



Dynamic effects of store promotions on purchase conversion: Expanding technology applications with innovative analytics[☆]

Leonardo D. Epstein^{a,*}, Ignacio E. Inostroza-Quezada^{a,b}, Ronald C. Goodstein^c, S. Chan Choi^b

^a School of Business and Economics, Universidad de los Andes, Santiago, Chile

^b Rutgers Business School, Newark, NJ, USA

^c Georgetown University McDonough School of Business, USA

ARTICLE INFO

Keywords:

Retail technology

Conversion

Promotional effectiveness

Counterfactual prediction

ABSTRACT

Retail technology provides novel ways to investigate the effects of marketing actions (MAs) on consumers' behaviors. To evaluate effects of MAs on conversion, we propose an approach that combines sensor-collected records of store entries with register-receipts within pre-determined time-bands. Studies of effects of MAs on conversion in brick-and-mortar settings, however, are surprisingly scarce. The approach has three stages: (a) build a model for the conversion probabilities with data outside the MA interval, (b) build a counterfactual baseline for the conversion probabilities during the MA, under the hypothesis that this action has not occurred, (c) compare the observed conversions to the corresponding counterfactual conversions to determine the effects of the MA. An analysis of a two-day promotion illustrates the method. It provides evidence that the promotion associates with increased counts of both visits and purchases, but the association with conversion remains unclear.

1. Introduction

A recent report by McKinsey & Company states that “Technology is rewriting operating models across every sector of the economy, but nowhere do we feel this more in our day-to-day lives than in retail” Balchandani et al. (2020, p. 1). For instance, Macy's newest mission statement specifically incorporates technology into their positioning: “Our goal is to be a retailer with the ability to see opportunity on the horizon and have a clear path for capitalizing on it. To do so, we are moving faster than ever before, employing more technology and concentrating our resources on those elements most important to our core customers”.¹ This retail technological revolution is the focus of special issues of several marketing/retailing journals around the globe, such as the Journal of the Academy of Marketing Science (2020), the International Journal of Electronic Commerce (2014), the Baltic Journal of Management (2018) and this one, among others. Several recent, seminal articles review and closely examine this topic (Grewal et al., 2018; Inman & Nikolova, 2017; Shankar et al., 2021). While most of the articles focus upon how technology will shape shoppers' and retailers' behaviors, they devote less attention to the manner in which technology can be used to analyze the effectiveness of marketing activities within a retail store.

Research analytics is vital for retailers to assess the success of their strategies and tactics, as well as for researchers to develop new models to explain consumers' behaviors. The purpose of our paper is to use retail technology to garner better insights into the effectiveness of marketing activities, for example store promotions, on conversion rates. Thus, our focus is not on the technology itself but on the utilization of that technology to improve analytics. Specifically, by utilizing in-store traffic sensors and electronically produced register receipts we propose and illustrate the use of a novel model for computing the predictive effects of promotional activities on conversion probabilities.

2. Relevant literature

2.1. Using technology to measure conversion rates

The marriage of technology and retailing actually began with the adoption of UPC (Universal Product Codes) scanner pricing systems in the early 1970s (Goodstein, 1994). This was the first time that technology promised, and delivered, to retailers a way for electronic machinery to bring improvements to the bottom line. Retailers were promised hard savings of 1%–2% of expenses from labor savings associated with item pricing, improved checkout efficiency, and lessening of other manual

[☆] The authors are grateful to FollowUp-Chile for providing the data and helpful discussions.

* Corresponding author.

E-mail addresses: lepstein@uandes.cl (L.D. Epstein), IInostroza@uandes.cl (I.E. Inostroza-Quezada), goodster@georgetown.edu (R.C. Goodstein), chanchoi@busines.rutgers.edu (S.C. Choi).

¹ <https://www.comparably.com/companies/macys/mission>

costs (Cutler & Rowe, 1990; Goodstein, 1994; Hardesty et al., 2014). In addition, they were afforded with soft savings of another 2%–3% because of the ability of managers to analyze the data from these systems to lower stockout rates, to implement automatic reordering, and to determine optimal locations and promotional effects within the store (Cutler & Rowe, 1990; Goodstein, 1994). Consumers also found these systems quite interesting and attractive because they made the shopping trip faster and more efficient.

For researchers, scanner data produced from UPC pricing systems are a virtual goldmine with scads of publications utilizing these data (see Bucklin & Gupta, 1999 for a review). As a result, scanner technology has improved not only managerial practice, but also insights based on researchers' mining of these data to test and develop new theories and models. We suggest that newer technologies offer some of these same benefits in terms of seeding research analytics. Such has been the case with the progression of technology adoption by retailers to date, including electronic shelving systems (Clodfelter, 2004; Goodstein & Escalas, 1995), video tracking (Epstein et al., 2016; Zhang et al., 2018), ambient music systems (Baker et al., 1992; Biswas et al., 2019), and sensors (Hui et al., 2009) to name just a few. This is not to mention the more sophisticated internet based technologies and artificial intelligence systems that are predicted to change the future face of retailing both managerially and in terms of research opportunities (Grewal et al., 2018; Inman & Nikolova, 2017; Shankar et al., 2021).

To determine conversion rates, analysts must be able to assess the number of people arriving into a store (traffic) during prescribed time bands, and the number of these prospects that actually purchase items, to wit, convert. Zhang et al. (2018) establish these numbers with video technology to assess traffic and point-of-sale (POS) transaction log data to capture purchases. Their findings suggest that social interactions have the ability to affect conversion rates. The authors suggest, however, that privacy concerns limited their ability to understand the exact content of the social interactions they study. It is also the case that video technology could raise privacy concerns for consumers (Inman & Nikolova, 2017). In their review of shopper-facing retail technology, Inman and Nikolova (2017) review how infrared sensors can be used to count the number of shoppers in a store. There are several companies that offer such technology and the authors mention that some retailers are combining sensor counts with POS data to determine conversion rates (p. 10). We were afforded the opportunity to do just that, to utilize a retailer's sensor counts and POS data to test a new model for assessing conversion rates.

2.2. The importance of conversion

Retailers use marketing activities, such as promotions, to improve store performance. Lam et al. (2001) suggest that store sales performance may be decomposed into three sources: arrival traffic or the number of store visitors in time-bands, the conversion rate of those visitors, and the monetary value of each sale. That is, a manager enacts a promotional campaign to attempt to increase at least one of the following: the traffic, the conversion rate, or the dollar value of sales (Epstein et al., 2016; Lam et al., 2001). An important question then is to examine the effect of promotions on each component of this model. While several important papers examine the cumulative effects of retail promotions on total sales (Gupta, 1988; Van Heerde et al., 2000; Van Heerde & Neslin, 2017) the models in these articles do not allow the analyst to identify the reason these effects occur; namely one can attribute them to an increase in store traffic, to an increase in the conversion rates of customers visiting the store, or to an increase in the monetary value of each sale. In an initial attempt to bring clarity to this issue, Epstein et al. (2016) develop a model that assesses the effect of promotions on foot traffic and find in their data that promotions meaningfully affect store visits. Similarly, Gauri et al. (2017) examine the impact of multiple features of promotions on traffic and sales. They find that feature promotions increase traffic in general and even more

so for high frequency, high penetration product categories, and for branded goods. Neither study, however, directly measures conversion rates as opposed to traffic or sales. Retailers also want to know whether visitors become buyers, which means that conversion rates are important. Prediction of conversion rates is precisely the focus of the current research.

