improved and people may be interested in recreational activities such as riding motorcycles. Thanksgiving seems to be a period in which cars are refreshed for the winter. On the other hand, the demand on Christmas and New Year are relatively low, possibly due to store closure, the winter season, and potential stockpiling from Thanksgiving. To account for such holiday effects, we include several variables that incorporate holidays linked with observations.




Figure 8 Average holiday week sales in the engine oil category in the 9 stores during 2012–2014.

Certain product characteristics can clearly affect demand. To capture the differences between products, we include many variables that indicate which product is associated with each observation. Even though we can include the product features directly into the model, we have enough data to estimate product-dependent parameters that capture the product-specific effects of these time-stationary product features.

As the focus of this article is to optimize the promotion planning, we evidently also consider the effect of pricing and promotions. First, there is the effect of the current price, as customers are

more likely to buy under a reduced price. We capture this effect through a current product price variable. On the other hand, there are cross-product price effects whereby a product’s demand might increase or decrease when another complementary or substitutable product is promoted. Based on our experience, cross-product price effects are weak among products of different brands or size. Hence, we include cross-product price variables only within each brand. Similarly, there are cross-time price effects, as a recent promotion may have induced customers to stockpile the product and purchase less in the future. In our datasets, cross-time price effects are limited to the most recent sales. Hence, we include past product price variables only for the most recent weeks. Finally, there are promotion vehicle effects whereby displays or features create awareness about a product and increase the likelihood for customers to buy. These effects can be included through indicator variables for whether a promotion vehicle was used.

Given the large number of demand factors, we use a stepwise selection process to estimate our demand model. We initially estimate several linear regressions that describe demand (and its non-linear transformations) as a function of all the aforementioned variables. We then iteratively remove variables based on their statistical significance, the Akaike Information Criterion (AIC) or the Bayesian Information Criterion (BIC). The variable selection process was also guided by managerial knowledge. Out of the models determined by these criteria, we select the final model based on having the best forecasting accuracy on a hold-out validation dataset. As our main interest is in prediction (as opposed to causal inference), stepwise selection allows us to generate good models quickly, which can then be tested out-of-sample for the best forecasting accuracy. We note that regularization is another good alternative for model selection.

As we are interested in the practical applicability of our approach, we focus on linear regressions. In practice, linear models are interpretable for retail managers, they are easier to estimate at a large scale, and they fit well into the optimization framework. In the end, we used a log-log demand model for each store and each product. A log-log model offers the advantage of interpreting the estimated coefficients as elasticities. We ultimately include the following variables: product intercept, current price, last week price, binary seasonality indicators for the month and for holidays. As the promotion plan is determined at the chain-level, we predict the aggregate demand jointly for all treated stores.

To estimate our demand model, we split the data into two parts: a training set composed of the first 104 weeks and a test set with the final 49 weeks. We then estimate the parameters by using the ordinary least squares regression on the training set. In Table 3, we present the product-specific parameter estimates for the 3 treated products (all estimates are statistically significant at the 0.05 level). We observe that the base sales of product 1 are smaller relative to products 2 and 3, but the estimates are relatively close. The estimated price elasticities have all a similar magnitude (between -5.6783 and -5.9812), but the past price effect is stronger for products 2 and 3.

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 on the training set (2012–2013).

Table 4 reports the time-specific parameter estimates of the trend, seasonality, and holiday effects (all the presented variables are statistically significant at the 0.05 level). In addition, the table includes the demand factor, which is the exponentiated estimate, to show to what extent the demand increases or decreases in a given month or during a specific holiday. The first row corresponds to the small positive estimate of the demand trend. The demand factor indicates that sales increase by approximately 0.17% every week, or equivalently a yearly increase of 9.24%. The second part of the table shows the parameter estimates and demand factors for the monthly sales. Three months (April, May, and June) are left out of the table, because the parameter estimates were statistically insignificant. The negative estimates for the 9 other months show that their estimated sales are lower than the estimated sales of the 3 base months. Especially, the demand factors of December, January, and February show that sales in the winter months are 36%, 38%, and 41% lower when compared to spring months. This confirms our earlier intuition that the winter period can admit lower sales, whereas spring sees the highest sales. Finally, the third part of the table reports the estimates for the holiday factors. Having corrected for the demand trend and monthly base demand, the only significant impact on sales can be seen during New Year, Martin Luther King day, and Christmas. Compared to other holidays, such as Father’s day and Thanksgiving, these three holidays lead to lower sales. In the New Year week, the demand drops by 32%, whereas in the Christmas week it decreases by 29%. We can attribute the large decrease in the Christmas week to the closure of most stores during this period.

