

Contents lists available at [ScienceDirect](#)

IJRM

International Journal of Research in Marketing

journal homepage: www.elsevier.com/locate/ijresmar

The effectiveness of different forms of online advertising for purchase conversion in a multiple-channel attribution framework

Evert de Haan^{a,*}, Thorsten Wiesel^{b,c,2}, Koen Pauwels^{d,3}

^a Department of Marketing, Goethe University Frankfurt, Theodor-W.-Adorno-Platz 4, 60323 Frankfurt am Main, Germany

^b Marketing Center Münster, Westfälische Wilhelms-Universität Münster, Am Stadtgraben 13–15, 48143 Münster, Germany

^c Marketing, Faculty of Business and Economics, University of Groningen, The Netherlands

^d Ozyegin University, Nisantepe Mah Orman Sk 13, 34794 Istanbul, Turkey

ARTICLE INFO

Article history:

First received on November 11, 2014 and was under review for 9½ months

Available online 17 December 2015

Keywords:

Marketing attribution
Content integration
Customer initiated
Structural vector autoregression
Restricted impulse response function

ABSTRACT

The Internet has given rise to many new forms of advertising. Scientific studies have focused on individual reactions to specific advertising forms in isolation and have offered little guidance for aggregate-level budget allocation decisions, which are typically based on simple rules. This article compares the long-term effectiveness of nine forms of advertising—seven online and two offline—by means of a structural vector autoregressive model and restricted impulse responses. For five product categories, we investigate how these forms of advertising generate traffic, affect conversion, and contribute to revenue. We find that content-integrated advertising is the most effective form, followed by content-separated advertising and firm-initiated advertising. Although online advertising forms have similar power to drive traffic, content integration dominates content separation in the area of progression toward purchase. Last-click attribution underestimates content-integrated activities and suggests online advertising budget allocations that yield 10%–12% less revenue than the status quo, whereas the model's proposed online advertising budget allocation yields a 21% revenue increase over the status quo. These results highlight the payoffs for companies that integrate content into online media.

© 2015 Elsevier B.V. All rights reserved.

Our success depends on our ability to attract customers on cost-effective terms. We are relying on the offline and online marketing initiatives as a source of traffic to our web site and new customers. If these initiatives are not successful, our results of operations will be adversely affected (Bluefly, 2011).

Online advertising comes in many forms and has grown from \$9.6 billion in 2004 to \$42.8 billion in 2013 (IAB, 2014). Despite this growth, companies remain unsure about the relative effectiveness of online advertising forms (Hanssens, 2009; Sethuraman, Tellis, & Briesch, 2011). A key reason for this knowledge gap is that existing research focuses on tactical decisions within a particular form of advertising. For example, Agarwal, Hosanagar, and Smith (2011) and Skiera and Nabout (2013) show how the position of search engine advertising (SEA) influences the revenues and profits of these advertisements. Blake, Nosko, and

* Corresponding author. Tel.: +49 69 798 33860.

E-mail addresses: dehaan@wiwi.uni-frankfurt.de (E. de Haan), thorsten.wiesel@wiwi.uni-muenster.de (T. Wiesel), koen.pauwels@ozyegin.edu.tr (K. Pauwels).

¹ The majority of the work on this paper has been conducted when Evert de Haan was a doctoral student at the Department of Marketing of the University of Groningen.

² Tel.: +49 251 83 28207.

³ Tel.: +90 216 5649532.

Tadelis (2015) show that brand-keyword paid search is ineffective for eBay, while non-brand keywords affect new and infrequent users. Lambrecht and Tucker (2013) find that the effectiveness of retargeting depends on whether the consumer has well-defined product preferences, whereas Braun and Moe (2013) show that the most effective ad depends on the person's ad impression history. The studies that do investigate several forms of online advertising typically omit offline marketing (e.g. Li & Kannan, 2014). However, firms must actually manage multiple forms of online and offline advertising simultaneously and make strategic budget allocation decisions across advertising forms (e.g. Dekimpe & Hanssens, 2007, Lehmann, 2004). Does sales effectiveness differ systematically by the advertising type?

In addition to sales impact, managers express keen interest in knowing whether online advertising effects differ by product category and stage in the purchase funnel (e.g. Court, Elzinga, Mulder, & Vetvik, 2009, Lecinski, 2011). These questions are largely unanswered by the current literature. Existing studies have examined online advertising effectiveness for only one product category at a time (e.g. Li & Kannan, 2014, Pauwels, Stacey, & Lackman, 2013, Wiesel, Pauwels, & Arts, 2011). It is also difficult to compare individual studies on advertising-form effectiveness because of differences in research setup, product categories, data types, and methodology. For example, some studies focus on short-term effects on sales (e.g. Agarwal et al., 2011, Blake et al., 2015), whereas others focus on a particular funnel stage (e.g. website visits; Ilfield & Winer, 2002). Such studies give little insight into the traffic-to-sales conversion. This is unfortunate because evidence shows that multiple touch points or funnel stages are typically required before purchase (e.g. Frambach, Roest, & Krishnan, 2007, Gensler, Verhoef, & Böhm, 2012) and that these touch points have different effects on purchase likelihood (Braun & Moe, 2013; Li & Kannan, 2014). Finally, many recent studies only consider online advertising (e.g. Blake et al., 2015, Li & Kannan, 2014), while many companies also engage in offline (mass) advertising.