Our research is not the first to discuss the importance of conversion rates to marketers and retailers. For instance, Hanssens et al. (2014) report that conversion is an important behavioral outcome for practitioners and explore how attitudinal measures might be used to predict conversion probabilities. In a more recent paper, Motyka et al. (2016) explore the effects of processing fluency on conversion rates. They find that altering the fonts in announcements of promotions can cause disfluency in processing and that, surprisingly, this disfluency can have positive effects on conversion rates. Conversion research also appears in related fields such as operations management to explore how analysis of foot traffic and conversion rates can be used to properly plan for staffing in a retail environment. Specifically, Perdikaki et al. (2012) examine the relationship between traffic and conversion. They find that increases in traffic can lead to decreases in conversion rates, and that these smaller conversion rates may lead to smaller future traffic. The authors suggest that improving customer service through increased staffing during peak shopping periods reduces traffic perceptions and increases conversion. Staffing is a very important component of retail atmospherics, a topic covered elsewhere in the retailing literature (Kim & Kim, 2012; Puccinelli et al., 2009).

3. The model

3.1. Overview

In a retail store visitors follow a sequence of stages that eventually leads a portion of them to make a purchase before they leave. Typically, this customer journey includes stages such as a store window catching the eye of a passer-by, browsing the window, and deciding whether to enter the store (Grewal & Roggeveen, 2020). If a person enters the store, stages such as trying an item or requesting information from a sales person may follow. A visitor may complete the visit with a purchase. The observed conversion fraction is the proportion of visitors who make a purchase to the number of visitors.

In the marketing literature conversion models are mostly associated with e-commerce (De Haan et al., 2016; Li & Kannan, 2014; Moe & Fader, 2004; Park, 2017; Xu et al., 2014). However, research on conversion models in the context of brick-and-mortar stores has been limited. First, in terms of behavior of arrivals and purchases, online and brick-and-mortar shopping differ markedly due to hours of operation. Online shopping is available 24/7/365, but brick-and-mortar shopping is available only during business hours. Second, browsing at a brick-and-mortar store has a much higher cost than browsing online, and hence the number of arrivals to online stores is typically larger than to brick-and-mortar stores. Third, space constraints in brick-and-mortar stores may also lead to discomforting situations such as crowding. Fourth, beyond the differences between online and brick-and-mortar shopping environments, conversion figures differ markedly between them. Indeed, conversion fractions are smaller in online stores where they seldom exceed 5% (Moe & Fader, 2004), than in brick-and-mortar stores where they are in the range 10–20% (Matzler et al., 2010).

One of the earliest articles to model conversion in brick-and-mortar stores is Lam et al. (2001). They use log-normal approximations for the distributions of the number of arrivals and of purchases per hour. Their use of standard regression models requires additional assumptions, namely independent observations with constant variance. The analysis of Lam et al. (2001) requires two pieces of data, the number of visitors entering the store and receipts within each time band (TB) to indicate conversion of traffic into sales. Scanners or registers record the purchase process automatically (Goodstein, 1994). Thus, counting the

number of visits in a store seems to be the main obstacle to estimate conversion. In recent years imaging technologies have made inroads into the retail industry to provide accurate sampling of stores' arrival or exit processes, thus making it possible to determine the number of store visits in TBs. Our approach combines data from sensors placed at store entrances and data from cashier registers to compute conversion. This novel use of these technologies offers promise for analyzing not only traffic and conversion, but could be used to measure attention within the store as well (Grewal et al., 2020).

Abraham and Lodish (1987) propose to evaluate the effect of promotions on sales against a baseline estimate. For this baseline they suggest to “estimate a baseline of what sales would have been had a promotion not been run” (p. 101). Epstein et al. (2016) recently combined this idea with formal statistical methods to make detailed evaluations of the effects of promotions on arrival traffic. Thus, the core of both Abraham and Lodish (1987) and of our paper is the “conditional on the counterfactual hypothesis” baseline prediction (CCFH) that no promotion had occurred for the conversion process during the promotion interval. The CCFH baseline consists of conversion predictions during the promotion interval.

To compute this CCFH baseline our approach fits a Beta–Binomial model with the visitors and the converters counts in each TB outside the promotion interval. Thus, the CCFH baseline probabilities represent conversion values one would expect to observe if the promotion had not occurred. Then the method compares conversion proportions actually observed during the promotion to the CCFH baseline conversion probabilities. Specification of a metric for these comparisons requires additional notation. Let (t, b) designate time band b of day t . Call $p(t, b)$, $t = 1, \dots, T$, $b = 1, \dots, B$, the conversion probabilities during the study. The baseline probabilities during the promotion interval are $\hat{p}(t, b)$, $b = 1, \dots, B$, with $t = T_S, \dots, T_E$, where T_S (resp. T_E) is the day-index for the promotion start (resp. end), with $T_S \leq T_E$. Call $p_O(t, b)$, $t = 1, \dots, T$, $b = 1, \dots, B$, the observed conversion proportions. For the sake of continuity we defer to Section 3.2 the formulas for $p(t, b)$, and $p_O(t, b)$, $t = 1, \dots, T$, $b = 1, \dots, B$.

With the observed and predicted baseline probabilities $\hat{p}(t, b)$, we proceed with the prediction of the difference $\text{logit } p_O(t, b) - \text{logit } \hat{p}(t, b)$, also called log odds-ratio,

$$\ln \widehat{OR}(t, b) = \ln \left\{ [p_O(t, b)/(1 - p_O(t, b))] / [\hat{p}(t, b)/(1 - \hat{p}(t, b))] \right\}. \quad (1)$$

We use the term “prediction” of conversion probabilities instead of “estimation” to stress that their associated measure of uncertainty must compound the sampling variability of model estimators with the distribution of observations, counts of purchase receipts, conditional on the parameter estimates. This procedure agrees with the standard rationale to use prediction intervals in linear models. To compute the prediction distribution associated with the CCFH baseline, this article uses a simple implementation of the bootstrap scheme. The method can be modified to incorporate additional stages of the customer journey (Grewal & Roggeveen, 2020).

The reason the approach uses the prediction distribution for the baseline, but it does not use the prediction distribution of the observed conversion rests on two points. First, after the promotion, the baseline has not been observed and hence is subject to uncertainty. The appropriate representation of this uncertainty is its prediction distribution. Second, after the promotion, the actual TB-specific conversions have been observed and hence, at that point they are not subject to uncertainty.

3.2. Models and baseline predictions

The approach, with adaptations, can be used to evaluate the effects of interventions, marketing actions and promotions in particular, and in settings where a control group is unavailable, including the evaluation of the effects on sales or foot traffic and sales. Broadly stated, our approach uses observations outside the intervention interval to

build a CCFH prediction baseline. After the intervention, the approach compares the observations during the intervention with the CCFH predictions. In particular, we are interested in measuring the effects of a promotion on conversion probabilities.

For each day t , $t = 1, \dots, T$, the observations are TB-specific counts of converting store visitors, that is, visitors who make a purchase before leaving the store. Each day has B TBs, and b , $b = 1, \dots, B$, indexes them. Let the random variable $V(t, b)$ be the number converters and the random variable $N(t, b)$ be the number of visitors in the store during TB (t, b) , $t = 1, \dots, T$, $b = 1, \dots, B$. Let $v(t, b)$ and $n(t, b)$ be the respective observed values.

We build a statistical model for the number of visitors who convert, conditionally on the number of visitors in the store. As is the case with binomial data, the conditional distribution of conversion counts depends on the probabilities of conversion in each time band, which are precisely the parameters of interest. The standard model assumes (a) the conditional conversion counts are binomial, that is $[V(t, b)|N(t, b) = n] \sim \text{Bin}(n, p(t, b))$, (b) conditionally on $\{N(t, b), t = 1, \dots, T, b = 1, \dots, B\}$, the conversion counts $\{V(t, b), t = 1, \dots, T, b = 1, \dots, B\}$ are independent, and (c) a logistic regression function parametrizes the conversion probability $p(t, b)$,

$$\text{logit } p(t, b) = \beta_0 + \mathbf{x}_p(t, b)' \beta_1, \quad (2)$$

which models the effects of covariates in the row-vector $\mathbf{x}_p(t, b)'$, for instance TB, or day of the week.

