Available online at www.sciencedirect.com
ScienceDirect
Journal of Interactive Marketing 42 (2018) 32 – 45
JOURNAL OF INTERACTIVE MARKETING
www.elsevier.com/locate/intmar

The Effect of House Ads on Multichannel Sales☆

Marcel Goic * & Rodolfo Álvarez & Ricardo Montoya

Department of Industrial Engineering, University of Chile, Chile

Abstract

Similar to in-store displays in brick-and-mortar stores, house ads promote a set of specific products for customers who have reached the website. In contrast to general display advertising whose primary goal is to bring traffic to the website, these self-promotional ads are aimed to highlight specific products and enhance conversion. We analyzed more than 300 house ad campaigns to study the effect of this type of promotional display on customer behavior across channels. We included not only direct effects on SKU sales in all channels, but also the promotional effect at the category level. Our model uses aggregated data that are easy to collect for most multichannel retailers, facilitating its implementation in similar settings. We found that (1) despite observing positive cross-channel effects, the primary effect occurs on online sales, (2) the effects are usually short-lived, and (3) there are no spillover effects on the corresponding category. We characterize the effects that house ads have on the whole system in terms of design variables such as type of display, and scope and duration of the campaign. Our evidence suggests that the effectiveness depends on the product category and that regular banners are the most effective in generating traffic. Interestingly, the depth of the promotion plays no role on the effect of the house ads' effectiveness. Based on the results, we provide suggestions for improving routine promotional planning.

Keywords: House ads; Online display advertising; VAR models; Multichannel retailing

Introduction

Sales promotions have been historically one of the most studied elements in the retail industry. A common distinction made when analyzing different branches of the promotional mix is between *feature* and *display*. While *feature* is typically distributed in newspaper inserts, *display* is carried out in the store to highlight some special value offerings for specific products. These elements can have different effects in boosting sales, and may even have a combined effect that goes beyond the simple sum of the two (Neslin 2002). In online retailing, the vast majority of research has been conducted to analyze the effect of *display advertising* on browsing and purchasing behavior. The primary goal of display advertising (e.g., banner ads on different websites) is to bring customers to a retailer's website, just as *feature* is devoted to increasing store patronage. For instance, social networks such as Facebook and Twitter that have jointly accounted for 33% of the digital display advertising in the US in 2017,[1] provide a complete advertising platform to third parties.

We use the names *house ads*[2] or *internal display* to classify any promotional information presented on the retailer's website that is devoted to signaling some specific attributes of a single product (or a narrow family of them), just as *display* does in brick and mortar stores. In this research, we investigate the effect of house ads on the multichannel sales of a department store at the SKU level empirically. In contrast to the abundant research on *display advertising*, the effect of *house ads* remains largely unexplored.

☆ The first author gratefully acknowledges partial funding by FONDECYT (Project No 11121434). The second author gratefully acknowledges partial funding by FONDECYT (Project No 1151395) The first and second authors also acknowledges Complex Engineering Systems Institute, ISCI (ICM-FIC: P05-004-F, CONICYT: FB0816).
* Corresponding author at: Beauchef 851, Santiago, Chile.
E-mail addresses: mgoic@dii.uchile.cl (M. Goic), rmontoya@dii.uchile.cl (R. Montoya).

[1] *eMarketer*. March 26, 2017.
[2] In the industry jargon, house ads are also called *self-promotion ads*, *on-site display*, or *internal links*.

https://doi.org/10.1016/j.intmar.2017.12.001
1094-9968/

The use of house ads is a common practice in online retailing. Nowadays, virtually all retailers' websites include a series of banners on their homepages highlighting specific products, complete product categories, or some special events. Although it is possible for firms to include external banners, in this investigation we restricted our attention to digital displays that help customers to navigate within the website once they have landed at the retailer's web page. Apart from the identity of the advertiser, there are important differences between internal (in-house) and external display advertising. For internal displays, the retailer has direct control over the content, size, and frequency of the ads, and the destination of the underlying hyperlinks, and even the context in which those banners are displayed. Thus, the content of house ads typically focuses on informing customers about specific promotions of either a specific SKU or a complete product category. Additionally, the selection of which banner to show at each point in time can be determined based on internal business rules, and those changes have almost no cost beyond that of the graphic design.

A clear understanding of the effect of house ads is useful because it can lead to designing more effective campaigns that take the temporal length of the effect into account, promoting product categories in which customers are more sensitive to this promotional vehicle, and leveraging any interaction with other elements of the marketing mix, such as direct marketing or price discounts. From a strategic point of view, it is important to characterize the nature of the effect of internal promotions on sales performance. Do house ads affect only the promoted product, or is there a positive spill-over effect at the category level? Answering this question can change how firms decide on the set of products to be displayed. If internal displays benefit only the sale of the promoted products, retailers can charge manufacturers to feature them, just as they are charged for certain shelf positions in brick and mortar stores. This can lead to a new source of monetization of the online channel. Indeed, according to Fortune Magazine, audience monetization can generate a 10% incremental revenue for e-commerce players.[3]

One of the major changes that the retail industry has faced in recent decades is the advent of additional channels to complement product offerings in traditional stores. We therefore consider it important to adopt a multichannel perspective. Currently, customers use many of these channels actively even on a single purchase occasion. For example, before going to an offline store, they might research online (Verhoef, Neslin, and Vroomen 2007) and be exposed to house ads. The literature has documented some cross-channel effects from other promotional tools (Dinner, Van Heerde, and Neslin 2014; Liaukonyte, Teixeira, and Wilbur 2015). In our investigation, we evaluate how various internal campaigns, occurring only on the website, affect both online and offline sales. Quantifying the cross-channel effect could provide important insights into how firms should balance multiple retail channels.

In this paper, we analyze the cross-channel effect of house ads on a set of endogenous variables that are related temporally. Specifically, we attempt to answer the following questions: (i) How do house ads affect website visitation and sales on all available channels? (ii) How long do these effects last? (iii) What combinations of displays and products are most effective? and (iv) Is there any spill-over effect at the category level?

When the retailer places a website banner for a given SKU, customers might be motivated to visit the corresponding product page inducing various multichannel behaviors. Some customers might want to compare alternatives, which may drive more online category sales. Other buyers might go to the closest store to purchase the product. In both cases, the customers' responses might require some time to materialize. The delay between the activation of the online display and the purchase varies considerably depending on the response to the promotion. To deal with this dynamic we use a time series approach.

To conduct our empirical investigation, we used eight months of transactional data from a multichannel department store that sells through both brick-and-mortar and online stores. In our dataset, along with sales on all channels and the number of online visits at the SKU level, we considered more than three hundred products that have been displayed on the homepage of the website, and the exact size and location of the display.[4] Considering that we observed activity from diverse SKUs, we can characterize the nature of the relationship, identifying the type of display that is most effective in stimulating sales, and determining what type of product it makes the most sense to display. To incorporate customer dynamics, we first used a vector autoregressive approach for each SKU in the sample, and then ran linear regressions to characterize how the impact of house ads depends on the product category, the attributes of the banners, and the interactions with other promotional vehicles. Our results indicate that, for most products, the effect of the online displays on sales lasts for only two days or less. However, an important portion of the campaigns exhibit longer effects (in 17% of the campaigns the effect lasted for more than a week). In terms of the effects on sales, we found evidence of both online and offline increments, but the effect is, in general, greater for online sales. Interestingly, we found almost no effect at the category level and therefore only promoted items benefit from house ads. Our results suggest that the company should consider the multichannel effect of house ads when evaluating the success of each campaign. The brief effect of these promotions indicates that the company needs to change these promotions dynamically. Also, given that most of the effect is at the SKU level, the company should select the items to be displayed carefully, and perhaps involve manufacturers in the decision-making process.

The rest of this paper is organized as follows: In the next section, we present a brief literature review regarding online display and multichannel retailing. We then describe our modeling framework explaining how we capture the most relevant customer dynamics. Next, we discuss our empirical analysis including a detailed description of the data we used and the results derived

[3] http://fortune.com/2014/06/20/online-retailers-are-selling-more-than-just-stuff-theyre-selling-eyeballs-and-audiences/.

[4] For simplicity, we focus on the inclusion of digital ads on the homepage of the transactional website of the firm. In our empirical setting, the front page receives a large percent of customer visits and most of the internal display activity.

from the model. We conclude with a discussion of the main results and propose some avenues for future research.