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 Day	-0.3834	0.6815
MLK Day	-0.0973	0.9073
Christmas Day	-0.3408	0.7112

Table 4 Estimated time-specific parameters on the training set (2012–2013).

Having estimated the demand model, we can now test how accurately it forecasts demand. We can apply our estimated model to the test set and compute out-of-sample forecasting metrics that assess

the model fit. Specifically, we consider three metrics: R^2 (coefficient of determination), $MAPE$ (mean absolute percentage error), and $MdAPE$ (median absolute percentage error). Generally, we would like the R^2 to be close to 1, while the $MAPE$ and $MdAPE$ close to 0.

The out-of-sample forecasting metrics are presented in Table 5. In this case, the in-sample R^2 for the oil category is 0.92, while the out-of-sample R^2 is 0.89. Note that this is considered as a very good prediction accuracy in the retail industry (e.g., see Ali et al. 2009, Ferreira et al. 2016, Cohen et al. 2017), especially for products with a less stable sales rate such as engine oils. In addition, the fact that the in-sample and out-of-sample R^2 are close together indicates that there is no strong overfitting and that the model generalizes well. Similar results are observed for the prediction accuracy at the brand level, at the individual product level, and when looking at the $MAPE$ and $MdAPE$.

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).

In Figure 9, we present a comparison of the actual and predicted sales for one of the treated products during the testing period. One can see that the predictions follow the same pattern as the actual sales, often with a similar magnitude. Only in some of the highest selling periods, our model under-predicts. Nevertheless, this difference is small in relative terms with a $MAPE$ of 22.76% and $MdAPE$ of 19.68% during all the promotion periods. Overall, we conclude that our proposed demand model results in a high prediction accuracy.




Figure 9 Time series of actual and predicted sales for one treated product during 2014.

4. Promotion Optimization

After estimating the demand forecasting model, we can formulate the optimization problem to prescribe promotions. The objective is to maximize the expected profits during the upcoming selling season by deciding which products to promote, the promotion depth, and when to schedule the promotions. For a typical retailer, the number of products to plan for is around 250 (for a representative category), the number of prices to choose from is around 20 (e.g., several prices ending in 99 cents), and the number of time periods to plan for is 13 (a quarter of 13 weeks). This means that there are 65,000 binary decision variables indicating whether each product is offered at one of the prices during each period. Since this retailer was primarily interested in price decisions, we do not consider promotion vehicles here, though we refer to Baardman et al. (2019) for methods to schedule promotion vehicles. As our objective, we use the expected profits, which is equal to the sum of the profit of each product in each period (the profit of a given product equals the difference between the unit price and the cost multiplied by the expected estimated demand).

The business rules set by retailers can be included as constraints. Many retailers only offer prices from a pre-determined price ladder for each product. We incorporate a constraint to the model to ensure that the recommended prices come from this price ladder. Often, there is a limitation on the number of promotions for each product to preserve the image of the brand or store. To capture such a rule, we include a constraint that allows at most a fixed number promotions to be used for each product. Similarly, there is a norm against two promotions following each other immediately. To satisfy this “no-touch” rule, we add a constraint that ensures at least a fixed number of separating periods between successive promotions for each product. One can naturally include additional business rules, depending on the requirements of the retailer.

The resulting formulation is a non-linear integer optimization problem, which is proven to be difficult to solve (Cohen et al. 2020a). However, by using the methods developed in previous work (Cohen et al. 2017, 2020b), we can generate approximate optimized promotion plans. The machinery relies on using a linear approximation to the original non-linear problem. Interestingly, the profit of this approximation is often very close to the optimal profit. Additionally, the running time of this approximation method is significantly faster than the optimal method, allowing us to solve the problem within seconds on a standard computer.