These assumptions may be unduly simplistic and if so, lead to distorted inferences for the promotion effects. In particular, the binomial assumption (a) in the paragraph above precludes modeling extra-binomial dispersion in terms of covariates. This limitation may be overcome by fitting Beta–Binomial models (Rigby & Stasinopoulos, 2005). The R-function `gamlss()` (Rigby & Stasinopoulos, 2005; Stasinopoulos & Rigby, 2007) uses a linear parametrization (2) for $\text{logit } p(t, b)$, and a separate linear parametrization for $\ln \sigma(t, b)$,

$$\ln \sigma(t, b) = \gamma_0 + \mathbf{x}_\sigma(t, b)' \gamma_1, \quad (3)$$

where $\sigma(t, b)$ is the extra-binomial dispersion parameter for the count $V(t, b)$. Thus the `gamlss` framework suits well the analysis of conversion probabilities with TB-specific counts of buyers.

Our approach estimates the parameters in (2) and (3) only with data outside the promotion interval. Call $(\hat{\beta}_0, \hat{\beta}_1)$ and $(\hat{\gamma}_0, \hat{\gamma}_1)$ the parameter estimates in the model. Call $\hat{p}(t, b)$, $t = 1, \dots, T$, $b = 1, \dots, B$, the predicted probabilities corresponding to the estimates $(\hat{\beta}_0, \hat{\beta}_1)$ and $(\hat{\gamma}_0, \hat{\gamma}_1)$. The predicted conversion probabilities $\hat{p}(t, b)$, during the promotion, namely $t = T_S, \dots, T_E$, $b = 1, \dots, B$, provide the CCFH baseline. For the promotion period, these are the conversion probabilities one would have expected to observe had this promotion not taken place. Next, the approach compares the baseline predicted probabilities with the observed proportions during the promotion. The TB-specific conversion proportions are $p_O(t, b) = v(t, b)/n(t, b)$, $t = 1, \dots, T$, $b = 1, \dots, B$.

4. The effect of a promotion on conversion: an illustrative example

4.1. The example

The example uses proprietary data from a clothing retail store. The data span a 6 month interval from day $t = 1$ to $T = 182$ (May 1 to October 29, 2016). The analysis evaluates the effect of a two-day promotion that starts on $T_S = 181$ and ends on $T_E = 182$. The promotion consists of a 15% discount on two clothing brands. A confidentiality agreement precludes us from disclosing additional information.

Fig. 1 shows the data collected. The number of TBs in the data set is $N_{TB} = 1831$. There are 172 days that contribute observations. Some days contribute fewer than $B = 11$ TBs (29 days have $B = 10$ and 32 days have $B = 9$). The light (dark) line represents the counts of visitors in the store (converters) by TB. Exploratory analyses suggest




Fig. 1. Number of store visits (dark line) and number of receipts (light gray line) by time-band. The data set starts on May 1, 2016, and ends on October 29, 2016. The two-day promotion takes place on October 28 and 29. The vertical gray band between July 27 and August 3 identifies an interval where the number of visits are persistently small. FollowUp analysts confirm that between during this interval sensors malfunctioned.

that counts of visitors $n(t, b)$, represented with the light gray lines, from $t = 88$ through $t = 95$, are unusually small compared with counts just to the left or to the right of this interval. The plot in Fig. 1 highlights this interval with a vertical gray band. This observation corresponds to the pattern in Fig. 2, top panel, where the gray dots, representing observed logits of conversion values in this interval, appear upward-shifted relative to the gray dots just to the left or to the right of this interval. FollowUP Customer Experience analysis confirm that sensors malfunctioned during this interval. For this reason our analysis omits data in this interval. Fig. 2, top panel, shows the observed logits for the example data $p_O(t, b)$. It also shows the estimated values $\hat{p}(t, b)$, prior to the promotion (i.e., $t = 1, \dots, T_S - 1$, $b = 1, \dots, B$, and the prediction for the promotion interval $b = 1, \dots, B$, with $t = T_S, \dots, T_E$. The bottom panel of Fig. 2 zooms in the data and the predictions for the two days of the promotion.

However, other separation metrics between $p_O(t, b)$ and $\hat{p}(t, b)$, may also be used. For instance, the relative risk $p_O(t, b)/\hat{p}(t, b)$, or the odds-ratio itself $\hat{O}(t, b)$ can measure the shift of the (transformed) observed proportions relative to the CCFH conversion probabilities. The solid line in Fig. 2, bottom panel, shows $\logit \hat{p}(t, b)$, $t = 1, \dots, T$, $b = 1, \dots, B$, the series of predicted conversion probabilities in the logit scale. These probabilities are the basis to evaluate the effect of the promotion on conversion. Since their computation uses the fitted model, we discuss first the model building process and then we return to the use and interpretation of the predicted baseline (Fig. 5) and predicted effects (Fig. 6).

In this example, the data for the analysis are rows $[(t, b), v(t, b), x_p(t, b), x_\sigma(t, b), n(t, b)]$, $t = 1, \dots, T$, $b = 1, \dots, B$. Our analysis starts off with a model that includes main effects, and two-factor interactions and autoregressive terms, as well as simple functions of t , the day-index that accounts for trends.

The autoregressive terms are of the form $\logit p_O(t, b-l)$, where $l > 0$ is an integer lag. Inclusion of same day autoregressive terms entails losing the first l observations of each day. Including autoregressive terms in the logit scale has the purpose of matching the scales of the left and the right-hand sides in Eq. (2), but model fitting diagnostics should guide the choice of the scale that one applies to $p_O(t, b-l)$. The inclusion of terms in the final model follows the standard model building process with likelihood-ratio tests.

A description of these factors is as follows. The first level of each factor serves as a baseline which R sets to zero.

- $TB(\cdot)$ is the factor that represents the TB of store operation, with levels $b = 1, \dots, B$, and $B = 11$ in this example. Level $b = 1$ represents the baseline 1-h time interval $[10 : 00; 11 : 00]$, and the following levels represent subsequent 1-h time intervals. The right-extreme bracket “ $]$ ” of the interval $[10 : 00; 11 : 00]$ indicates that it excludes the upper-extreme point, namely $11 : 00$. The last level, $b = 11$, represents the 1-h time interval of operation $[21 : 00; 22 : 00]$.
- $D(\cdot)$ is the factor with levels $d = 1, \dots, 7$, that represents the successive days of the week Monday through Sunday.
- $W(\cdot)$ is the factor with levels $w = 1, \dots, 6$, that represents the successive weeks of a month (October 2016 has six weeks).
- $M(\cdot)$ is the factor with levels $m = 5, \dots, 10$, that represents the successive months May through October.

In this model the levels b, d, w , and m of their respective factors depend on the day index t , $t = 1, \dots, T$. Thus, the factors-levels corresponding to a count $v_{(t,b)}$ are $TB(b(t))$, $D(d(t))$, $W(w(t))$, and $M(m(t))$. To simplify the notation, we omit the argument t in the factor-level in the following. Thus, we use d , instead of $d(t)$, for the “day of the week level” that corresponds to day t . The notation makes analogous omissions for the other factors.

We assess significance of factors and covariates with standard likelihood-ratio tests. The final model has, (a) a parametrization for $\logit p(t, b)$ with main effects for factors TB , D , W , M , and no interactions, and (b) a parametrization for $\ln \sigma(t, b)$ with factor D only. The final model has 34 parameters, 27 for the mean (1 for the intercept, 10, 6, 5, and 5 for factors TB , D , W and M , respectively), and 7 additional parameters in the log σ -parametrization, 1 for the intercept and 6 for the only factor D .