Literature Review

Although the effects of house ads are still largely unexplored in the literature, there are a few related streams of research that are relevant for our investigation. As we stated in the introduction, an important motivation for this research comes from the well-studied effect of in-store display on sales in brick and mortar retail stores. Since the introduction of scanner panel data in the early 80s, display has been recognized as an effective tool for influencing brand choice (Abraham and Lodish 1987; Kumar and Leone 1988). Later, display has been included in a variety of demand models including structural equilibrium models (Besanko, Gupta, and Jain 1998) and hierarchical Bayesian models (Montgomery 1997) among others. As pointed out by Neslin (2002), the effect of display is strong and consistently superior to featured advertising. Importantly, display can have a positive effect on sales even in the absence of a price discount (Inman, McAlister, and Hoyer 1990). Accordingly, we expect that internal banners also would have positive effects on customer attention and, therefore, sales.

In the online channel literature, there is a large body of research analyzing the effect of banner advertising on sales. Manchanda et al. (2006) was one of the first papers to connect exposure to online advertising with actual sales, finding a positive relation. Since then, a variety of studies have analyzed different aspects of online advertising, such as the identification of its differential effect at different stages of the purchase funnel (Hoban and Bucklin 2015), and the evaluation of short and long-term effects of online advertising (Breuer, Brettel, and Engelen 2011).

This article also contributed to the growing literature on the multichannel effect of promotion and advertising. In this regard, Dinner, Van Heerde, and Neslin (2014) explored the effect of advertising in multichannel campaigns and found that cross effects in advertising are appreciable, particularly from online advertising to offline sales. In their research, they looked at both traditional advertising and new online mechanisms such as display and paid search, and used market level data. Here we focus on internal display, and we conduct our analysis at the campaign level, enabling us to shed some light on the specific attributes that make on-site banners more effective.

In terms of the design of the campaigns, Lohtia, Donthu, and Hershberger (2003) investigated the use of interactivity, color, and animation in banner click-through rates. Similarly, Braun and Moe (2013) found that different creatives can lead to differing levels of effectiveness on visiting and conversion behavior. Along the same line, we investigate how the type of banners and the nature of the associated campaigns affect the performance of the promotion providing a more comprehensive understanding of the conditions favoring the effectiveness of house ads.

One notable exception to the studies exploring the role of on-site banner advertising was that conducted by Rutz and Bucklin (2012) in which the authors found that the effect of house ads is short-lived, and that these banners affect behavior only during a current browsing session. As in our study, Rutz and Bucklin connected internal display empirically with future browsing behavior. However, unlike in our research, they did not track sales, and restricted their attention to customer decisions on the next page to visit. Moreover, they did not track customer activities at the category level, making it impossible to quantify cross-product spillovers, which is one of the key questions in our study.

The use of house ads is also related to the implementation of recommendation systems in which retailers select a reduced set of products to be suggested to customers based on their historical data (see e.g., Ansari, Essegaier, and Kohli 2000, and Senecal and Nantel 2004). In both cases, the firm chooses a subset of attractive products to be displayed on the website. However, when using recommendation systems, these decisions are made at the customer level, while in our application all customers face the same set of highlighted products. If the retailer were using an automated recommendation system, we could expect a significant lift in the conversion to the recommended products, but it would be difficult to observe a market level impact for any given item since various products would be shown to different customers.

To deal with the dynamic relationship between internal display, browsing behavior, and multichannel sales, we use a multivariate time series approach. The use of time series models to evaluate promotional effectiveness has a long tradition in marketing. It has been used to study complex dynamic effects, and it has been proven to be a powerful analytical tool for evaluating the long-term impact of marketing spending (e.g., Dekimpe and Hanssens 1999, and Nijs et al. 2001). For example, it has been stated that when the sales pattern is evolving, managers should pay attention not only to the short-term responses, but also to the long-run consequences of their actions (Dekimpe and Hanssens 2000). The interaction among marketing vehicles was investigated by Nijs et al. 2001 who found that price promotions accompanied by advertising have a strong short-term impact on sales, but these marketing actions rarely show persistence over time. Joshi and Hanssens (2010) used time series analysis to describe customer responses, and evaluated the impact of advertising on sales revenues and profits.

In the past few years, time series analysis has been used to describe several facets of multichannel dynamics. For example, Pauwels et al. (2011) investigated the effect of online information on offline revenues. In their setting, they analyzed the existence of a pure informational online channel that affects offline sales, and consequently, they did not study how different elements of the display might affect customer behavior. Moreover, they studied a website that has no transactional capabilities, thus all gain in sales was materialized through offline stores. In this regard, Wiesel, Pauwels, and Arts (2011) is closer to our study because they analyzed how expenditures in various marketing activities contribute to modifying consumer behavior on online and offline channels. Their investigation considered a wider range of promotional vehicles than we do, but internal display was not included. Moreover, they analyzed aggregated expenditures and, therefore, there was no description of how different campaign characteristics may have influenced their effectiveness.

In summary, this paper contributes to the literature by providing an empirical evaluation of the effectiveness of internal

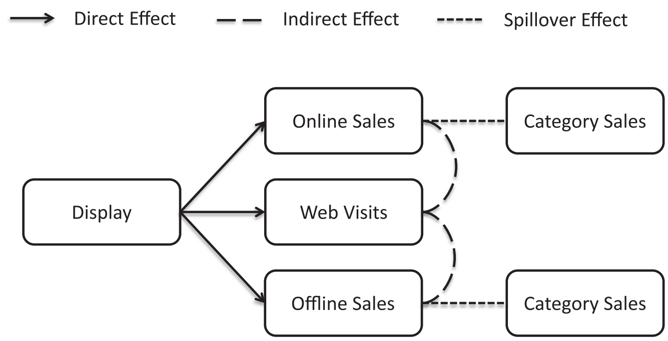

Fig. 1. Graphic representation of theoretical framework.

display that, to the best of our knowledge, has not been analyzed before. In our evaluation, we consider not only the direct effect on sales, but also cross-category and cross-channel effects. Moreover, our analysis at a campaign-level enables us to describe not only the aggregate impact of house ads, but also to characterize the conditions that are most likely to exhibit superior performance.

Modeling Framework

We propose to study the direct and indirect effects of house ads on customer behavior across channels. More specifically, we study the direct effects on online sales, offline sales, and website visits, and the indirect effects among these variables. The indirect effects are primarily due to potential sales cannibalization or complementarity across channels. We also investigate potential spillover effects at the category level, both online and offline, caused by incremental sales or substitution within product categories. Fig. 1 illustrates the proposed theoretical framework.

The interaction between channels generates dynamic relationships posing methodological challenges in the evaluation of the nature and magnitude of the underlying phenomena. For example, a shopping pattern that is widely discussed in the industry is the research shopping phenomenon (Verhoef, Neslin, and Vroomen 2007)[5] that occurs when customers do research on their potential purchases in the online environment, but when deciding to buy, they make the purchases in traditional offline stores. Considering that the online activity usually occurs before customers go to the store to make their purchases, it is necessary to explore the customers' online activity before the occasion of the actual purchase when trying to identify this behavior using transactional data. Naturally, the period in which the activity on the website can affect offline sales depends on the nature of the purchase itself. For instance, for high involvement purchases, customers tend to engage in longer external search periods (Beatty and Smith 1987), and therefore it is plausible that they search online several days before purchasing the product. Similarly, for purchases motivated by short-term promotions, or purchases associated with impulsive behavior, the purchase could occur just a few minutes after online browsing.

To investigate these dynamic relationships, we use vector autoregressive models with exogenous variables (VARX). The availability of more detailed historical data has enabled a significant growth in the use of time series analysis in marketing in the past few years (Dekimpe and Hanssens 2000). In our application, the choice of these models is motivated by several factors. First, in the VARX methodology we can analyze the evolution of several endogenous variables simultaneously, and, therefore, we can accommodate all the effects we described in the theoretical framework, including direct, indirect, and category spillover effects. Second, as we have argued, the nature of the relationship between the activities across channels can vary substantially among campaigns in the number of time periods required to describe the interactions. Therefore, the chosen approach provides flexibility in accommodating different time lag structures across campaigns. Furthermore, our choice of developing a model using aggregated data is justified by the limited availability of individual data. As is typical in the retail industry, only about 20% of website visitors in our dataset could be identified at the time of the navigation, and, among them, only a small fraction could be properly identified across all channels. Finally, our modeling framework provides a straightforward mechanism for evaluating the overall impact across channels of product display taking place in the online channel because online and offline activities are endogenous in the model.