Overall, the existing studies provide fruitful insights into ways to improve the effectiveness of individual forms of online advertising but are of limited use for strategic decision making with respect to budget allocation across multiple forms of advertising. As a result, when making budget allocation decisions, managers tend to rely on gut feelings, trial-and-error, and relatively simple rules, such as last-click methods (Econsultancy, 2012; Jordan, Mahdian, Vassilivtskii, & Vee, 2010). Instead, we offer the conceptual classification of different online advertising forms as content-integrated versus content-separated types and show that this distinction relates to their relative sales effectiveness. Content-integration means that the advertising message is integrated into the editorial content of the website (e.g. a price comparison site), while content-separation implies it is not (e.g. retargeting prospects while they are visiting an unrelated website). While consultants have hinted to the higher effectiveness of content-integrated ads (e.g. Sarner & Herschel, 2008), no research has demonstrated (1) the magnitude of this difference, (2) why last-click methods do not pick up this difference and, relatedly, (3) where in the online funnel this difference in effectiveness materializes.

To the best of our knowledge, no study has compared the effectiveness of a wide range of online forms of advertising throughout the purchase funnel, over multiple product categories, and controlling for offline advertisements. In this study, our objectives are to examine (1) the relative effectiveness of different online advertising forms, (2) how long the effects last, and (3) where in the funnel the effects are strongest (i.e., bringing in more customers, increasing the conversion rate, and/or realizing higher average revenue per order). Furthermore, we investigate whether and to what extent simple rules, such as last-click, are biased as well as where in the funnel and for which types of online advertising such bias occurs. In addressing these questions, we contribute to current literature by analyzing the short- and long-term funnel-stage and revenue effects of a wide range of online advertising forms for five different product categories, while controlling for traditional forms of advertising. The studies with the closest parallels to our own are those of Wiesel et al. (2011); Li and Kannan (2014), and Blake et al. (2015). Our work differs from that of Wiesel et al. in that we include a wider range of online advertising forms, control for more traditional forms of offline advertising (e.g. television, radio), obtain data on several product categories, and analyze those categories using a structural vector autoregressive (SVAR) model. This methodology enables us to test whether or not to include specific backward and forward feedback loops, which are automatically included in standard vector autoregression models (Srinivasan, Vanhuele, & Pauwels, 2010; Trusov, Bucklin, & Pauwels, 2009; Wiesel et al., 2011). Our work differs from that of Li and Kannan and Blake et al. in our cross-category replication, the methodology of aggregate-level time-series analysis, and that we control for offline advertising forms. Time-series patterns of aggregate data allow analysts to observe the temporal sales impact of allocation shifts among online and offline advertising forms simultaneously (Dekimpe & Hanssens, 2007).

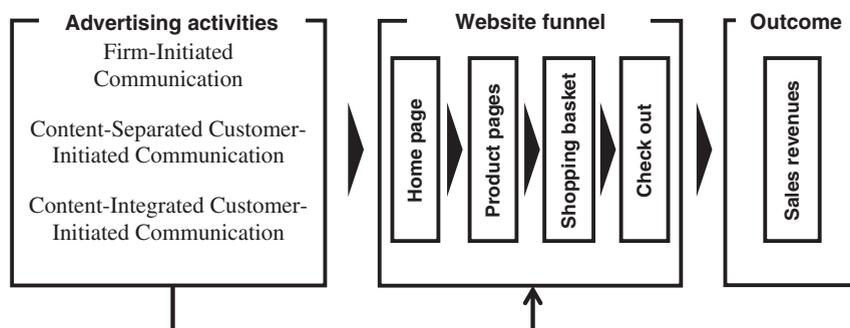


Fig. 1. Website purchase funnel and advertising activities.

Overall, the key contribution of our approach is that we provide guidance to researchers and managers on strategic decision making with respect to aggregate-level budget allocation across multiple forms of advertising and to potential biases when using simple rules, such as last-click methods, for budget allocation decisions. We find that advertising forms in which the message is integrated into the editorial content of a medium are the most effective overall. These so-called content-integrated activities have the same power to drive website traffic as content-separated activities (i.e., activities in which the advertising message is separate from the medium's editorial content). However, the difference is that content-integrated activities are more effective in increasing product-specific, shopping basket, and checkout page sessions. As a result of these funnel benefits, content-integrated activities are substantially more effective in increasing overall revenues. Last-click attribution methods severely underestimate the effectiveness of content-integrated activities and lead to sub-optimal (online) budget allocation, resulting in 10%–12% less revenue than the status quo and 25%–27% less than our proposed allocation.

1. Framework: Advertising effectiveness and funnel progression

Our conceptual framework (see Fig. 1) focuses on the revenue effect of several advertising forms. Advertising generates revenue by inducing a customer to begin the purchase process at a particular firm and helping him or her progress through the firm's purchase funnel until a product or service is finally purchased. Thus, different advertising forms may differ in their total revenue effect, may vary in when the effect occurs, and may affect different stages of the purchase funnel. Consequently, we investigate the differential effectiveness of different advertising forms on three levels. First, we compare *which* advertising form is the most effective, i.e., which form has the highest revenue elasticity. Second, we investigate *when* the effect of different advertising forms occurs, i.e., how long it takes an advertising form to become effective (wear-in) and how long it remains effective (wear-out). Third, we investigate *where* in the purchase funnel the effect occurs, i.e., whether a particular form of advertising attracts more website visitors and/or increases those visitors' conversion probability. By answering these *which*, *when*, and *where* questions, we can create a more nuanced understanding of how different forms of advertising drive bottom-line performance and budget allocation.

Fig. 1 displays a typically observed online funnel. The first stage of the funnel is the front end of the website (home page), on which visitors can explore various product-related pages (second stage). On the product pages, visitors can choose to place products into a shopping basket (third stage), followed by checkout (fourth stage), in which they purchase products and generate revenues.

Advertising may increase revenues in this funnel in several ways. First, an advertising activity can generate additional traffic to the website (i.e., to the home page or product pages). This additional traffic can follow the funnel in the same manner as visitors who have not been exposed to the advertising (i.e., "direct load" visitors). Second, advertising can attract prospects who are more likely to go deeper into the website (i.e., those who have a greater likelihood of moving from product pages to the shopping basket) and therefore are more likely to convert to purchasing. Finally, different advertising activities can influence the amount shoppers spend at checkout. A checkout always generates revenues, but revenues are higher in some cases than in others. By distinguishing among the different stages in the website funnel, we can identify whether advertising forms differ in generating additional traffic, increasing the likelihood of conversion, and/or enhancing revenue per conversion.