The first step is thus to formulate the promotion optimization problem by plugging the demand forecasting model and setting the design parameters (e.g., price ladder and business rules). The second step is to solve the optimization problem for 2014. By analyzing the solution, we observe that all the recommended promotions have the same substantial discount and are spaced out over the selling season. This insight can be useful to managers and is corroborated by results seen on previous real-world settings (Cohen et al. 2017). We can then compare our optimized promotion

plan to the actual promotion policy that was implemented in 2014 by the retailer. Altogether, this backtest provides an empirical validation on historical data of how much an optimized promotion planning could have improved the 2014 profits. Table 6 reports the potential improvement in total sales, revenue, and profit.

KPI for 2014	Actual Promotions	Actual Promotions	Optimized Promotions	Improvement $\frac{(c) - (b)}{(b)}$
	Actual demand	Forecasted demand	Forecasted demand	
	(a)	(b)	(c)	
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.

Column (a) reports the actual sales, revenue, and profit. In column (b), we use historical prices to compute the sales, revenue, and profit, but instead of using historically realized demand we use our demand forecasting model. The differences between columns (a) and (b) indicate the aggregate error in our demand prediction model. Note that the 12% difference between the yearly actual and forecasted sales is smaller than the 24% *MdAPE* in predicting demand. Column (c) reports the resulting sales, revenue, and profit obtained from our optimization model. To ensure a fair comparison, we compare Columns (b) and (c) and hence, use the same demand forecasting model for both policies. The final column shows that our optimized promotion policy could lead to nearly 10% additional profits, around 1% extra revenue, and a similar sales level. This backtest suggests that optimizing the retailer’s promotion planning using our tool can have a significant impact on the profits, while maintaining the level of revenues and sales. Interpreting this result, it is likely that the model is able to capture a larger portion of demand at full price by spacing out promotions further apart relative to current practice. As the retailer’s Chief Information Officer puts 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% for some of our products. This is a very significant improvement, considering that our margins are thin.”

In addition to optimizing the promotion policy for 2014, we also run multiple “what-if” scenarios. Since our methods run very fast (within milliseconds for instances with hundreds of products), we are able to rerun the model under a variety of parameter settings. Our promotion policy reported in Table 6 was designed under the business rule that the yearly revenues should not decrease relative to last year. In Figure 10, we compare the profits and revenues when this business rule is relaxed (Scenario 1) to the case where this business rule is imposed (Scenario 2). We also compare the performance relative to the revenue and profit of the actual promotion plan (Current). Ultimately, our discussions with the managers conveyed the importance of including such a business rule (Scenario 2), as a 12.54% loss in yearly revenues (Scenario 1) could be too risky.




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

In this article, we discussed the early stages of implementing our promotion planning approach at a large retailer. However, in this collaboration between industry and academia, we developed a general data-driven approach to optimize promotion planning, which can be applied to many retail settings, works for general demand models, captures a wide range of business rules, and is calibrated using transaction data. We see this work as one of the steps in improving retail operations through the use of data analytics. We conclude with a quote from our partner retailer: “We have been working with the team of researchers for a little over than a year and it is truly amazing to see the growth and value of this work, starting with understanding raw data to providing useful and significant insights.”

References

- Ali ÖG, Sayın S, Van Woensel T, Fransoo J (2009) Sku demand forecasting in the presence of promotions. *Expert Systems with Applications* 36:12340–12348.
- 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).
- Baardman L, Cohen MC, Panchamgam K, Perakis G, Segev D (2019) Scheduling promotion vehicles to boost profits. *Management Science* 65(1):50–70.
- Blattberg RC, Neslin SA (1990) *Sales promotion: Concepts, methods and strategies* (Prentice Hall).
- Cohen MC, Gupta S, Kalas J, Perakis G (2020a) An efficient algorithm for dynamic pricing using a graphical representation. *Production and Operations Management* Published online in Articles in Advance.
- Cohen MC, Kalas J, Perakis G (2020b) Promotion optimization for multiple items in supermarkets. *Management Science* Published online in Articles in Advance.
- 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.
- 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.