As pointed out by Pauwels et al. (2004), one challenge in using time series analysis is the identification and evaluation of structural changes in the model. However, in our case, given the nature of the interventions in which the structure of the website remains constant and only the identity of the displayed items is changing, we consider it to be unlikely that they would generate any permanent effect in purchasing behavior or channel choice.

Similar to Nijs et al. (2001) we organize the analysis in two phases. In the first stage, we estimate a VARX model for each campaign to characterize the structure and magnitude of the relationship in the activity across channels. These models can provide a measure of the effectiveness of the campaign for each variable. In the second stage, we use a regression analysis to summarize and characterize the effectiveness of the campaigns in terms of product characteristics and website display.

Vector Autoregressive Model

Several methodological steps are required to specify a vector autoregressive model. We start by defining the set of endogenous variables. We are interested in studying the evolution of sales on online and offline channels ($SalesOnline_{st}$ and $SalesOffline_{st}$) as a consequence of house ads. Consistent with our theoretical framework, we consider that this promotion does not necessarily affect sales directly, but, instead, through the customer's visit to the website. In fact, this internal display could generate product awareness that might motivate customers to search for more information about products (Kireyev, Pauwels, and Gupta 2016) and, as a consequence, those customers might increase their browsing activity before deciding to make the purchase. In our

[5] Some case studies providing complementary evidence of research shopping can be found at http://www.thinkwithgoogle.com/case-studies/online-research-driving-offline-purchase-for-gortz.html and http://www.thinkwithgoogle.com/articles/proof-online-ads-increase-offline-sales.html.

model we include the variable $WebVisits_{st}$ that considers the number of visits to the pages associated with the products promoted in the campaign s at time t. Including website visits in the model allows us to have more flexibility in accommodating complex customer dynamics, and it also enables us to identify situations in which promotional activity influences product search, but does not affect sales in a significant way.

To complete the set of endogenous variables in the model, we include category sales across channels ($CatSalesOnline_{st}$ and $CatSalesOffline_{st}$). Our rationale behind these effects is that visitors do not necessarily have a particular product in mind when reaching the website and, therefore, if they are attracted by the highlighted product, they can initiate a new journey that could boost sales for products in the category that were not displayed on the front-page of the website. The inclusion of these category variables is important because it allows us to determine whether the potential variations in sales are generated only by substitution within the category, or if a promotional display generates incremental sales for the retailer. With this evaluation, we can identify who benefits from internal displays, and address one of the managerial questions that motivates this study.

In our model, we consider the promotional display on the front-page as an exogenous variable. Our interviews with the marketing department of the company indicated that the performance of previous promotions is not explicitly considered when deciding the set of SKUs to promote. In fact, the lack of a rigorous evaluation of the promotional effectiveness at the SKU level was one of the company's main motivations for getting involved in this project. Even though the managers stated that they consider seasonality when selecting the categories to display on the website (for example, they do not highlight heating devices in the summer-time), the timing of the promotions is only weakly related to sales patterns. While high demand seasons last for several months, the display of a promoted SKU lasts only on average for a few days.

We operationalize the promotional display on the front-page as a dummy variable, $DISP_{st}$, taking the value 1 if the SKU s is displayed in any form at day t. Notice that there are many different ways in which a display could be executed in terms of size, style, and location, and thus we could have used the activity in each type of display as an independent exogenous variable. However, we discarded such an approach because it would imply that, at the SKU level, many variables would have constant zero values. In fact, during the observational period, any given SKU is promoted at most a few times, with only the use of one or two types of displays.

Most of the campaigns are accompanied by emails sent to the customer database announcing the products under promotion. We evaluated the possibility of including the number of emails sent to customers as another exogenous variable. Unfortunately, we have only the total number of messages sent per campaign, but do not have the exact dates when those emails were sent. Not knowing how these e-mails are distributed over time prevents their use in the VARX model. From our conversations with the company executives, we learned that they acknowledge that the traffic of emails coming to the company via the website during these campaigns accounts for only about 6–8% of the total number of visits, and it is plausible to believe its influence is not a first order concern. Consequently, we dropped it from the model, and postponed the empirical description of the influence of emails for the second phase of the analysis. The resulting model is given in Eq. (1)

$$\begin{pmatrix} \text{WebVisits}_{st} \\ \text{SalesOnline}_{st} \\ \text{CatSalesOnline}_{st} \\ \text{SalesOffline}_{st} \\ \text{CatSalesOffline}_{st} \end{pmatrix} = \phi_{s0} + \sum_{i=1}^{p} \phi_{si} \begin{pmatrix} \text{WebVisits}_{st-i} \\ \text{SalesOnline}_{st-i} \\ \text{CatSalesOnline}_{st-i} \\ \text{SalesOffline}_{st-i} \\ \text{CatSalesOffline}_{st-i} \end{pmatrix} + \beta_s DISP_{st} + \varepsilon_{st} \tag{1}$$

where ϕ_s and β_s are parameters to be estimated and $\varepsilon_{st} \sim \mathrm{N}(0,\Sigma_s)$. Note that our modeling approach does not include lags for display. We made this simplification mainly to facilitate the exposition of the results. In addition, penalized fit criteria justify our decision. Specifically, when allowing for a more flexible structure of lags for the exogenous variable, the average AIC and BIC improve by less than 5% and 2%, respectively. These results suggest that most of the dynamics are captured by a single term.

To evaluate whether or not a time series is stationary, we applied the Dickey–Fuller unit root test (Shumway and Stoffer 2006) finding that 91.5% of all estimated time series pass the test without taking differences. For those that failed the test, we used the difference resulting in only one time series not passing the test after taking the differences. The failing time series was consequently discarded from the sample. For simplicity in Eq. (1) we present only the expression with no differences.

To evaluate the impact of an online display in the system, we focus on two sets of parameters. First, we analyze the direct effect through the parameter β. Then, we consider the differential impulse response function (DRF) providing an aggregate evaluation of the impact of the display on the evolution of the system.

Linear Regression on Differential Response Functions

After estimating the VARX parameters, we conducted a regression analysis to characterize the relationship between the impact of the online display and the design variables. Specifically, we consider an assessment of the magnitude of the effect of online display in the system as a dependent variable. Nijs et al. (2001) proposed two performance metrics, one capturing the net effect over the dust-settling period, and the other devoted to evaluating the persistent long-term impact of the promotion. In our application, the relative importance of the marketing intervention is devoted mainly to having a short-term impact, and, therefore, we focus on the net effect considering only the short-term. The net impact corresponds to the differential response function (DRF) with respect to the baseline situation where no display is present. We are interested in analyzing the effect in each behavioral component, and, therefore, we construct a separate regression for each endogenous variable. In these models, we take a logarithmic transformation and remove campaigns with a negative net impact. If, instead, we keep the negative values, the results are directionally similar, but the fit of the model is considerably worse.

As independent variables, we consider different factors that the retailer can use to design the campaign. More specifically, we selected the following components:

- **Department:** A dummy variable taking the value 1 if the promoted SKU belongs to the corresponding department. We consider 8 departments: Home Decor, Sports, Bedding, Appliances, Kitchen & Dining, Furniture, Kids, and Perfume.
- **Discount:** Some of the campaigns offer a price discount with respect to the regular price. Here we consider the percentage discount.
- **Number of emails:** The retailer periodically selects a subset of the customer base and sends those customers an email highlighting current promotions with an invitation to visit the website. We include this variable because it could enhance the incremental number of visits to the website.
- **Multichannel Price Promotion:** When a campaign includes a price discount, the discount could be applied to purchases made through any channel, or could be exclusively for those on the online channel. If the discount applies to any channel, the variable takes the value 1, and 0 otherwise.
- **Short Campaign:** The online displays of products vary in their time duration. Even though they could last for any number of days, we observed empirically that only two time horizons were implemented. Short campaigns last for one or two days whereas longer campaigns last for approximately a week. To accommodate this fact, we introduced a dummy variable, *Short*, which takes the value 1 if the campaign lasts at most for two days.
- **Display Type:** The company has many degrees of freedom when deciding the size and location of the display. As is common practice in the industry, the company works with a web template that has a relatively fixed structure in which the style and colors are constant. In the template, there is a lateral menu on the left side and a search bar at the top, leaving the center right part of the website for the display of the selected product, as is depicted in Fig. 2. The structure of the website provides several cells where particular information can be allocated. Among them we identify three main types:
  - **Carousel:** Located at the top of the available area, it covers its whole width providing a large display. An important feature of this type of display is that it is visible to the customer for only a few seconds and then automatically changes to another product. At any point in time, there is a short list of the products rotating in this display, and customers can go back and forth using a (numbered) navigation panel at the top.
  - **Cover:** It is a medium-sized display and it is the only one that is not SKU specific. Instead, it refers to a set of related products that are displayed only after the customer clicks on them. For example, it might be devoted to a special event of photography for which several cameras are displayed in the inner page.
  - **Regular:** All other available spaces are static and SKU specific banners, in the sense that when customers click on the display, they land on the SKU page where they can add the product to their shopping cart. Even though there are differences in terms of the locations, we cluster all of them in a single category.
  - **Previous Sales:** To control for potential differences in the effects depending on the popularity of the items, we include the level of online and offline sales before the display.