The illustration assumes that receipt counts are independent. Model diagnostics do not provide evidence contradicting this assumption. Indeed, the fitted model has quantile residuals (Rigby & Stasinopoulos, 2005) with marginally significant autocorrelations and partial autocorrelations only for lags 3 and 10, the same lags for each autocorrelation type. The analysis proceeds with the working assumption that conditional on covariate-values, the TB-specific receipt counts are independent.

Next we examine the effects of factors on conversion during the pre-promotion interval. The plots in Fig. 3 use a model fit with pre-promotion data only. Hence, these plots inform about effects of factors TB , W and M , on baseline conversion in the logit scale. These interpretations are not the main goal of this article, but may be useful




Fig. 2. Top panel: Observed conversion fractions (gray dots) and estimated conversion probabilities with pre-promotion data only (light gray line) during the pre-promotion period. The two-day promotion interval is at the extreme right. In this period predicted conversion probabilities (black line) represent likely conversion values one would have observed had the promotion not taken place. These serve as baseline values to evaluate the promotion effect. The vertical gray band between July 27 and August 3 marks the interval where the sensors malfunction. To ease the examination of the plot solid vertical segments indicate the first day of each month and broken vertical segments indicate the first TB of each Monday. Since the logits of 0 or 1 are not defined, we have changed the values of such observed proportions $p_O(t, b) = v(t, b)/n(t, b)$ as follows: (a) for a TB (t, b) with $v(t, b) = 0$, the plot uses the value $p(t, b) = (0+1)/(n(t, b)+1)$. The plot identifies these TBs with a “0” on the horizontal line with ordinate 0. (b) for a TB (t, b) with $v(t, b) = n(t, b)$, the plot uses the value $p(t, b) = v(t, b)/(n(t, b)+1)$. The plot identifies these TBs with a “+” in the horizontal line with ordinate 0. Bottom panel: zooms in the data and the predictions during the promotion.

to managers. The additive parametrization, with no interactions, facilitates interpretations (see Table 1). If there were interactions, some panels would include one curve for combinations of levels of the factors in each interaction. These plots use the logit scale, but they translate readily into probabilities. In Fig. 3, the top-left panel shows the logits of conversion versus TB. Time band [10 : 00; 11 : 00] is the baseline for this factor, set at zero, and does not appear in this plot. The logit is positive, approximately constant at the level 0.3 in the first six TBs of the panel, but then increases monotonically toward the closing hour. Thus, within a day, conversion is the smallest in the first TB.

The top right panel shows the effect of day of the week where Monday is the baseline. The effect of Tuesday is positive, indicating that one expects conversion to be higher on Tuesdays than on Mondays. On Wednesdays the effect is zero, indicating that one expects conversion to be approximately the same as on Mondays, but smaller than on Tuesdays. During the following days of the week conversion appears

to decrease. Sunday is the day with smallest conversion. Interpretation of the next two plots proceeds in a similar fashion. In these plots it is advisable to report confidence intervals with center at the corresponding effects.

Fig. 4, shows the effects of factor D on $\ln \sigma$, where σ represents the extra-binomial dispersion. The only statistically significant factor in this parametrization is D . Monday is the baseline day, and the intercept estimate is -4.5465 , thus $\hat{\sigma}(\text{Monday}) = 0.01$, a small but statistically significant value. This result indicates that the counts of receipts issued on Mondays are nearly binomial, but with significant extra-binomial over-dispersion. The extra-binomial dispersions on Tuesdays, Wednesdays and Thursdays are even smaller, and for Fridays and Saturdays these variabilities increase to the levels of Tuesdays and decrease for Sundays. Only the effect of Friday is statistically significant, indicating that the extra-binomial variability of conversion counts on Fridays is slightly, but statistically significantly, smaller than the extra-binomial




Fig. 3. Estimated effects on the logits of conversion by factor. Top row: Factors are Time-band (left panel) and Day of the week (right panel). Bottom row: Factor are Week of the month (left panel) and Month of the year (right panel). Vertical segments are 95% point-wise confidence intervals.

variability for Mondays. An overall likelihood ratio statistic to compare the fit of models with and without factor D has p -value = 0.043. Prior to the analysis, it is unclear whether the binomial distribution is appropriate to model the conversion counts. Only after this analysis with Beta-Binomial distributions, one concludes that extra-binomial dispersion may be negligible and that an analysis with binomial distributions would probably lead to similar inferences. This analysis highlights one advantage of Beta-Binomial distributions: they allow the analyst to assess whether extra-binomial dispersion affects inferences.

4.2. Illustrative example: promotion effects on conversion

Our approach measures the effects of a promotion on conversion comparing observed conversions with CCFH conversions. Our illustration uses the logit metric for these comparisons. More specifically, we use the difference $\text{logit } p_O(t, b) - \text{logit } \hat{p}(t, b)$, for TBs (t, b) in the promotion period. In this difference $p_O(t, b)$ is the observed conversion fraction and $\hat{p}(t, b)$ is the point-wise CCFH prediction for TB (t, b) . We note in passing that these differences are the standard log odds-ratios $\ln \text{OR}(t, b) = \text{logit } p_O(t, b) - \text{logit } \hat{p}(t, b)$, of common use in logistic regression, although we stress that our inferences use `gamlss`, not logistic regression.

Fig. 5 shows predictions (broken line) with 95% point-wise prediction bands (shaded with gray) with the observed conversion fractions




Fig. 4. Estimated effects days of the week on the log extra-binomial dispersion. Vertical segments are 95% point-wise confidence intervals.

Table 1

Output from gamlss fit with pre-promotion data only.

Inferences: parameters for logit conversion				
	Estimate	Std. Error	t-value	p-value
Intercept	−2.0025358	0.0730375	−27.418	$<2 \cdot 10^{-16}$
BH2	0.2250159	0.0627450	3.586	0.000344
BH3	0.2854843	0.0609075	4.687	$2.98 \cdot 10^{-6}$
BH4	0.2721667	0.0601282	4.526	$6.39 \cdot 10^{-6}$
BH5	0.2911449	0.0603335	4.826	$1.51 \cdot 10^{-6}$
BH6	0.2080906	0.0594591	3.500	0.000477
BH7	0.2920950	0.0587462	4.972	$7.25 \cdot 10^{-7}$
BH8	0.4017690	0.0580729	6.918	$6.32 \cdot 10^{-12}$
BH9	0.4093332	0.0583731	7.012	$3.3 \cdot 10^{-12}$
BH10	0.6623811	0.0598910	11.060	$<2 \cdot 10^{-16}$
BH11	1.0855990	0.0796435	13.631	$<2 \cdot 10^{-16}$
D2	0.0477564	0.0444569	1.074	0.282869
D3	−0.0006289	0.0438642	−0.014	0.988563
D4	−0.0103384	0.0424837	−0.243	0.807762
D5	−0.0209926	0.0424510	−0.495	0.621004
D6	−0.0075560	0.0415674	−0.182	0.855778
D7	−0.1369125	0.0461090	−2.969	0.003024
S2	0.0527884	0.0356543	1.481	0.138898
S3	0.0172916	0.0355844	0.486	0.627074
S4	−0.0725118	0.0366553	−1.978	0.048057
S5	0.0183353	0.0428525	0.428	0.668797
S6	−0.2300774	0.1330510	−1.729	0.083937
M6	−0.0374390	0.0341398	−1.097	0.272945
M7	0.0965486	0.0337993	2.857	0.004332
M8	−0.2131029	0.0376631	−5.658	$1.78 \cdot 10^{-8}$
M9	−0.0345496	0.0354935	−0.973	0.330483
M10	−0.0491174	0.0368105	−1.334	0.182264

Inferences: parameters for the log extra-binomial dispersion				
	Estimate	Std. Error	t-value	p-value
Intercept	−4.5465	0.2945	−15.441	$<2 \cdot 10^{-16}$
D2	−0.9454	0.6557	−1.442	0.1496
D3	−1.5854	1.1406	−1.390	0.1647
D4	−2.3064	1.4451	−1.596	0.1106
D5	−1.3920	0.6506	−2.140	0.0325
D6	−0.8632	0.4344	−1.987	0.0470
D7	−1.3285	0.7930	−1.675	0.0940

$p_O(t, b)$ (solid lines). The plots in this figure use the 11 TBs in each day of the two-day promotion. The solid and broken lines are in the scale of conversion, the quantity of interest, which has a direct interpretation in the marketing application.