In doing so, we distinguish between advertising forms in two ways. The first key distinction in advertising forms is the difference between firm-initiated contacts (FICs) and customer-initiated contacts (CICs). FICs, such as television, radio, and e-mail, focus on pushing the message to the consumer (Shankar & Malthouse, 2007). However, these FICs (defined as any contact with a customer that is initiated by the firm) are increasingly unwanted (Blattberg, Kim, & Neslin, 2008). Conversely, CICs are triggered by (prospective) customers' actions (e.g. Li & Kannan, 2014, Shankar & Malthouse, 2007, Wiesel et al., 2011). CICs include search (organic and paid), price comparison sites, referrals, and retargeting. While firms often pay to show their offer in CICs (e.g. paid search, retargeting), they only enjoy the benefits and (typically) costs of such communication if and when prospective customers take an action that triggers the CIC. The Internet has empowered consumers to interact with companies on their terms, and CICs show great potential and have become a substantial component of firms' marketing (Ghose & Yang, 2009). CIC response rates are projected to be much higher than traditional FICs (Sarnier & Herschel, 2008) in that they require a level of interest from the customer because they are based on customers' own queries and are considered far less intrusive (Shankar & Malthouse, 2007). In addition, the greater effectiveness may be due to the situation a customer is in, including the closeness of the customer to the purchase decision (e.g. Alba et al., 1997). A central idea in marketing is that customers move toward a purchase in a series of stages, including need recognition, information search, evaluation of alternatives, and, ultimately, choice (purchase). While FICs can reach consumers that have not yet recognized a need for the product category, CICs assist consumers who have — as evidenced by them been searching for more information (e.g. search engines), evaluating alternatives (e.g. price comparisons) or getting close to buying the product (retargeting). As a result, CICs should be more (directly) sales effective than FICs. Indeed, previous studies (e.g. Li & Kannan, 2014, Shankar & Malthouse, 2007, Wiesel et al., 2011) show that CICs are more effective in generating revenue, which is the final part of the online purchase funnel (see Fig. 1) and is related to our *which* question. Regarding our *where* question, we expect that CICs attract visitors who are more likely to go deeper into the funnel than FICs. Regarding the *when* question, we expect that because customers must undertake an action that triggers the CIC and CICs are more related to the later stages of the purchase funnel, the effect will be more immediate than that of FICs, for which the firm undertakes the first action and must stimulate the prospective customer; this action is aimed more at the earlier need recognition and information stages and therefore will take relatively more time.

As a second dimension, we investigate the difference between content-separated and content-integrated activities. Our inspiration derives from Russell and Belch's (2005, p. 74) definition of brand placement as "the purposeful incorporation of brands into editorial content." We define "content-integrated" marketing communication as advertising or other promotional (paid or unpaid) activities on third-party websites that are an integral part of the medium's editorial content. For example, consumers specifically access price comparison sites to obtain information on the searched-for items. The prices that different retailers set for a product or service are the actual content that comparison sites offer. In contrast, "content-separated" advertising means that the information, message, or offering of the advertisement is not part of the medium's editorial content. Retargeting messages, traditional banners, and paid search advertising (typically appearing at the right, top, or bottom of the loaded search engine's site) are examples of content-separated activities. A similar rationale for CIC versus FIC applies to content integration versus content separation: by accessing e.g. a price comparison site, the prospect indicates interest in the product/category and thus is further down in the purchase funnel. However, we do not envision a direct mapping from our advertising classification to purchase funnel stages given the widely varying paths possible in the "online decision journey" (e.g. Anderl, 2015, Srinivasan, Rutz, & Pauwels, 2015, Verhoef, Neslin, & Vroomen, 2007). Moreover, two forms of content-separated ads may target different stages, e.g. paid search advertising is most likely to be helpful in the information search stage, while retargeting requires a previous visit to the page displaying the product.

Why could content-integrating advertising be more persuasive and relevant than advertising separated from the site's editorial content? Traditional content-integrated forms of advertising, such as product placement, tend to be more informative and amusing but less irritating than content-separated activities (Tutaj & van Reijmersdal, 2012), which might explain why product placements in films and on television increase brand awareness, positive attitudes, and purchase intentions (e.g. D'Astous & Chartier, 2000, Russell, 2002). This higher customer response and greater purchase intent likely translate into greater effectiveness of content-integrated forms of advertising. All else being equal, we thus expect greater effectiveness of content-integrated advertising activities than of content-separated activities, answering our *which* question. In terms of the *where* question, given that the messages are more informative (Tutaj & van Reijmersdal, 2012) and lead to higher purchase intentions (e.g. D'Astous & Chartier, 2000, Russell, 2002), we expect that content-integrated forms of advertising attract customers who are more likely to go more smoothly through the online funnel than customers who are attracted by content-separated advertising forms. Regarding the *when* question, because they induce positive attitudes and higher purchase intentions, content-integrated activities are probably more likely to attract customers immediately. Combining these two dimensions enables us to classify advertising forms in our current data and to position other established and future forms of advertising.

2. Data

A major European online retailer provided daily data from July 2010 to August 2011, covering 385 days. The data provider is a pure online retailer with no physical stores that mainly sells well-known brands that are also available in other online and physical stores. Both its positioning and budget allocation are similar to other larger online retailers (e.g. Amazon.com, Zalando). For our data provider, branded and product-related SEA yields 63.5% of the total online advertising budget. In comparison, in 2011 SEA had a 46.5% share in online advertising expenditures in the United States (IAB, 2013). E-mail accounts for 1.6% of the retailer's online advertising budget, which is somewhat higher than the .7% share of e-mail in online advertising expenditures in the United States (IAB, 2013). Our data also show that 33.4% of basket sessions lead to an actual sale for our data provider, indicating that in 66.6% of the cases, the customer abandons the shopping basket. This finding is in line with the 60%–70% rate of shopping-basket abandonment indicated by Bronto (2012).