Considering these variables, we use the following econometric specification:

$$\ln(\mathrm{DRF}_s) = \sum_d \lambda_d \delta_{sd} + \sum_k \mu_k \delta_{sk} + \theta_1 Disc_s + \theta_2 Em_s + \theta_3 Mult_s + \theta_4 Short_s + \theta_5 POnS_s + \theta_6 POffS_s + \varepsilon_s \quad (2)$$

In Eq. (2) DRF_s is the differential response function when a display is made, δ_{sd} is a dummy variable taking the value 1 if SKU s belongs to category d, and δ_{sk} is another dummy variable indicating the banner type (carousel, cover, or regular). $Disc_s$ corresponds to the percentage price discount of product s; Em_s corresponds to the total number of emails sent during the promotional campaign of product s; $Mult_s$ is a dummy variable equal to 1 if the price discount is multi-channel; and $Short_s$ is equal to 1 if the corresponding campaign lasts fewer than 2 days. Finally, the variables $POnS_s$ and $POffS_s$ in our data set are average online and offline sales of that product before it is displayed for the first time. Considering that we have five endogenous variables in the VARX model, we run an independent regression model obtaining a separate set of coefficients for each behavior.

Empirical Analysis

Data Description

In this project we partnered with a multichannel retailer who is competing in the department store market with almost fifty brick and mortar stores, and has an important presence in the online channel. The company has a leading competitive position with a market share close to 25%, and approximately 10% of its sales are made through the online channel.

Our dataset tracks daily online and offline sales for all products displayed on the homepage of the retailer's website in a time span of eight months. In this period, we observed 337 SKUs being displayed. We removed 16 SKUs corresponding to new product introductions that did not have enough sales history before the promotional display, or because they had very infrequent online sales (generating a sparse series of only a few non-zero observations). When selecting products to be displayed on the homepage, the company does not make personalized decisions, but rather shows the same set of products to all visitors. To select the SKUs to be featured on the website the retailer applies some simple filtering criteria such as the attractiveness of the product and the value of the deal, product novelty, or exclusivity with respect to other retailers, and the variety of products to be displayed over time. After applying these criteria, a large number of products are still admissible in a wide set of categories, and the final selection among the acceptable products follows a non-systematic

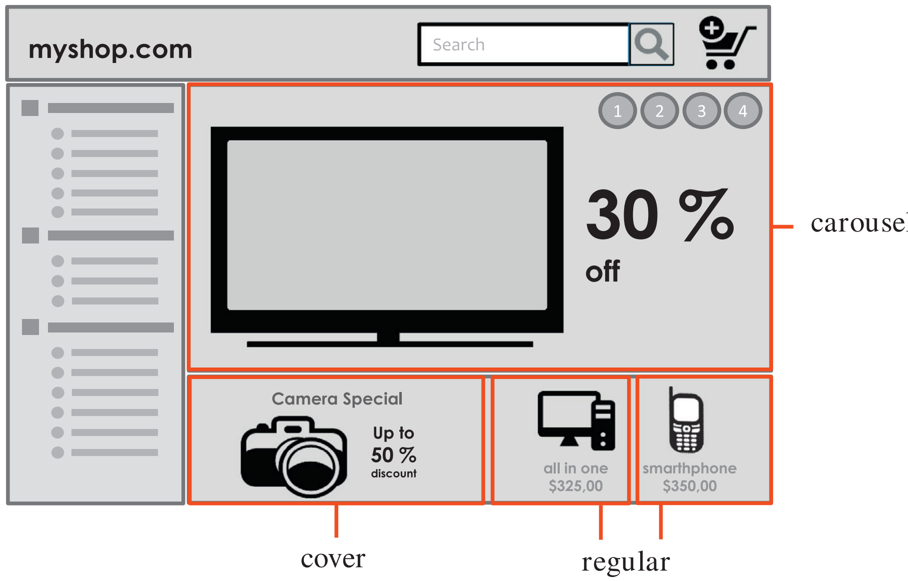

Fig. 2. Schematic representation of website homepage.

procedure. Thus, the SKUs considered in our study correspond to a subset of all possible SKUs. However, as we show in Table 1, there is still ample variation within this group of SKUs in terms of the departments they belong to, the price levels, the duration of the campaigns, and the promotional depth during the campaign. The average time span in which each SKU was promoted on the web site was 5.4 days, with a maximum of 34 days, and a minimum of 1 day. In terms of the prices, while the mean price was $350.74, we observed products ranging from $8.75 to $2,615.37 confirming that the set of products displayed is quite diverse. A large part of the campaigns (81.3%) have price discounts with the average discount being 34%.

We also tracked visits to the product page on the website, and total sales in the corresponding category in both online and offline sales. Table 2 exhibits descriptive statistics for each endogenous variable in our study. We compare periods both with and without house ads for the analyzed products. Table 2 shows that there is a large increment in the daily average of web visits, offline sales, and online sales in the periods being affected by the online display. The resulting dispersion is also large which demonstrates the need for a formal model in which the environmental factors are properly controlled.

In Fig. 3, we display the time series for all five endogenous variables for two representative SKUs. In these plots, we marked the periods associated with house ad campaigns with a dark circle. These examples illustrate that internal display can lead to very different responses. On the left side panel we present the example of a 750 Gb Toshiba external hard drive with a 42% discount. In this case, the internal display is accompanied by a large instantaneous increment in web page visitation and online sales. The impact on offline sales and category sales on the online channel also appears to be positive, but small. Category sales on the offline channel do not seem to be affected by the internal display. On the right side panel, a Samsung 40” LED TV with a 27% discount was displayed during two consecutive days, but there is no obvious impact on any of the performance metrics analyzed. If there is any impact on sales, it is not instantaneous, but, rather, is delayed for a few days after the display.

Before analyzing the data, we need to determine the number of days to consider for each SKU in the sample. For each SKU

Table 1
Descriptive statistics of house ad campaigns at the department level.

Department	N	Prices [USD]			Duration [days]			Discounts	
		Min	Max	Mean	Min	Max	Mean	Incidence	Depth
Home	10	9.22	76.91	29.68	1	27	14.4	20.0%	35.0%
Decoration	4	19.97	61.52	45.37	2	27	8.3	50.0%	56.3%
Sporting goods	21	19.98	1,076.91	314.84	1	34	7.3	81.0%	30.2%
Bedding	44	30.75	1,399.98	532.40	1	27	6.6	90.9%	39.7%
Appliances	157	12.29	2,615.37	377.06	1	34	4.6	85.3%	29.1%
Kitchen and dinning	10	19.98	123.06	59.11	1	2	1.5	60.0%	35.9%
Furniture	40	138.31	1,538.45	565.58	1	27	4.1	100.0%	41.7%
Kids	29	46.14	276.91	138.23	1	1	1	100.0%	38.0%
Perfume	22	8.75	115.37	52.11	1	14	12.2	13.6%	40.3%
Total	337	8.75	2,615.37	350.74	1	34	5.4	81.3%	34.0%

Table 2
Descriptive statistics for the five endogenous variables: web visits, offline sales, online sales, category offline sales, and category online sales.