Fig. 6 shows the vertical separation in the logit scale, between the solid line and its corresponding broken line in the plots of Fig. 5. This separation in the logit scale represents TB-specific effects. The gray bands represent the variability of the effects that compounds the two sources: sampling variability of estimates and the conditional variability of number of converters given the values of the parameter estimates.

Figs. 5 and 6 are helpful to inform managers the evolution of the promotion effect. These analyses could help adapting, modifying, and perhaps interrupting or extending, a promotion according the evolution of the effects. In particular, when planning a promotion managers can use predicted baselines to set conversion goals. These displays could be useful to identify time bands, or combinations of factor levels, where conversions is small and offer sales potential, for instance when small conversions co-occur with large traffic. After, or during a promotion managers and analysts can monitor the evolution of conversion and possibly decide on implementing changes to improve performance.

4.3. Illustrative example: discussion

In related settings, measuring effects of marketing actions on arrival traffic, earlier studies propose methods that model the observations during the promotion interval (Ando, 2008; Fok et al., 2012). One may be tempted to adapt these methods, namely those that model

the data during the promotion, to evaluate effects on conversion. We caution against adopting such an approach because such models have at least two limitations that are difficult to overcome. First, usually the promotion interval is much shorter than the pre-promotion interval and consequently has many fewer observations to model complex effects during the promotion. Second, the promotion introduces a sudden distribution change that is similar to a shock or a perturbation. Thus, the observations exhibit erratic patterns, usually a combination of level and dispersion changes, that make it very difficult to model, particularly with short promotions.

Sound evaluations of promotion effects require analysts to be mindful of two issues: one needs to avoid omitting important covariates such as holidays and promotions are assumed to be independent from past data, such as sales slowdowns. Indeed, our approach predicts a counterfactual baseline on the basis of the conversion history. Thus, this baseline is appropriate as long as there is no association between the promotion and this baseline. More concretely, if management sees an unexpected, below baseline slump either in conversion or other performance metrics, and at that moment decides to start a promotion to under grid performance, then the historic baseline would not be appropriate; rather the analysis should use a locally contemporaneous measurement.

To take advantage of visitor and conversion counts in small TBs, we use covariate-factors to represent them, in addition to differential effects for day of the week, week of the month, and month, and possibly interactions among these and other factors or covariates. We first build a baseline model and then extend it to incorporate TB-specific predicted effects. This is an important advantage of our approach as it allows managers to identify combinations of TBs, days, weeks, and months, when conversion is exceptionally low. This identification is useful to guide marketing actions to increase conversion probabilities, particularly in intervals with large potential, such as when arrival traffic is large.

The analysis in Lam et al. (2001) shows that price promotions positively affect both store entry and the likelihood that a consumer will make a purchase. In contrast, in our illustration the data show a marked increase of arrival traffic, but our analysis shows no evidence that the promotion has an effect on conversion probabilities. Once the promotion starts, both the number of visitors and the number of receipts increase and this approximately monotone relationship holds for the entire study. This contrast between the correlation between traffic and receipts versus traffic and conversions illustrates an interesting instance where both counts of visitors and receipts increase but the conversion probability decreases. This observation shows the importance of conducting examinations of both of these performance measures.

Joint analyses of conversion probabilities and visitor count data opens new research problems and new opportunities to improve promotional profitability. For example, plots such as those in Fig. 3 allow managers to identify combinations of factor levels with small conversion probabilities. With the same factors, one can build similar models and plots for TB-specific visitor counts to identify combinations of factor levels with large expected visitor counts. Factor combinations with small conversion probabilities and large expected visitor counts have a potential for a large profit increase, by identifying promotion activities that increase conversion a manager can leverage large visitor counts to increase the counts of receipts.

5. General discussion

5.1. Advantages and implications

This research contributes to the literature on the effects of promotions on conversion by developing and illustrating a predictive model that allows one to develop a counterfactual baseline of conversion to compare observed promotional conversion to the counterfactual




Fig. 5. Bootstrap replicates of prediction conversion probabilities along time bands (500 curves). The left panel shows the replicates of curves for the first promotion day, October 28. The right panel shows the replicates of curves for the second promotion day, October 29. The solid lines are observed conversion fractions. The dotted lines are point-wise baseline predictions. The region with gray lines mark the bootstrap replicates.




Fig. 6. Bootstrap replicates of conversion effects along time bands (logit scale) relative to observed conversion fractions (logit scale, 500 curves). The broken lines are the point-wise baseline logit predictions minus the observed logit conversion fractions. The left panel shows the replicates for the first promotion day, October 28. The right panel shows the replicates for the second promotion day, October 29.

conversion. Our approach has several advantages worth highlighting. First, it takes advantage of time-band specific counts of visitors and receipts, obtained through sensor technology and registers, respectively. With the available data, we build a model that incorporates covariate information, such as date, time, and additional environment covariates. Were the data available, the model could also incorporate marketing actions of competitors or “complementors” that is, stores selling complementary items. Second, our approach unlike others in the literature, does not require modeling the conversion data during the promotion (Ando, 2008; De Haan et al., 2016; Lam et al., 2001). In fact, the approach uses the data observed during the promotion to compute standard conversion proportions, and it computes the counterfactual predictions with a model that uses only the observations outside the promotion. Most often, the interval outside the promotion has a much larger number of observations than the promotion interval, enabling analysts to build solid counterfactual predictions for the shorter promotional period. Third, our approach measures the effects of the promotion on the scale of the quantity of interest, namely the conversion probability or an equivalent metric such as the logit of conversion. Thus, in the illustration for instance, the approach captures the effect of conversion with time-band specific differences between observed and predicted measures of conversion. These differences are

more informative and easier to interpret than a parameter coefficient associated with a promotion indicator covariate. It is worth noting that prior studies on the effects of promotions, either on traffic or on sales, use hypothesis testing to decide on the statistical significance of a promotion (Ando, 2008; De Haan et al., 2016; Lam et al., 2001).

Unlike these other approaches, our methodology compares the evolution of observed TB-conversions to the corresponding TB-predicted conversion in the promotional interval. These bands allow the analyst to identify time-bands along the promotion where the observed conversion is outside the prediction bands, representing deviations likely associated with the promotion. This comparison is more informative than inferences about a promotion-effect parameter, a summary about the overall effect of the promotion. The hypothesis-based research summarizes the promotion effect with a statistical test (Ando, 2008; De Haan et al., 2016; Lam et al., 2001). However, although a test may declare an effect statistically significant, its size may not affect the business in a meaningful manner.

This study extends prior work discussing the importance of conversion rates to marketers and retailers (Hanssens et al., 2014; Motyka et al., 2016). Though the outcomes associated with the illustration are specific to the data provided by the retailer, there are several findings that are worthy of discussion. For example, the evidence here supports

prior research that finds that although promotions may have a significant effect in increasing traffic (Epstein et al., 2016), they may also be associated with a decrease in conversion (Perdikaki et al., 2012). Of course, additional data across a variety of retailers and product categories would allow one to draw more general conclusions about this relationship (Gauri et al., 2017). However, the fact that traffic, conversion, and sales do not always move in the same direction speaks to the importance of measuring each of these behavioral outcomes in assessing the effectiveness of marketing actions (Lam et al., 2001). One could argue that as long as sales increase at the bottom line, then the behavioral path leading to those increases is trivial. That would be an unfortunate argument because specific marketing actions may be directed to affect different behaviors or processes during the customer journey (Grewal & Roggeveen, 2020).