We have category-level⁴ data for five different product categories: (1) fashion; (2) electronics, entertainment, and hardware (EE&H); (3) home & gardening (H&G); (4) sports and leisure (S&L); and (5) beauty and wellness (B&W). We use daily data because expenditures on advertising activities vary daily, making it inappropriate to aggregate to, for example, the weekly level (Tellis & Franses, 2006). Moreover, shoppers can and do progress through website funnel stages within 1 or a few days. Table 1 provides an overview of the operationalization of variables. For each of the advertising variables, we know the expenditure per day per product category. The exceptions are radio, television, and branded SEA (i.e., SEA related to the online retailer itself), for which expenditures are made at the firm level. Referrals (non-paid links from other websites to the retailer's website) represent free advertising to the retailer and are measured in terms of how many people per day visit the retailer's site through a referral. Referrals involve no (direct) costs, which makes this form of online advertising extremely attractive.

For the website funnel, home page sessions (i.e., the number of sessions during which the online retailer's home page was viewed) are not product-category specific; instead, they occur at the firm level. All the other parts of the website funnel, including revenues, are product-category specific; thus, the values vary among the five product categories. A session is measured by the online retailer through its web server logs. The session ends when the customer has been inactive for more than 30 min or has actively left the website.

In addition to the seven online advertising activities and the website funnel stages, we have data on how many people visit the website through organic search (i.e., clicking on the non-advertised results returned by search engines). We also know how many people visit the website through affiliate programs (i.e., third-party websites that promote the retailer's website in exchange for a percentage of the generated revenues). These data are available for each of the five different product categories. Finally, the online

⁴ The data provider performs most of its budget allocation at this broad product-category level, making this level appropriate for examining differences between advertising effectiveness and providing relevant insights for online retailers. If we had product-level data, we could investigate intra-category differences.

Table 1
Variable operationalization.

	Variable	Operationalization
Advertising	E-mail	Daily costs of e-mails consisting of products in the product category (€1 for every 2500 e-mails sent), which are sent to people who have signed up
	SEA product	Daily cost of SEA for product-related search terms (i.e., can be divided by product category)
	SEA branded	Daily cost of SEA related to retailer-related search terms. These are measured at the retailer level and cannot be divided into individual product categories.
	Retargeting	Daily cost of retargeting (i.e., banners that are displayed on a third-party website and show advertisements that are related to a customer's previous visits); costs are for the category that is being advertised/clicked.
	Referrals	Daily number of website sessions started by clicking links on third-party websites (for which the retailer has not paid)
	Portals	Daily cost of portals (i.e., websites that provide links to other websites with specific information). In contrast with referrals, the online retailer pays for inclusion on portal websites.
	Comparison	Daily cost for inclusion on price comparison websites for a specific product category
	Television	Daily cost of television advertising. These are measured at the retailer level and cannot be divided into individual product categories.
	Radio	Daily cost of radio advertising. These are measured at the retailer level and cannot be divided into individual product categories.
	Website funnel stages	Home page
Product		Daily number of sessions that include (one or more) product pages
Basket		Daily number of sessions that include the shopping basket
Outcome	Checkout	Daily number of sessions that include a checkout/payment
	Sales revenues	Daily gross sales revenues
Control	Organic search	Daily number of website sessions started from organic search results
	Affiliates	Daily cost of affiliates (i.e., third-party websites that provide an overview of the retailer's products). Affiliates receive a share of the sales revenues when the customer buys a product through them.

Table 2
Descriptives and correlations of the endogenous variables of interest.

	CV	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1. Sales revenues	.377	1.000													
2. Checkout	.332	.965**	1.000												
3. Basket	.290	.943**	.978**	1.000											
4. Product	.203	.843**	.864**	.897**	1.000										
5. Home page	.263	.857**	.828**	.830**	.840**	1.000									
6. E-mail	.890	.093 ^{ns}	.093 ^{ns}	.109*	-.017 ^{ns}	.097 ^{ns}	1.000								
7. SEA product	.742	.408**	.428**	.484**	.370**	.230**	.138**	1.000							
8. SEA branded	.812	.688**	.643**	.651**	.613**	.580**	.123*	.643**	1.000						
9. Retargeting	.316	.405**	.379**	.380**	.436**	.401**	.039 ^{ns}	.257**	.217**	1.000					
10. Referrals	.348	.453**	.445**	.414**	.557**	.554**	-.001 ^{ns}	.025 ^{ns}	.287**	.326**	1.000				
11. Portals	.183	.559**	.525**	.571**	.618**	.533**	.081 ^{ns}	.365**	.630**	.244**	.182**	1.000			
12. Comparison	.280	.428**	.443**	.520**	.476**	.347**	.103*	.586**	.490**	.178**	.070 ^{ns}	.551**	1.000		
13. Television	1.290	.389**	.354**	.398**	.303**	.298**	.099 ^{ns}	.454**	.367**	.090 ^{ns}	-.046 ^{ns}	.321**	.376**	1.000	
14. Radio	.941	.315**	.315**	.311**	.246**	.304**	-.074 ^{ns}	.136**	.161**	.106*	.199**	.022 ^{ns}	.090 ^{ns}	.237**	1.000

Note: CV = Coefficient of variation (the standard deviation divided by the mean).

** $p < .01$.

* $p < .05$.

^{ns} $p > .05$.

retailer spends money on firm-level (not product-category-level) advertising on television and radio. An external company collects information regarding the daily reach of each commercial, which the online retailer then uses to distribute the cost of each campaign over the days, which is the data we have.