Campaign period	Statistic	Web visits	Offline sales [USD]	Online sales [USD]	Category offline sales [USD]	Category online sales [USD]
Displayed	Average	190.4	2,226.1	858.12	925,443.2	121,877.6
	Std. dev.	465.0	7,422.9	4,674.8	911,681.2	128,759.6
Not displayed	Average	81.8	1,230.1	258.8	1,009,710.2	131,697.1
	Std. dev.	248.5	5,389.9	1,551.6	1,059,334.6	131,628.5

we observed 243 days of sales. However, for our analysis we considered those visits occurred after the first visit, and before the last visit to the corresponding page on the website. This is justified because, according to the managers, some products are not always available. Additionally, for those products that are promoted for only a few days on the time horizon, we consider a maximum of 90 days before the first display, and 90 days after the last one. Finally, we scaled the series to have all of them of similar magnitude and to avoid ill-conditioned matrices.

The Effects of House Ads on Multichannel Retailing

Before presenting the results from the VARX models, we compare them against two alternative model specifications. The first model (M1) considers only the direct effect of a display on the five endogenous variables for different time lags. This can be conceptualized as a "semi-nested" version of our main model where $\phi_{si} = 0$, but where $DISP_{st}$ is allowed to have a time-varying impact. In the second model (M2) the endogenous variables have only a lagged effect on the same metric, but cross effects are not allowed. This model is similar to the VARX model we propose, but in M2 the matrix of coefficients, ϕ_{si}, is restricted to being diagonal. These two benchmarks are estimated by using a seemingly unrelated regression approach (SUR) in which the number of lags was chosen based on AIC and BIC. In all cases, both criteria coincided in selecting the same number of lags. When comparing our proposed model with these two benchmarks, we found that the VARX is the preferred model in 89.54% of the analyzed campaigns, whereas the M1 and M2 models are preferred in 2.8% and 9.66% of the campaigns, respectively.[6] We conclude that the flexibility of the proposed model provides a substantially better representation of the effect of display on the analyzed variables. Consequently, we focus next on analyzing the results of the proposed VARX model. To determine the lag structure of the vector autoregressive model we ran the model considering different numbers of lags, and then compared them in terms of penalized fit. In our evaluation, we considered models ranging from zero up to 21 lags (three weeks). When analyzing the results, the BIC criterion tends to favor a very limited number of lags while the AIC suggests a slightly larger number of lags. Therefore, we used the Hannan and Quinn criterion (Quinn 1980) that serves as a compromise between a consistent criterion (e.g. BIC) and an asymptotically efficient criterion (e.g., AIC). The left panel of Fig. 4 shows the histogram of the preferred number of lags under the HQ criterion. Consistent with previous research related to in-store display, we saw that the effect of house ads is short-lived for many products, and does not last for more than one or two periods. However, there are a substantial number of products with longer-lasting effects.

The right panel of Fig. 4 shows boxplots of the β coefficients in Eq. (1) that capture the direct effect of online displays on the system. Boxplots are skewed to the positive values for all variables, indicating that, in general, online display tends to increase web visits and sales. Additionally, the direct positive effect tends to be more consistent and greater in magnitude for the promoted SKU (as opposed to the whole category) and for behaviors occurring within the same channel, such as web visits and online sales.

Considering that the VARX model can capture the delayed effects of internal displays on the endogenous variables, it is useful to complement the direct effect with a description of how the internal display affects the evolution of the system. This implies studying not only the direct effect captured by the β parameters, but also the correlation of the endogenous variables. Fig. 5 displays the differential response function computed as the difference between model forecast with and without house ads for the first five days, which captures most of the dynamics for the majority of the series analyzed. Given that the magnitude of the time series can be very different across SKUs (some products can sell five times more than others), we divided the forecast by the average value of the correspondent series. Therefore, we can interpret the ratio as a raw estimate of the percentage of change in the variable due to online display. In each plot, we display the median effect across all promoted SKUs, and the corresponding 90% percentiles.[7] The shaded areas, therefore, depict the heterogeneity of these effects across campaigns with different time lags.

From the differential response functions, we observe that web visits are affected by internal display. In fact, most of the analyzed campaigns exhibit a positive effect with daily increases in visitation that are as high as 15%. Although it is possible to observe a positive impact on multichannel sales without a significant increment in browsing behavior, we believe the mechanism by which sales are affected is caused by customers being motivated to gather more information about the promoted product, taking a step forward in the purchase decision process.

Whereas online sales exhibit a large boost in sales due to internal display, the influence on offline sales tends to be less, but is still quite significant for several of the analyzed campaigns providing additional support to the phenomenon of research shopping. In terms of the timing, the effect on online sales lasts for several days, but the effect on offline sales is noticeably shorter. This is somewhat surprising because it suggests that

[6] On average, the AIC values are 6,633.15, 7,381.72, and 7,223.65 for VARX, M1, and M2, respectively; whereas the BIC values are 6,720.281, 7,447.984, and 7,337.36 for VARX, M1, and M2, respectively.

[7] We report the median and not the mean because there are a few campaigns with very large responses with a greater than 200% increase in the number of web visits and online sales. In general, these items correspond to niche products with small sales volumes.

(a) (b)

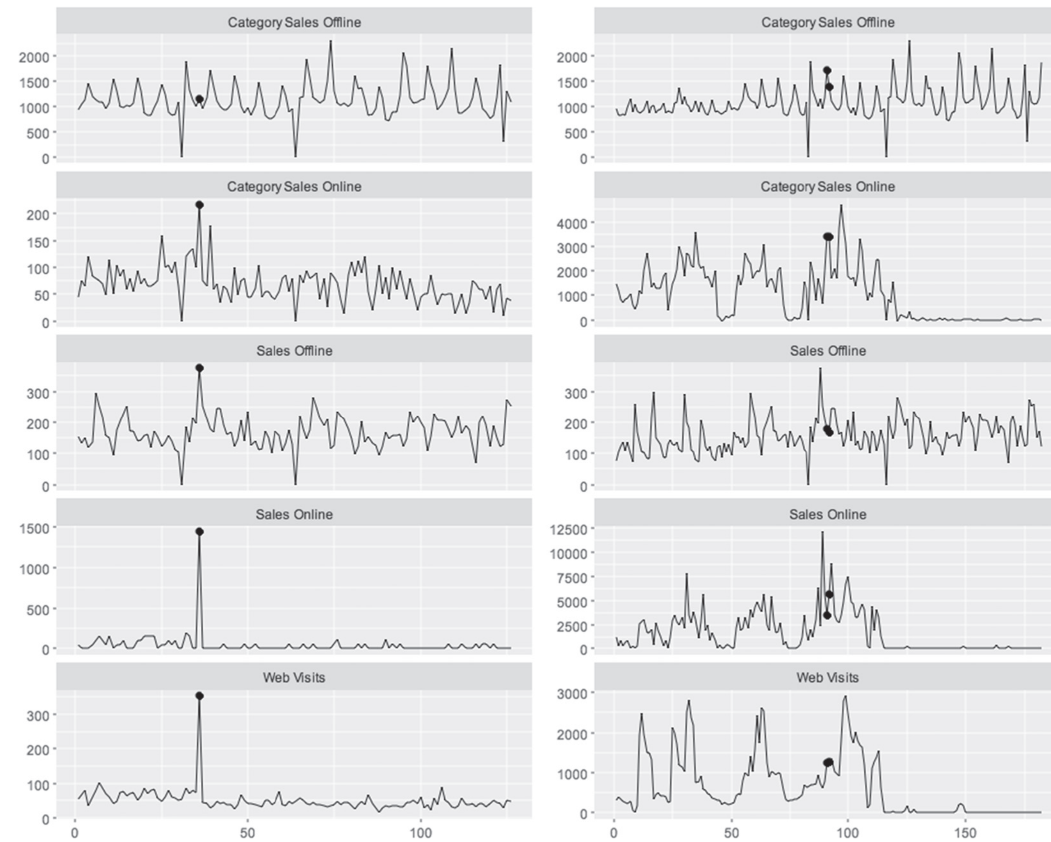

Fig. 3. Example of time series for endogenous variables for two selected SKUs. (a) External hard drive. (b) 40” LED TV. The dark dots represent the periods associated with house ad campaigns.