There is evidence that promotions may affect sales beyond the end of the promotional period (Doyle & Saunders, 1985, 1986). For instance, prospective buyers accelerating purchases to buy on deal may lead to post-promotion drops in sales. The converse effect, namely post-promotion sales increases may take place if the promotion leads to increased product awareness. On the other hand, effects may occur before a promotion, for instance when prospective buyers postpone purchases to the promotion interval or when supply is limited, which leads to demand increases before the promotion from customers who attempt to avoid a stock-out. Time shifts of customer behavior may affect conversion probabilities outside the promotion intervals, and therefore change the counterfactual model. Pre and post-promotion effects on conversion probabilities have been barely studied. Our method may assist researchers to develop models for these effects.

When is it worthwhile to invest in a promotion to increase conversion probabilities? This is a question that managers and researchers should discuss on a case-by-case basis. Often, strategic issues weigh on the answer together with considerations of profitability. There are instances when a conversion would be worth the investment in a promotion: (a) when the retailer is interested in increasing profitability but he cannot increase store traffic substantially, perhaps due to limited store capacity or limits on arrival traffic, such being the case during the COVID-19 pandemic (Aydinli et al., 2020), or (b) when traffic is already large but conversion is comparably small (Perdikaki et al., 2012). This is often the case in environments of intense comparison shopping.

Pornpitakpan and Han (2013) document that in brick-and-mortar stores shoppers make most of their decisions on the basis of environmental cues, some static such as store displays, and some dynamic such as store traffic. Our research could allow managers to connect some dynamic environmental cues with the effects of a promotion on conversion. For instance, during the promotion that our paper examines, the number of visits increases but the conversion fraction decreases relative to the predicted baseline. Our analysis is consistent with others finding that the presence of other people may adversely affect conversion probabilities (Perdikaki et al., 2012). To mitigate, or even reverse adverse effects during promotions, store managers can conduct activities such as modifying the location of products to keep the density of visitors at levels comfortable to all. Interestingly, some of these activities may help avoid two undesirable behaviors: unintended interpersonal touches (Martin, 2012) or excessive manipulation of products (Argo et al., 2006), which often lead to negative brand evaluations of store products. Store managers may also consider increasing the staff to make the store appear less crowded (Perdikaki et al., 2012). Thus, our approach can assist managers to inhibit situations that lead customers to engage in undesirable behaviors and increase conversion. Furthermore, the method used also suggests that managers can use simple technologies to expand their abilities to analyze consumers and to estimate the effects of many marketing actions beyond promotions.

In a dynamic, competitive market environment, a store manager formulates pricing policies, determines needed sales promotions, and authorizes clearance sales, to achieve and maintain levels of success of his store. Evaluating effects of these marketing activities on store

performance can help this manager to allocate resources to these activities wisely. However, computing these effects can be difficult because several factors, such as focal store's location and competitors' clearance promotions, may explain simultaneously changes in store performance during a specific time interval.

An approach to accurately compute these effects is a marketing experiment that involves an intervention, an outcome, an experimental design, and a control condition (Palmatier & Sridhar, 2017). Broadly, to compute the effect of a price promotion on conversion fraction, one usually considers two stores that are identical in all aspects, except that one of the stores performs the price promotion and the other one does not. The latter store represents the control condition. During the experiment, visitors arrive randomly to one of these stores only. When the price promotion period ends, one compares the conversion fraction between both stores to assess the price promotion effect on conversion fraction.

Conducting a marketing experiment may seem straightforward, but actually it is not. Finding two stores that are identical in all aspects except in the price promotion condition is difficult. On the one hand, if both stores are located far from each other, then they could operate under different environmental conditions. On the other hand, if both stores are close each other, then a shopper can be exposed simultaneously to both promotion and no-promotion conditions during the experiment. In addition, creating promotion and no-promotion conditions when both stores are close each other is unfeasible from a managerial perspective because it does not guarantee random assignment of shoppers to stores.

Our approach provides a feasible analytical alternative for store managers to conduct marketing experiments when evaluating the effect of an intervention on an outcome. With data outside the intervention period, our approach predicts the outcome for the control condition during the intervention. Thus, a store manager can compute the intervention effect on the outcome by comparing the observed and predicted outcome for the same store without the above-mentioned complexities. In addition, the dynamic nature of our approach allows a store manager to track this effect during the intervention period.

One may think that although the promotion does not affect conversion, increased traffic still results in increased sales (Gauri et al., 2017). This conclusion is not warranted without further analyses about the effect of the promotion on the expenditure of converting visitors. Indeed, a promotion may result in an increase of such expenditures, as one often expects, or may result in a decrease of them if, for instance, store crowding associated with the promotion leads customers to leave the store with small purchases. Indeed, a promotion may lead a large increase in arrival traffic and subsequent crowding that can alter adversely shoppers' experience, such as non-intentional touches among strangers (Martin, 2012), excessive manipulation of products in the store (Argo et al., 2006), or alter their basket composition (Aydinli et al., 2020). Our research provides information useful to store managers to identify unexpected shoppers' experiences associated with promotions that promote or inhibit conversion.

5.2. Limitations and future research

As with most studies, both the model and the data have limitations. Some of these prevent us from generalizing our results, but offer pathways for future research. For instance, the data we use in our analyses are from a single store with a promotional period of two days. Although our results are consistent with prior research indicating that promotions may increase store traffic but decrease conversion (Perdikaki et al., 2012), we do not have a broad enough data set to conclude that this effect is a general finding and not unique to the specifics of our data. Future research might utilize our model to examine promotions across an array of stores, time frames, promotional discount percentages, and items to find a broader spectrum of normative managerial recommendations. An additional limitation is that our data do not distinguish

between sales of promoted and non-promoted items. To draw specific conclusions as to how an item-based promotion affects sales of that item, one would need register receipts with records indicating the sales of the promoted item and non-promoted items. For our purposes, we focus on changes in store conversion based on a brand promotion rather than assessing the more narrow issue of whether the promotion was effective to increase conversion with the sale of the item promoted (Blattberg et al., 1995; Guha et al., 2018). In the future, one might utilize the prediction model to evaluate whether promotional effects are due to store-level, category-level, or item-level changes in revenues.

Though our study does not include environmental factors such as store displays, shelf locations, or signage, research provides evidence that these atmospheric elements influence consumers' behavior (Pornpitakpan & Han, 2013; Puccinelli et al., 2009). Our research could allow managers to connect some dynamic environmental cues with the effects of a promotion on conversion. As the analysis of the illustrative example indicates, the promotion increases the number of visitors and decreases conversion fractions during some time band. Thus, the data suggest that the presence of other store visitors might adversely affect conversion probabilities. Indeed, earlier research suggests that this might be due to shoppers feeling crowded (Perdikaki et al., 2012). It would be interesting to investigate whether managers can alter atmospheric factors to counteract shoppers feeling crowded (Machleit et al., 2000). Another environmental factor that might affect conversion is the depth of the promotional offer. In this research example, the retailer offered a 15% discount on two brands without substantially increasing conversion. As prior research indicates, promotional discounts often must pass a threshold to affect behavior (Gupta & Cooper, 1992). Perhaps a 15% discount was not sufficient to generate a meaningful number of additional conversions. Thus, exploring how the depth of a promotion might affect traffic and conversion is an open opportunity for future research.

The model has limitations that can be addressed with more detailed data or by augmenting the model itself. In its current instantiation, the model does not incorporate descriptive variables, such as customers' demographic or psychographic characteristics, though others find that these consumer characteristics affect the success of promotions (Ailawadi et al., 2001). Access to specific customer characteristics is certainly possible given the advances in technology (Grewal et al., 2018). For instance, store traffic sensors can easily be linked to mobile devices with access obtained at the time customers sign up for loyalty cards (Graeff & Harmon, 2002) or when they download an app (Keusch et al., 2019). This is common practice, though such access concerns privacy researchers and public policy makers (Inman & Nikolova, 2017).