Table 2 provides the coefficient of variation and the correlation matrix for each variable of interest at the firm level.⁵ The coefficients of variation show substantial changes in spending for each advertising form over time, which enables us to examine their temporal effect on sales revenues. This variation over time is highest for radio and television, for which advertising is conducted in pulses (i.e., there are periods of many commercials followed by periods of no commercials), and lowest for portal advertising.⁶ As expected, high correlations exist between sales revenues and funnel metrics, but correlations among advertising forms are moderate. This information enables us to distinguish the separate effects of different advertising forms on performance. The highest correlation among advertising forms is .643 (product- and brand-related SEA). As mentioned, the correlations among the funnel

⁵ Online Appendix A shows the correlations and coefficients of variation at the product-category level.

⁶ For e-mail, the variation over time is much higher at the product-category level than at the firm level, as Online Appendix A shows. A product category may be the topic of an e-mail campaign on one day and then receive no e-mail coverage the next day, when another product category is covered. At the firm level, these product-category differences cancel each other out, resulting in a lower coefficient of variation.

stages are high (e.g. a day with a surge in home page visits also experiences a surge in checkouts) but not perfect, indicating that not all traffic is equally likely to lead to purchase. This enables us to investigate the funnel progression of different forms of advertising. Section 4 details how we address the issue of highly correlated funnel stages.

Fig. 2 shows the time series of overall sales revenues for the five product categories (scaled for confidentiality reasons). The two peak periods, indicated by the shaded bars in the graph, are the two 5-day discount periods, which we control for using five dummy variables per product category. During the rest of the year, there are no product-category-wide discount activities. It is therefore not necessary to control for price for these broad product-category levels. In addition, the assortment size for each product category is stable over the period for which we have data; that is, in no category has there been a major increase or decrease in the amount of products offered. Fig. 2 also shows a clear weekly pattern, which we control for by day-of-the-week dummy variables.

Fig. 3 shows the time series of the total online advertising expenditure for each of the five product categories. The expenditures vary more strongly and have a less clear weekly pattern than the revenues.

The data provider allowed us to share only the relative allocation of marketing actions, not absolute numbers across different product categories. Substantial variation occurs over time with all marketing actions and across different product categories. The analyzed retailer focuses most of its efforts on product-related SEA (32.0% of the total advertising budget and 53.1% of the online advertising budget), followed by television (23.2%) and radio (16.5%). Price comparison sites (7.3%), portals (7.0%), retargeting (6.8%), and branded SEA (6.3%) all receive an approximately equal share of the total advertising budget. The remaining 1.0% of the total budget (i.e., 1.6% of the online budget) is spent on e-mail, which is relatively inexpensive at €1 per 2500 e-mails sent.

In summary, the dataset is rich in advertising forms and in time-series expenditures and revenue measures, which are required to assess the relative effectiveness of advertising forms. Furthermore, information on the online funnel stages and across the five product categories allows us to test whether advertising effectiveness differs by stage and/or by product category.

3. Model-free indications of different advertising effectiveness

For model-free indications about the effectiveness of the different advertising forms in the different product categories, we visualize in Fig. 4 for all five product categories the spending on product-related SEA (the data provider's largest content-separated advertising form in terms of advertising expenditure) and on comparison site advertising (the largest content-integrated advertising form in terms of expenditure). Together these two advertising forms make up about two-thirds of the data provider's online advertising budget.

For fashion, we observe a substantial revenue hike in early March. Both product SEA and comparison site advertising have started increasing before this revenue hike, so there seems to be a relation, but it is unclear which action may be responsible for the revenue hike from eyeballing the data. In contrast, the revenue hike in mid-March for EE&H clearly coincides with the hike in comparison advertising, while hikes in product SEA have little noticeable revenue impact. Next, both H&G and S&L categories show little noticeable effects of either advertising form. Large hikes in product SEA around mid-March do not appear to move revenue, while comparison site advertising remains substantially above revenue in some periods (e.g. after mid-February for H&G and after mid-March for S&L) and below revenue in other periods (e.g. mid-February to mid-March for S&L). Finally, B&W shows mostly decreasing revenue after mid-March, which coincides with the prior day drop in comparison advertising and occurs despite substantial hikes in product SEA.

In sum, these model-free indications reveal that (1) content-integrated comparison site advertising appears to have a stronger revenue relation than content-separated product SEA on average, (2) different categories likely show different advertising effectiveness, and (3) we need a model to separate the (immediate and dynamic) effects of different advertising forms.

4. Methodology

4.1. Modeling approach considerations and justification

Our modeling approach must be able to answer the questions of *which* advertising forms are most effective, *when* the effect takes place (i.e., wear-in and wear-out), and *where* in the website funnel the effect occurs. In addition, it must capture dynamic dependencies among aggregate-level time-series metrics of marketing, funnel stages, and revenues.

Vector autoregressive (VAR) models are designed for aggregate-level time-series data such as advertising expenditures, funnel stages, and revenues (e.g. Srinivasan et al., 2010, Wiesel et al., 2011). Thus, they include the proper dynamics that we want to investigate (i.e., *when* does the effect take place, thus what is the wear-in and wear-out). These models offer a flexible and unified treatment of short-term, long-term, and indirect effects and address endogeneity by incorporating lagged effects and complex feedback loops typical of these types of data (Dekimpe & Hanssens, 2007). Daily advertising, revenues, and funnel data enable us to test for temporal Granger causality (Granger, 1969). The Granger causality of each variable in the model is tested by investigating, for each combination of variables, whether they have Granger-caused each other. For each variable that has a Granger-causal relationship to another variable, the variable is treated as endogenous in the VAR model.