(a) Histogram of preferred number of lags under Hannan and Quinn Criteria

(b) Boxplot of the coefficient of the exogenous variable on each regression equation

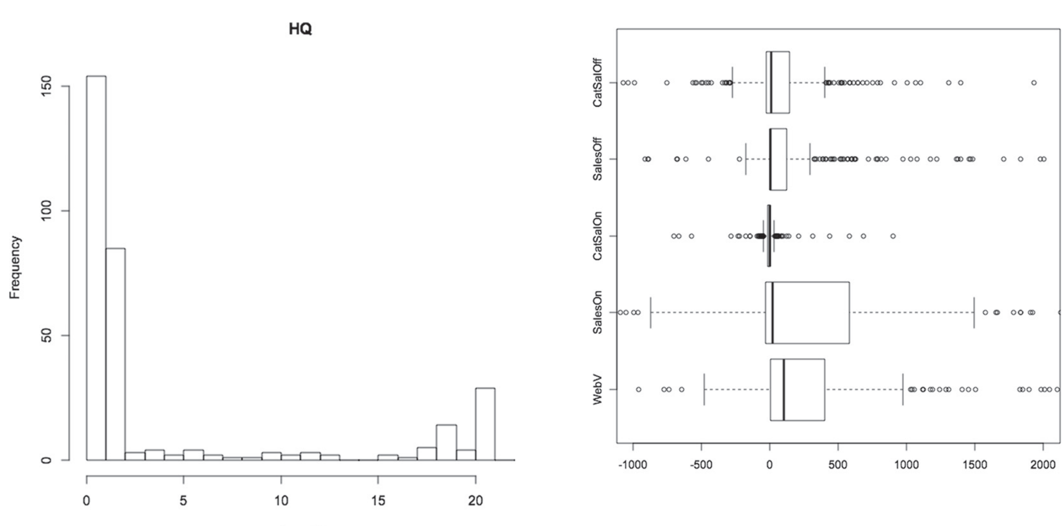

Fig. 4. VARX model results.

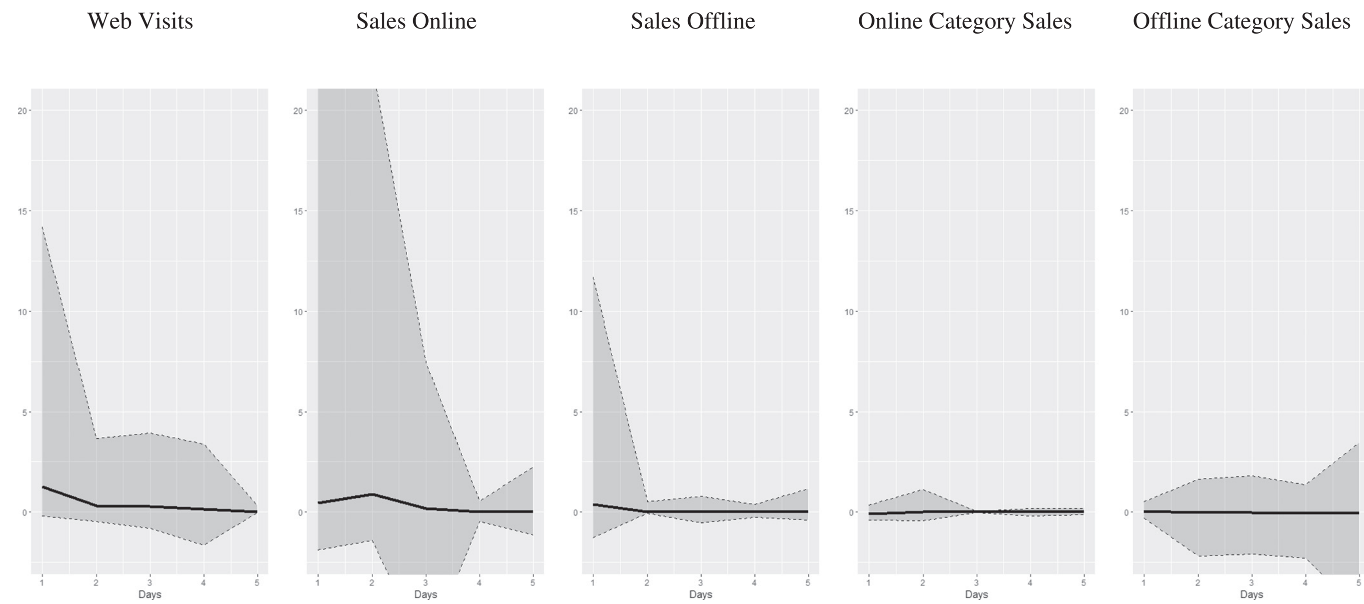

Fig. 5. Differential response function for all endogenous variables. The shaded areas depict the heterogeneity in these effects across campaigns for different time lags.

when customers search online before purchasing offline, they do it predominantly just the day before making the purchase.

Interestingly, we found little evidence that online display has the ability to generate incremental sales at the category level. Category sales on the online channel exhibit almost no effect for most campaigns. However, on the offline channel, there is more variation, with some campaigns having a moderated positive effect that is offset by other campaigns with negative effects. Although some alternative explanations for these behaviors exist, we interpret this pattern simply as a noisy measure of offline category sales.

We have shown that most of the noticeable effects occur in the first few days after the display. However, there might be an important fraction of the promotions generating influence in the system for longer periods. In fact, as we illustrated on the left panel of Fig. 4, our selection of lag structure suggests an effect that may last for several periods for a subset of campaigns. Thus, we consider necessary to derive the total effect by integrating the differential response function for the prediction horizon of 5 days. This measurement of the influence of online display exhibits a large variation across SKUs as shown in Table 3.

Consistent with previous results, we observed a robust impact on the number of web visits and online sales, as well as some evidence of impact on offline sales. We also confirmed that there is no clear effect at the category level. By comparing mean and median, we found that these effects are typically moderated; however, there are some items with a larger impact. To understand what types of products and display characteristics are more positively related to the effect of online display, we regressed the net impact on design variables and product categories following Eq. (2).

Characterizing the Magnitude of the Effect of House Ads

To identify what products are more sensitive to internal display and what type of campaigns are most effective in driving sales, we regressed the net impact of the system on the characteristics of the campaign including the department, the type of display, and the duration of the campaign. We also included the percentage of the discount, if any, the number of emails sent to inform customers, and whether or not the promotion applied to all channels (see Eq. (2)). Considering that we have five different metrics for measuring the net impact of internal display, we ran five separated regressions. Table 4 shows the parameter estimates for all five regressions. Based on the adjusted R^2, the regression models explain a large amount of the variation on the impact. An

Table 3
Net effect of internal promotional display on response functions for the five endogenous variables: web visits, offline sales, online sales, category offline sales, and category online sales.

Response	1st quartile	Median	Mean	3rd quartile
Web visits	7.05	176.62	1,553.96	549.55
Sales online [USD]	−10.63	201.54	3,544.68	1,653.38
Sales offline [USD]	−23.35	312.55	4,431.75	2,858.26
Category sales online [USD]	−21,490.00	−96.51	14,289.92	20,373.85
Category sales offline [USD]	−87,262.46	12,538.90	−14,116.62	259,889.23

Table 4
Results from regression analysis of the net effect of internal promotional display on response functions for the five endogenous variables: web visits, offline sales, online sales, category offline sales, and category online sales.

Parameter		Web visits	Online sales	Category sales online	Offline sales	Category sales offline
λ	Home Deco	0.405	4.553 *	−2.654	−1.585	0.973
	Sporting Goods	1.695	6.141 ***	0.486	1.694	2.911 *
	Bedding	2.377 *	5.391 ***	3.257 **	2.047 ·	3.943 ***
	Appliances	2.707 *	6.839 ***	2.628 *	2.849 *	4.539 ***
	Kitchen & Dinning	1.163	4.414 **	−2.1	1.195	1.026
	Furniture	1.984 ·	5.199 ***	0.254	0.902	1.734
	Kids	1.291	4.848 ***	−1.532	0.459	2.754 *
	Perfume	2.109	4.645 **	−1.546	0.464	2.588 *
μ	Carrousel	−0.615	−0.339	−1.406 ·	−0.187	−1.203
	Regular	2.235 *	−1.511	1.68	2.47 *	2.056 *
	Cover	−1.934 *	−2.337 **	−0.776	−4.333 ***	−2.155 **
θ	Price discount	0.381	0.283	0.44	−0.214	0.323
	No email	0.454	0.087	0.526 ·	0.122	0.378
	Multichannel	−1.03 *	−1.551 **	0.08	−0.574	−0.479
	Short	2.353 **	3.192 ***	−0.701	0.81	−0.872
	Previous online sales	−0.51 ***	−0.317 *	0.099	−0.198	0.325 *
	Previous offline sales	0.217 ***	0.138 *	−0.04	0.098 ·	−0.119 *
	Adj. R^2	0.885	0.898	0.762	0.849	0.910

· $p < 0.1$.
* $p < 0.05$.
** $p < 0.01$.
*** $p < 0.001$.

alternative model without the specific department shows a significantly worse fit with the data, with adjusted R^2 ranging from 0.12 for Offline Category Sales to 0.42 for Web visits, suggesting that a large part of the variation in effectiveness is department specific.