Retailers often use a combination of marketing activities to influence consumers though our model assesses the effects of a single marketing action (Naik et al., 2005). For instance, Naik et al. (2005) investigate how advertising and price promotions interact in determining market share. They find that these two activities have a negative interaction effect on the market share of five detergent brands, despite the fact that their respective main effects are positive. The interaction is such that a deep price promotion attenuates the positive effect of advertising on the product's market share. The prediction model permits one to conduct similar analyses on the joint effects of marketing activities on conversion rates, or more specifically, on the joint effect of any short-term marketing activity combined with another long or short-term marketing activity. Returning to the advertising and price promotion example, one would start analyzing the joint effects problem by collecting data on the number of visitors and converters during four types of periods: (1) periods with neither advertising nor promotion; (2) periods with advertising but no price promotion; (3) periods with no advertising but with a price promotion; and (4) periods with both advertising and price promotion. In this case, period 1 establishes the baseline, periods 2 and 3 the main effects, and period 4 the interaction

effect. Indeed, with the prediction approach, one can just as easily estimate the effect of a single marketing action on conversion as they can multiple effects. We suggest that doing so would test the external validity of the model because the marketplace is often described by multiple actions happening simultaneously.

Finally, retailers often do not remain idle when their competitors launch promotions (Nijs et al., 2001). In a dynamic, competitive market environment, a store manager formulates pricing policies, determines promotions, schedules advertising, and develops other marketing activities by considering how competitors might react (Ailawadi et al., 2001). Again, with the prediction approach, one can augment the model to account for competitive reactions by following a similar four period model, but this time substituting the focal store and its competition into the time periods describing marketing actions. We suggest that taking this approach would be a fruitful area for future research.

5.3. Conclusion

This research suggests that advancements in retailing technology provide promising opportunities for data generation and analytics to garner deeper insights into the effect of marketing actions on the success of the retailer. Here, we specifically utilized data provided by a retailer employing store entry sensors and register receipts to illustrate a novel model assessing the influence of price promotions on customer conversion. The model and research developed offer promise in establishing baselines for comparisons that can be used in a variety of common retail scenarios.

References

- Abraham, M. M., & Lodish, L. M. (1987). Promoter: An automated promotion evaluation system. *Marketing Science*, 6(2), 101–123.
- Ailawadi, K. L., Neslin, S. A., & Gedenk, K. (2001). Pursuing the value-conscious consumer: Store brands versus national brand promotions. *Journal of Marketing*, 65(1), 71–89.
- Ando, T. (2008). Measuring the baseline sales and the promotion effect for incense products: a Bayesian state-space modeling approach. *Annals of the Institute of Statistical Mathematics*, 60(4), 763.
- Argo, J. J., Dahl, D. W., & Morales, A. C. (2006). Consumer contamination: How consumers react to products touched by others. *Journal of Marketing*, 70(2), 81–94.
- Aydinli, A., Lamey, L., Millet, K., ter Braak, A., & Vuegen, M. (2020). How do customers alter their basket composition when they perceive the retail store to be crowded? An empirical study. *Journal of Retailing* (in press).
- Baker, J., Levy, M., & Grewal, D. (1992). An experimental approach to making retail store environmental decisions. *Journal of Retailing*, 68(4), 445–460.
- Balchandani, A., Hancock, B., Phillips, S., & Wachinger, T. (2020). Rebooting Retail: How technology will shape the future of retail. Retrieved on December 31, 2020, from <https://www.mckinsey.com/industries/retail/our-insights/rebooting-retail-how-technology-will-shape-the-future-of-retail>.
- Biswas, D., Lund, K., & Szocs, C. (2019). Sounds like a healthy retail atmospheric strategy: Effects of ambient music and background noise on food sales. *Journal of the Academy of Marketing Science*, 47(1), 37–55.
- Blattberg, R. C., Briesch, R., & Fox, E. J. (1995). How promotions work. *Marketing Science*, 14(Suppl. 3), G122–G132.
- Bucklin, R. E., & Gupta, S. (1999). Commercial use of UPC scanner data: Industry and academic perspectives. *Marketing Science*, 18(3), 247–273.
- Clodfelter, R. (2004). An examination of pricing accuracy at retail stores that use scanners. *Journal of Product & Brand Management*, 13(4), 269–278.
- Cutler, K., & Rowe, C. (1990). Scanning in the supermarket: for better or worse? A case study in introducing electronic point of sale. *Behaviour & Information Technology*, 9(2), 157–169.
- De Haan, E., Wiesel, T., & Pauwels, K. (2016). The effectiveness of different forms of online advertising for purchase conversion in a multiple-channel attribution framework. *International Journal of Research in Marketing*, 33(3), 491–507.
- Doyle, P., & Saunders, J. (1985). The lead effect of marketing decisions. *Journal of Marketing Research*, 22(1), 54–65.
- Doyle, P., & Saunders, J. (1986). Measuring the true profitability of sales promotions. *Journal of the Operational Research Society*, 37(10), 955–965.
- Epstein, L. D., Flores, A. A., Goodstein, R. C., & Milberg, S. J. (2016). A new approach to measuring retail promotion effectiveness: A case of store traffic. *Journal of Business Research*, 69(10), 4394–4402.
- Fok, D., Paap, R., & Franses, P. H. (2012). Modeling dynamic effects of promotion on interpurchase times. *Computational Statistics & Data Analysis*, 56(11), 3055–3069.