Significant for our research purposes and for answering the *where* question, recent studies have emphasized the importance of allowing customers to skip stages (e.g. the home page) and of marketing to affect conversions throughout the funnel (e.g. Wiesel et al., 2011). Such studies typically have empirical frameworks that allow for all possible feedback loops and routes. However, several routes may be implausible. For example, customers cannot enter at any point in the funnel (e.g. the shopping-basket stage is

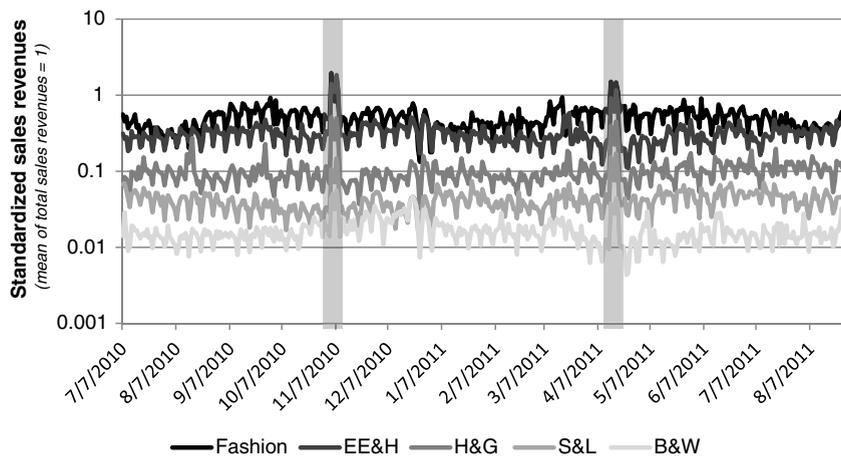


Fig. 2. Standardized sales revenues for the five product categories*. *The vertical bars represent discount sales periods.

necessary before checkout and can only be accessed after viewing at least one product page). Furthermore, it is unlikely that later stages in the funnel significantly drive earlier stages (e.g. that the number of product sessions affects the number of home page sessions) and ultimately generate revenues. Therefore, we need a modeling approach that allows us to test for specific dynamic interactions and to potentially restrict them on the basis of the test results. Testable restrictions can be based on sample-based information (e.g. in-sample test of model fit, out-of-sample forecasting accuracy), managerial judgment or technical limitations (e.g. the website is designed so that one stage must follow another), and/or marketing theory (e.g. brand awareness comes before brand liking in the hierarchy of effects).

SVARs are “specifically designed to supplement sample-based information with managerial judgment and/or marketing theory” (Dekimpe & Hanssens, 2000, p. 185). Although SVARs have received considerable attention in the field of economics (e.g. Litterman, 1983, Pesaran & Smith, 1998), they are not commonly used in marketing (Pauwels, 2004; Pauwels et al., 2005). In our context, SVARs enable us to investigate whether we can restrict parameters associated with specific routes to zero, which greatly reduces the number of parameters to be estimated in the final model and reduces multicollinearity. Owing to the high correlation among funnel stages (Table 2), including all stages in each equation would severely hinder the model’s ability to distinguish between the effects per stage. By restricting impossible routes, we substantially reduce this multicollinearity.

4.2. Model estimation steps

We estimate a SVAR model for each of the five product categories. Our approach extends the persistence modeling steps in Dekimpe and Hanssens (1999), as shown in Table 3. In the first step, we determine which variables are Granger-caused by other variables. This analysis determines which variables we must include as endogenous in the system. When two or more variables are endogenous, we must capture their dynamic interactions in a dynamic system model (Lütkepohl & Krätzig, 2009). We include variables that do not have a Granger-causal relationship to any other variable as exogenous in the model.

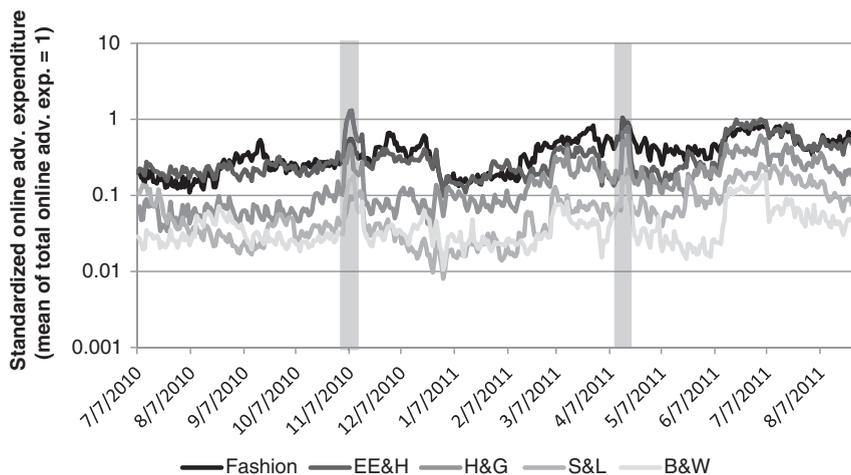


Fig. 3. Standardized total online advertising expenditure for the five product categories*. *The vertical bars represent discount sales periods.