In terms of the department, there are important differences in the responses. Bedding and Appliances are the departments that are affected most by online display, showing positive intercepts for all variables. This implies that, regardless of the design of the campaign, the products in these departments tend to respond favorably, even for offline sales. In these departments, we observe incremental sales even at the category level. When looking at online sales, we find that all departments have positive baselines, revealing that the products of any department have the potential for increasing sales with a well-designed promotional display on the website.

When looking at the effect of the discount, we were surprised to find that the size of the price cut was not significant, which motivated us to explore the effect of the depth of the price promotion further. Fig. 6 displays the increase in online and offline sales as a function of the size of the discount. Although, in general, the tendency is positive, indicating that larger mark-downs are associated with larger incremental sales, there are several other factors that are worth discussing. First, we note that

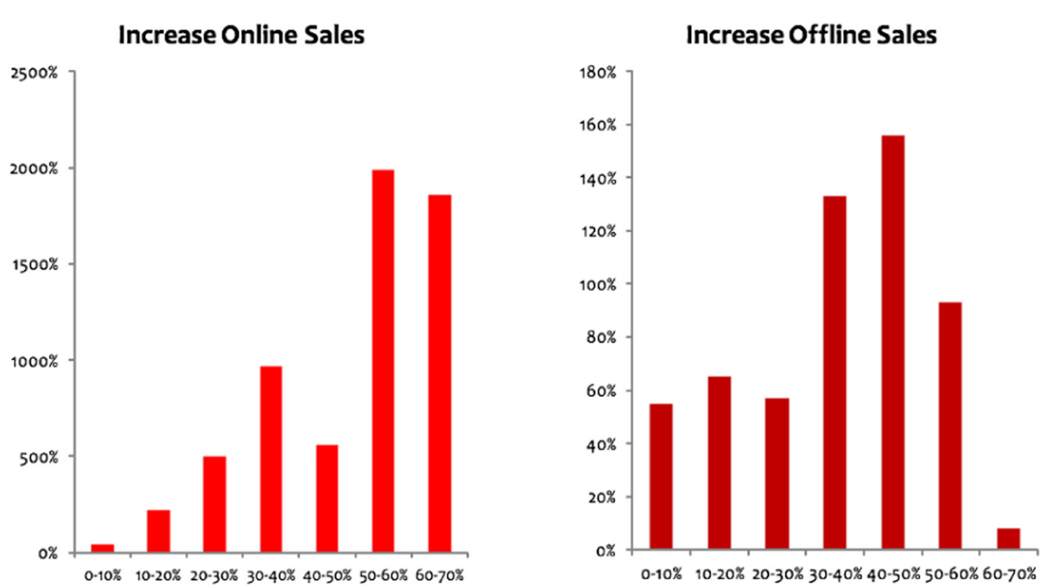

Fig. 6. Online and offline lift on sales. Note that the effects are different in magnitude but are scaled for illustration purposes.

mild discounts on the online sales are not as effective as they are in the offline stores. This is probably due to the low search cost that online shoppers encounter, which reduces the perceived value of a rather routine discount. Second, we observe that large discounts are not very effective in the brick and mortar stores. Interviews with store managers indicate that deep discounts only occur during the end of the season clearances whose occurrence in stores frequently involves shelves with a reduced number of items and limited availability of sizes and colors. Third, we observe that most of the incremental sales occur for discounts of over 50% of the regular price. However, these promotions are relatively scarce. In the dataset, only 10% of the campaigns offered more than a 50% price discount.

Another factor that helps to explain the lack of significance of the size of the discount is that markdown policies tend to be category specific, and, therefore, part of the effect is captured by the department intercepts. For example, the size of the discount in the Kids and Sports departments tends to be larger than the corresponding one for Furniture and Appliances. Indeed, when running a regression with no differential intercept, by department, we find a positive effect for web visits and online sales (with 90% confidence).

In terms of the design of the display, there are some interesting dynamics at play. While regular banners appear to be the most effective in generating traffic, cover is found to be detrimental. Considering that this type of banner is used predominantly to display non-SKU-specific information, the mechanism might divert customer attention resulting in having no direct benefits on sales. One possible explanation that might help to explain the relatively large effect of regular banners on web visits is that those banners are not used for the most popular items. Therefore, regular banners can have a larger potential for growth than more popular products displayed on the most visible regions of the website. This is consistent with what we found when analyzing the effect of previous sales: the larger the number of previous online sales, the smaller the marginal impact on web visits and online sales. Interestingly, we find complementarity between channels from the online sales perspective. That is, it is better to display products with strong sales in the offline stores, but relatively few sales on the online channel. From the offline sales perspective, the effect is weaker because we get significance mainly at the category level. This suggests that brick and mortar store sales are more affected when the displayed product has greater sales online, but lesser offline sales.

Analyzing the nature of the campaign, we found that the length and the channel scope of the promotion only have a direct effect on web visitation and online sales at the SKU level, but not at the category level. Neither did we find any effect on offline sales. Parameter estimates suggest promotions that can be redeemed exclusively on the online channels are more effective in bringing customers and generating online sales. Multichannel campaigns do not perform better on the offline channel. Similarly, short campaigns are more effective in generating e-commerce activity, but the influence does not spread to offline sales.

Finally, we found no evidence of email playing a significant role in helping online display to increase sales. This could be explained by a lack of sophistication in targeting and personalizing the value offering of the emails. In the current implementation, there is little customization of the messages, and most customers receive the same email with all the products that are being displayed online. Moreover, the selection of the customers who receive each communication is very broad, resulting in little dispersion in terms of the number of effective emails sent for each campaign.

Discussion and Future Research

In this research, we explored the effect of internal display of products on web traffic and sales across channels at the SKU and category levels. Unlike most of the previous research related to online display, we focused on promotional display exhibited on the website. This promotional vehicle has goals that are different from those of external display, and therefore it is expected to have a different impact on consumer behavior. Using a flexible-vector, autoregressive approach, we found that the effect of online display is short-lived for the majority of the analyzed campaigns. However, for a fraction of the products, the promotions appear to have a more complex lag structure. In terms of the magnitude of the effect, as expected, we found a larger positive effect in the domain of direct influence. This corresponds to an increase in online sales at the SKU level with little evidence of cross-category spillover. In our analysis, we included more than three hundred house ad campaigns and we found that the impact of internal display is heterogeneous especially on the online channel. We see that there are some items for which internal display generates a positive impact across channels and even at the category level.

The large dispersion motivates conducting a meta-analysis to identify which factors are the most important for this promotional vehicle in influencing the performance of the retailer. To do so, we built aggregated measures of the net effect of online display, and regressed them on some characteristics of its execution, such as the depth of the price cut, the multichannel nature of the promotion, and its duration. Results of the meta-analysis indicate that an important source of variation is related to the department where the promotion is carried out suggesting that the nature of the products being offered plays an important role in the effectiveness of the display. We also found that both the design and layout of the display matter. From the types considered, static regular banners for a single product are the most effective. In terms of the length and scope of the promotions we find that short-term promotions are more effective, and multichannel campaigns decrease the effect online, but, interestingly, they do not increase the effect offline.

Our findings can help retailers to make better advertising decisions. Indeed, our results suggest that firms should conceptualize house ads as short-term promotional vehicles that have effects on the promoted products across channels with very limited spillover effects. Specifically, from an operational perspective, our results can help retailers plan which products should be displayed on the website, and the most effective type of display that should be chosen. Also, these results shed light on how house ads can be coordinated better with other elements of the marketing mix, such as price discounts and direct marketing communications. From a strategic point of view, the evaluation of cross-category and cross-channel effects is important for marketing attribution and budget allocation. Our analysis suggests that even though the largest effect occurs on the online channel itself, there is some noticeable effect in the offline stores. Therefore, displays on the website should be coordinated between the channels. If complete integration is not possible, at least online and offline divisions should consider this cross-channel effect, and incorporate it in their resource allocation processes.