- Gauri, D. K., Ratchford, B., Pancras, J., & Talukdar, D. (2017). An empirical analysis of the impact of promotional discounts on store performance. *Journal of Retailing*, 93(3), 283–303.
- Goodstein, R. C. (1994). UPC scanner pricing systems: Are they accurate? *Journal of Marketing*, 58(2), 20–30.
- Goodstein, R. C., & Escalas, J. E. (1995). Improving pricing accuracy at the supermarket: Electronic shelving systems and public policy. *Journal of Public Policy & Marketing*, 14(2), 216–224.
- Graeff, T. R., & Harmon, S. (2002). Collecting and using personal data: Consumers' awareness and concerns. *Journal of Consumer Marketing*, 19(4), 302–318.
- Grewal, D., Ahlborn, C.-P., Beitelspacher, L., Noble, S. M., & Nordfält, J. (2018). In-store mobile phone use and customer shopping behavior: Evidence from the field. *Journal of Marketing*, 82(4), 102–126.
- Grewal, D., Hulland, J., Kopalle, P. K., & Karahanna, E. (2020). The future of technology and marketing: A multidisciplinary perspective. *Journal of the Academy of Marketing Science*, 48(4), 1–8.
- Grewal, D., & Roggeveen, A. L. (2020). Understanding retail experiences and customer journey management. *Journal of Retailing*, 96(1), 3–8.
- Guha, A., Biswas, A., Grewal, D., Verma, S., Banerjee, S., & Nordfält, J. (2018). Reframing the discount as a comparison against the sale price: Does it make the discount more attractive? *Journal of Marketing Research*, 55(3), 339–351.
- Gupta, S. (1988). Impact of sales promotions on when, what, and how much to buy. *Journal of Marketing Research*, 25(4), 342–355.
- Gupta, S., & Cooper, L. G. (1992). The discounting of discounts and promotion thresholds. *Journal of Consumer Research*, 19(3), 401–411.
- Hanssens, D. M., Pauwels, K. H., Srinivasan, S., Vanhuele, M., & Yildirim, G. (2014). Consumer attitude metrics for guiding marketing mix decisions. *Marketing Science*, 33(4), 534–550.
- Hardesty, D. M., Goodstein, R. C., Grewal, D., Miyazaki, A. D., & Kopalle, P. (2014). The accuracy of scanned prices. *Journal of Retailing*, 90(2), 291–300.
- Hui, S. K., Fader, P. S., & Bradlow, E. T. (2009). Path data in marketing: An integrative framework and prospectus for model building. *Marketing Science*, 28(2), 320–335.
- Inman, J., & Nikolova, H. (2017). Shopper-facing retail technology: A retailer adoption decision framework incorporating shopper attitudes and privacy concerns. *Journal of Retailing*, 93(1), 7–28.
- Keusch, F., Struminskaya, B., Antoun, C., Couper, M. P., & Kreuter, F. (2019). Willingness to participate in passive mobile data collection. *Public Opinion Quarterly*, 83(S1), 210–235.
- Kim, J.-E., & Kim, J. (2012). Human factors in retail environments: A review. *International Journal of Retail & Distribution Management*, 40(11), 818–841.
- Lam, S. Y., Vandenbosch, M., Hulland, J., & Pearce, M. (2001). Evaluating promotions in shopping environments: Decomposing sales response into attraction, conversion, and spending effects. *Marketing Science*, 20(2), 194–215.
- Li, H., & Kannan, P. (2014). Attributing conversions in a multichannel online marketing environment: An empirical model and a field experiment. *Journal of Marketing Research*, 51(1), 40–56.
- Machleit, K. A., Eroglu, S. A., & Mantel, S. P. (2000). Perceived retail crowding and shopping satisfaction: What modifies this relationship? *Journal of Consumer Psychology*, 9(1), 29–42.
- Martin, B. A. (2012). A stranger's touch: Effects of accidental interpersonal touch on consumer evaluations and shopping time. *Journal of Consumer Research*, 39(1), 174–184.
- Matzler, K., Mooradian, T. A., King, L., & Linder, A. (2010). Converting browser to buyers: An approach to measure and increase conversion rates in retailing. *Innovative Marketing*, 6(1), 34–38.
- Moe, W. W., & Fader, P. S. (2004). Dynamic conversion behavior at E-commerce sites. *Management Science*, 50(3), 326–335.
- Motyka, S., Suri, R., Grewal, D., & Kohli, C. (2016). Disfluent vs. fluent price offers: Paradoxical role of processing disfluency. *Journal of the Academy of Marketing Science*, 44(5), 627–638.
- Naik, P. A., Raman, K., & Winer, R. S. (2005). Planning marketing-mix strategies in the presence of interaction effects. *Marketing Science*, 24(1), 25–34.
- Nijs, V. R., Dekimpe, M. G., Steenkamps, J.-B. E., & Hanssens, D. M. (2001). The category-demand effects of price promotions. *Marketing Science*, 20(1), 1–22.
- Palmatier, R. W., & Sridhar, S. (2017). *Marketing strategy: Based on first principles and data analytics*. London, UK: Macmillan International Higher Education, ISBN: 9781137526236.
- Park, C. H. (2017). Online purchase paths and conversion dynamics across multiple websites. *Journal of Retailing*, 93(3), 253–265.
- Perdikaki, O., Kesavan, S., & Swaminathan, J. M. (2012). Effect of traffic on sales and conversion rates of retail stores. *Manufacturing & Service Operations Management*, 14(1), 145–162.
- Pornpitakpan, C., & Han, J. H. (2013). The effect of culture and salespersons' retail service quality on impulse buying. *Australasian Marketing Journal (AMJ)*, 21(2), 85–93.
- Puccinelli, N. M., Goodstein, R. C., Grewal, D., Price, R., Raghubir, P., & Stewart, D. (2009). Customer experience management in retailing: Understanding the buying process. *Journal of Retailing*, 85(1), 15–30.
- Rigby, R. A., & Stasinopoulos, D. M. (2005). Generalized additive models for location, scale and shape. *Journal of the Royal Statistical Society. Series C. Applied Statistics*, 54(3), 507–554.
- Shankar, V., Kalyanam, K., Setia, P., Golmohammadi, A., Tirunillai, S., Douglass, T., Hennessey, J., Bull, J., & Waddoups, R. (2021). How technology is changing retail. *Journal of Retailing* (In Press).
- Stasinopoulos, D. M., & Rigby, R. A. (2007). Generalized additive models for location scale and shape (GAMLSS) in R. *Journal of Statistical Software*, 23(7), 1–46.
- Van Heerde, H. J., Leeflang, P. S., & Wittink, D. R. (2000). The estimation of pre- and postpromotion dips with store-level scanner data. *Journal of Marketing Research*, 37(3), 383–395.
- Van Heerde, H. J., & Neslin, S. A. (2017). Sales promotion models. In *Handbook of marketing decision models* (pp. 13–77). Springer.
- Xu, L., Duan, J. A., & Whinston, A. (2014). Path to purchase: A mutually exciting point process model for online advertising and conversion. *Management Science*, 60(6), 1392–1412.
- Zhang, X., Li, S., & Burke, R. R. (2018). Modeling the effects of dynamic group influence on shopper zone choice, purchase conversion, and spending. *Journal of the Academy of Marketing Science*, 46(6), 1089–1107.

Leonardo D. Epstein is Professor of Statistics at the Business School of Universidad de los Andes, Chile, where he has been teaching courses in business statistics, management science and quantitative marketing. He received his M.Sc. and Ph.D. from Carnegie Mellon University. Dr. Epstein's recent research concentrates on inventory models for rental items, and methods and models for the evaluation of marketing and sales promotion campaigns. He is also interested in biostatistics and collaborates with public health researchers. Dr. Epstein research has been published in *Computers and Industrial Engineering*, *Annals of Operations Research* and *Journal of Business Research*, *The American Journal of Epidemiology*, *The Lancet* and *Statistics in Medicine*.

Ignacio Inostroza-Quezada is Assistant Professor at the Universidad de los Andes, Chile. He currently teaches statistical inference, applied econometrics, and marketing research. He earned his Bachelor's degree in Mathematics, a Master's in Statistics, and a Master's in Management Sciences at the Pontifical Catholic University of Chile. He is also a Ph.D. Candidate in Marketing at Rutgers Business School, Rutgers University. His dissertation research concerns the relationship between new product pricing strategies and the structure of connections among consumers in a social network. His research interests are quantitative models in marketing.

Ronald C. Goodstein is Associate Professor of Marketing at Georgetown University's McDonough School of Business, where he has taught since 1999. Dr. Goodstein received his BS from the McIntire School of Commerce at the University of Virginia and his Ph.D. from Duke University's Fuqua School of Business. His research includes topics related to retailing, advertising, emotions, and consumer psychology. Dr. Goodstein has published in many top journals including the *Journal of Business Research*, *Journal of Consumer Research*, *Journal of Marketing Research*, *Journal of Marketing*, and *Marketing Letters* among others.

S. Chan Choi is Professor of Marketing Sciences at Rutgers Business School (RBS), where he has taught since 1990. He is a former chair of Marketing Department and as the director of the Center for Marketing Research, supervises the MBA concentration of Marketing Research Insights and Analytics (MRIA). He received his BS (Pharmaceutics) and BA (Management) from Seoul National University, an MBA from Michigan State University, and a Ph.D. from the Wharton School. He was a Fulbright Visiting Scholar to KAIST in Korea. Dr. Choi's research focus is in modeling price competition in various marketing contexts and competitive product positioning. His current research topics include designing efficient coordination mechanisms within a distribution channel, quality-price competition in distribution channels, and marketing research in the pharmaceutical industry. He has published in such journals as *Marketing Science*, *Management Science*, *Journal of Consumer Research*, *Journal of Retailing*, *Journal of Interactive Marketing*, *Journal of Product Innovation Management*, and *Marketing Letters*.