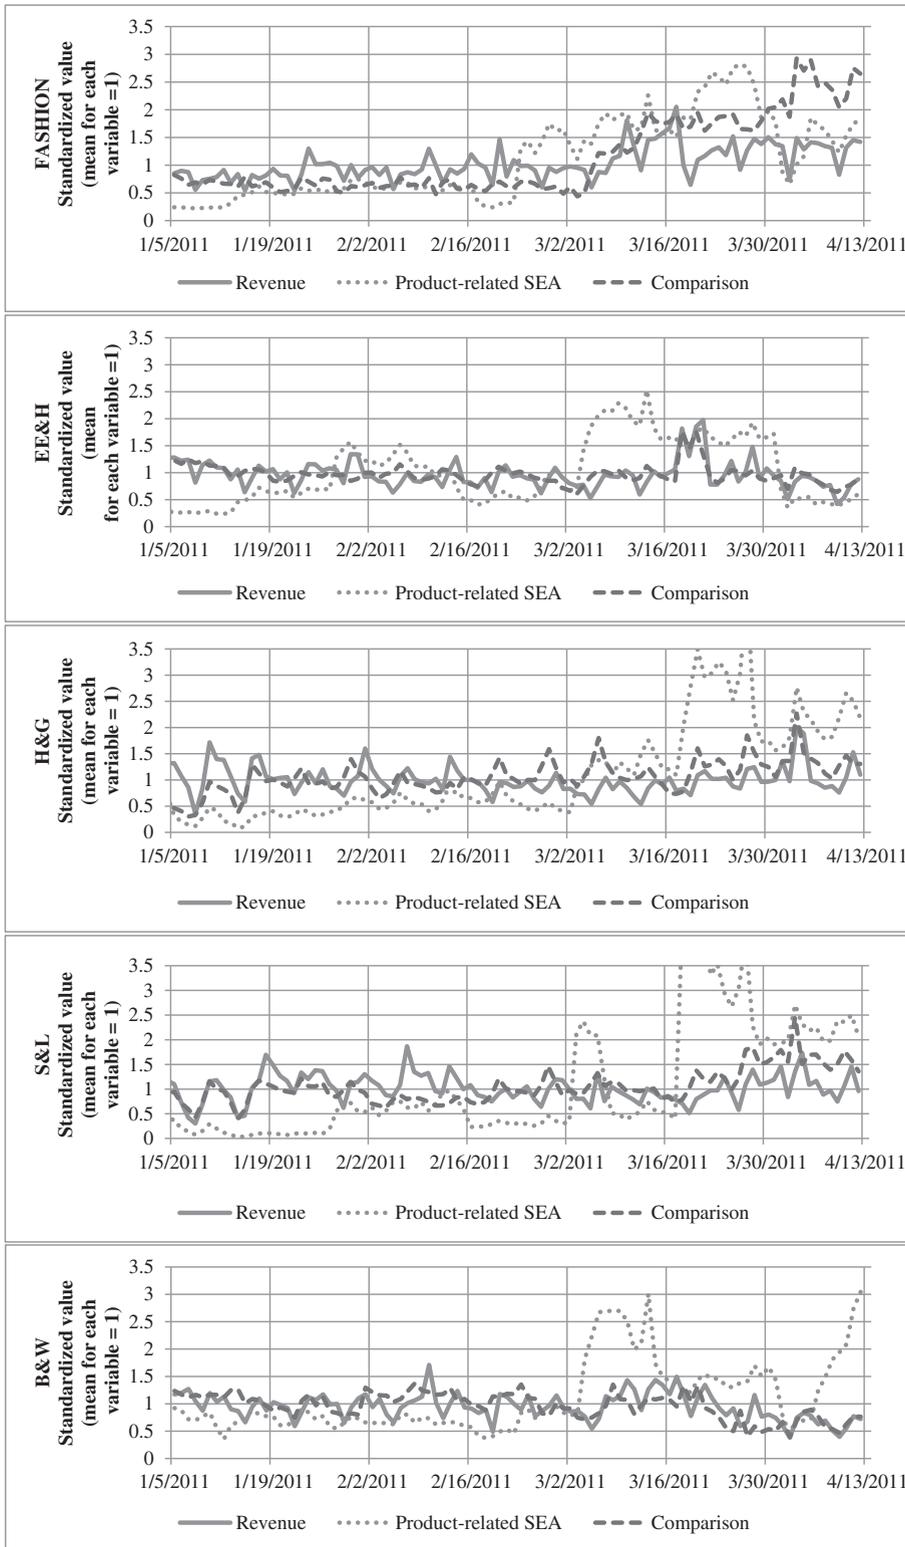


Fig. 4. Category revenues, product-related SEA, and comparison site ad spending.

Table 3
Overview of methodological steps.

Methodological step	Relevant literature	Research question
1. Granger-causality test	Granger (1969)	Which variables temporally Granger-cause which other variables?
2. Unit root and cointegration Augmented Dickey–Fuller test, Kwiatkowski–Phillips–Schmidt–Shin test, cointegration analysis	Dickey and Fuller (1979), Kwiatkowski, Phillips, Schmidt, and Shin (1992), Johansen (1991)	Are variables stationary, or do they evolve? Do evolving variables share a long-term equilibrium?
3. Dynamic interaction models Vector autoregression model, VAR in differences, vector error correction, s tructural VAR (SVAR) predictive validity	Dekimpe and Hanssens (1999), Pauwels, Srinivasan, and Franses (2007), this study, Leeflang et al. (2015)	How do performance and advertising interact in the long and short run, accounting for unit roots and cointegration and excluding specific paths? How does the model predict out-of-sample?
4. Policy simulation analysis Unrestricted impulse response functions (IRF) Restricted IRF	Sims (1980) Pauwels (2004)	What is the net dynamic impact of a change in advertising on performance? What is the dynamic impact of a change in one advertising activity, keeping other advertising constant?

In the second step, we perform unit-root and cointegration tests to determine the form of the variables to include in the models. Variables with a unit root are included in first differences. If some of these variables are co-integrated, however, a vector error correction-type model is appropriate (Lütkepohl & Krätzig, 2009).