When analyzing the cross-category effects, we found almost no evidence that displaying a given SKU internally has any spillover effect on the whole product category. Therefore, retailers should coordinate with manufacturers and select the identity of the product to display carefully. Moreover, for those banners with a large effect, the retailer can charge the manufacturer for the display, or include the display option when negotiating other promotional spending. Certainly, retailers can monetize their website by selling advertising space to any other firm, even if it does not sell through the retailer. A prominent example of this strategy is Amazon who sells advertising for firms not selling through amazon.com. However, the large impact of the analyzed retailer on online sales advises against this strategy. In summary, we advise companies to monitor the differential response value of each individual campaign closely to evaluate the success of the promotion.

Our investigation provides valuable information for designing more effective house ad campaigns, and for incorporating them in a multichannel retailing strategy. However, our methodology has several limitations. First, we focused on analyzing the effect on a subset of SKUs that were chosen by the company using an ad-hoc criterion. Thus, the extrapolation of these results to products outside this subset should be made with caution. The diversity of products analyzed, however, provides confidence that moderate extrapolations could be made. In terms of the set of variables we used, the availability of offline traffic might provide a more complete description of the potential effect of display in the system. Similarly, controlling for the day of the week, information on competitive products, or product availability could help to refine our characterization of what types of promotions are most effective. Since the department explains a significant part of the variation of effectiveness, the analysis of product characteristics into a more detailed set of attributes could provide additional insights regarding the products that are most likely to be displayed successfully. Among the attributes that we consider interesting to incorporate are the absolute value of the price, the seasonality, and whether the product is hedonic or utilitarian, to mention a few.

We found no evidence that larger discounts have a greater influence in generating sales. We have proposed some plausible explanations for its occurrence, but this phenomenon certainly requires further exploration. Such an analysis could be conducted by introducing a more complex structure to endogenize the magnitude of the price cut, or by adding more data to evaluate the likelihood of alternative mechanisms. Along this line, it would be useful to introduce controlled variations on price changes to identify under what conditions price can effectively play a role in modifying customer behavior. The use of an experimental approach could also help to disentangle the effect of the department, and the potential effect that a more precise selection of customers to whom to send emails that complement the promotional offering would have.

References

Abraham, Magid M. and Leonard M. Lodish (1987), “Promoter: An Automated Promotion Evaluation System,” *Marketing Science*, 6, 2, 101–23.

Ansari, Asim, Skander Essegaier, and Rajeev Kohli (2000), “Internet Recommendation Systems,” *Journal of Marketing Research*, 37, 3, 363–75.

Beatty, Sharon E. and Scott M. Smith (1987), “External Search Effort: An Investigation Across Several Product Categories,” *Journal of Consumer Research*, 14, 1, 83–95.

Besanko, David, Sachin Gupta, and Dipak Jain (1998), “Logit Demand Estimation Under Competitive Pricing Behavior: An Equilibrium Framework,” *Management Science*, 44, 11-part-1, 1533–47.

Braun, Michael and Wendy W. Moe (2013), “Online Display Advertising: Modeling the Effects of Multiple Creatives and Individual Impression Histories,” *Marketing Science*, 32, 5, 753–67.

Breuer, Ralph, Malte Brettel, and Andreas Engelen (2011), “Incorporating Long-term Effects in Determining the Effectiveness of Different Types of Online Advertising,” *Marketing Letters*, 22, 4, 327–40.

Dekimpe, Marnik G. and Dominique M. Hanssens (1999), “Sustained Spending and Persistent Response: A New Look at Long-term Marketing Profitability,” *Journal of Marketing Research*, 397–412.

——— and ——— (2000), “Time-series Models in Marketing: Past, Present and Future,” *International Journal of Research in Marketing*, 17, 2, 183–93.

Dinner, Isaac M., Harald J. Van Heerde, and Scott A. Neslin (2014), “Driving Online and Offline Sales: The Cross-channel Effects of Traditional, Online Display, and Paid Search Advertising,” *Journal of Marketing Research*, 51, 5, 527–45.

Hoban, Paul R. and Randolph E. Bucklin (2015), “Effects of Internet Display Advertising in the Purchase Funnel: Model-based Insights from a Randomized Field Experiment,” *Journal of Marketing Research*, 52, 3, 375–93.

Inman, J. Jeffrey, Leigh McAlister, and Wayne D. Hoyer (1990), “Promotion Signal: Proxy for a Price Cut?” *Journal of Consumer Research*, 17, 1, 74–81.

Joshi, Amit and Dominique M. Hanssens (2010), “The Direct and Indirect Effects of Advertising Spending on Firm Value,” *Journal of Marketing*, 74, 1, 20–33.

Kireyev, Pavel, Koen Pauwels, and Sunil Gupta (2016), “Do Display Ads Influence Search? Attribution and Dynamics in Online Advertising,” *International Journal of Research in Marketing*, 33, 3, 475–90.

Kumar, V. and Robert P. Leone (1988), “Measuring the Effect of Retail Store Promotions on Brand and Store Substitution,” *Journal of Marketing Research*, 178-185.

Liaukonyte, Jura, Thales Teixeira, and Kenneth C. Wilbur (2015), “Television Advertising and Online Shopping,” *Marketing Science*, 34, 3, 311–30.

Lohtia, Ritu, Naveen Donthu, and Edmund K. Hershberger (2003), “The Impact of Content and Design Elements on Banner Advertising Click-through Rates,” *Journal of Advertising Research*, 43, 4, 410–8.

Manchanda, Puneet, Jean-Pierre Dubé, Khim Yong Goh, and Pradeep K. Chintagunta (2006), “The Effect of Banner Advertising on Internet Purchasing,” *Journal of Marketing Research*, 43, 1, 98–108.

Montgomery, Alan L. (1997), “Creating Micro-marketing Pricing Strategies Using Supermarket Scanner Data,” *Marketing Science*, 16, 4, 315–37.

Neslin, Scott A. (2002), *Sales Promotion, Relevant Knowledge Series*. Cambridge, Massachusetts: Marketing Science Institute.

Nijs, Vincent R., Marnik G. Dekimpe, Jan-Benedict E.M. Steenkamp, and Dominique M. Hanssens (2001), “The Category-demand Effects of Price Promotions,” *Marketing Science*, 20, 1, 1–22.

Pauwels, Koen, Imran Currim, Marnik G. Dekimpe, Dominique M. Hanssens, Natalie Mizik, Eric Ghysels, and Prasad Naik (2004), “Modeling Marketing Dynamics by Time Series Econometrics,” *Marketing Letters*, 15, 4, 167–83.

———, Peter S.H. Leeflang, Marije L. Teerling, and K.R. Eelko Huizingh (2011), "Does Online Information Drive Offline Revenues?: Only for Specific Products and Consumer Segments!," *Journal of Retailing*, 87, 1, 1–17.

Quinn, Barry G. (1980), "Order Determination for a Multivariate Autoregression," *Journal of the Royal Statistical Society: Series B: Methodological*, 182–5.

Rutz, Oliver J. and Randolph E. Bucklin (2012), "Does Banner Advertising Affect Browsing for Brands? Clickstream Choice Model Says Yes, for Some," *Quantitative Marketing and Economics*, 10, 2, 231–57.

Senecal, Sylvain and Jacques Nantel (2004), "The Influence of Online Product Recommendations on Consumers' Online Choices," *Journal of Retailing*, 80, 2, 159–69.

Shumway, Robert H. and David S. Stoffer (2006), *Time Series Analysis and its Applications: With R Examples*. Springer Science and Business Media.

Verhoef, Peter C., Scott A. Neslin, and Björn Vroomen (2007), "Multichannel Customer Management: Understanding the Research-shopper Phenomenon," *International Journal of Research in Marketing*, 24, 2, 129–48.

Wiesel, Thorsten, Koen Pauwels, and Joep Arts (2011), "Practice Prize Paper — Marketing's Profit Impact: Quantifying Online and Offline Funnel Progression," *Marketing Science*, 30, 4, 604–11.