With the outcomes of these first two steps, we estimate in the third step the appropriate dynamic models with the appropriate number of lags, determined by the Bayesian information criterion (BIC). As exogenous control variables, we include in each equation an intercept, a linear time trend, six day-of-the-week dummies, quarter-of-the-year dummies, and five dummies for discount periods. Eq. (1) shows the general model form in matrix notation:

$$Y_t = A + B_t Y_t + \sum_{i=1}^p \Phi_i Y_{t-i} + \psi X_t + \Sigma_t, \quad t = 1, 2, \dots, T, \tag{1}$$

where Y_t is the $m \times 1$ vector of endogenous variables (where m is the number of endogenous variables), A is an $m \times 1$ vector of intercepts, X_t is the vector of exogenous control variables, Σ_t is the $m \times 1$ matrix of residuals, B is the $m \times m$ matrix of (restricted) parameters with direct effects, and Φ_i are the $m \times m$ matrices of (restricted) parameters for lag i . The BIC determines the value of p (the amount of lags to be included in the model). We log-transform all variables so that the parameters can be interpreted as elasticities. Because on some days the expenditure on some of the variables is zero, we use Manchanda, Rossi, and Chintagunta's (2004) procedure—that is, in all cases we use $\ln(y_t + 1)$ instead of $\ln(y_t)$. Provided that the elasticities have a value between 0 and 1, an advantage of this procedure is that we can capture the diminishing returns of advertising (Leeflang, Wieringa, Bijmolt, & Pauwels, 2015).

For Eq. (1)'s B and Φ_i , respectively the direct (same-period) and lagged effects, we can restrict some of the parameters to zero to capture specific relationships (Pauwels, 2004). These restrictions transform the VAR model into an SVAR model, as shown in Eq. (2) (for clarity of exposition, we group the advertising and endogenous control variables).

$$\begin{bmatrix} Advertising_t \\ Homepage_t \\ Product_t \\ Basket_t \\ Checkout_t \\ Revenue_t \end{bmatrix} = A + \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ \beta_{21}^0 & \beta_{22}^0 & \beta_{23}^0 & \beta_{24}^0 & \beta_{25}^0 & \beta_{26}^0 \\ \beta_{31}^0 & \beta_{32}^0 & 0 & \beta_{34}^0 & \beta_{35}^0 & \beta_{36}^0 \\ \beta_{41}^0 & \beta_{42}^0 & \beta_{43}^0 & 0 & \beta_{45}^0 & \beta_{46}^0 \\ \beta_{51}^0 & \beta_{52}^0 & \beta_{53}^0 & \beta_{54}^0 & 0 & \beta_{56}^0 \\ \beta_{61}^0 & \beta_{62}^0 & \beta_{63}^0 & \beta_{64}^0 & \beta_{65}^0 & 0 \end{bmatrix} \begin{bmatrix} Advertising_{t-1} \\ Homepage_{t-1} \\ Product_{t-1} \\ Basket_{t-1} \\ Checkout_{t-1} \\ Revenue_{t-1} \end{bmatrix} + \sum_{i=1}^p \begin{bmatrix} \beta_{11}^i & \beta_{12}^i & \beta_{13}^i & \beta_{14}^i & \beta_{15}^i & \beta_{16}^i \\ \beta_{21}^i & \beta_{22}^i & \beta_{23}^i & \beta_{24}^i & \beta_{25}^i & \beta_{26}^i \\ \beta_{31}^i & \beta_{32}^i & \beta_{33}^i & \beta_{34}^i & \beta_{35}^i & \beta_{36}^i \\ \beta_{41}^i & \beta_{42}^i & \beta_{43}^i & \beta_{44}^i & \beta_{45}^i & \beta_{46}^i \\ \beta_{51}^i & \beta_{52}^i & \beta_{53}^i & \beta_{54}^i & \beta_{55}^i & \beta_{56}^i \\ \beta_{61}^i & \beta_{62}^i & \beta_{63}^i & \beta_{64}^i & \beta_{65}^i & \beta_{66}^i \end{bmatrix} \begin{bmatrix} Advertising_{t-i} \\ Homepage_{t-i} \\ Product_{t-i} \\ Basket_{t-i} \\ Checkout_{t-i} \\ Revenue_{t-i} \end{bmatrix} + \psi X_t + \Sigma_t \tag{2}$$

The first restriction (Block 1) implies that (changes in) the metrics of website funnel stages cannot influence (changes in) the metrics of earlier funnel stages on the same day. For example, when we observe an increase in both shopping-basket and checkout sessions, we attribute the latter to the former, not vice versa. The second restriction (Block 2) implies that consumers do not typically skip stages in the website funnel on the same day. For example, an increase in the number of product-page sessions does not directly increase checkout sessions; this effect is fully mediated by an increase in shopping-basket sessions. Finally, restrictions 3 and 4 generalize the backward restriction (Block 3) and the skipping restriction (Block 4), respectively, to dynamic effects on subsequent days. If all restrictions hold, the resulting model is much simpler than the standard VAR structure that allows for all contemporaneous (same-day) and dynamic effects among all the endogenous variables. We test for each of the restrictions in two ways. First, the BIC yields an in-sample test of model fit. Second, we assess out-of-sample forecasting accuracy, which is crucial

to both managers and scholars (Franses, 2005). In general, more complicated models may perform better in-sample but tend to perform worse out-of-sample (Leeflang et al., 2015). This is especially important given the large number of endogenous variables and the dynamic model structure. We split the dataset into the usual two-thirds for estimation and one-third for validation and calculate the mean absolute percentage error (MAPE) in the validation sample (Leeflang et al., 2015). Our final model is a model with restrictions that results in the lowest BIC and MAPE. To account for heteroskedasticity and cross-equation contemporaneous correlation in the error terms, we estimate the models using seemingly unrelated regression (Pauwels, 2004; Zellner, 1962).

We tested whether certain advertising forms enhance each other by including each possible interaction between the advertising variables one-by-one in the models but found no significant improvements in terms of the BIC. While some studies have found significant interactions (e.g. Li & Kannan, 2014, for a U.S. lodging chain), others report no significant interactions (e.g. Wiesel et al., 2011, for a Dutch furniture seller). In addition to differences in application settings, a likely reason for the absence of interactions is that managers use different advertising forms to reach different types of consumers, as confirmed by our data provider. In addition, our data provider does not specifically design campaigns to gain synergies.

4.3. Restricted policy simulation and the website funnel effects of advertising

To answer our *which* question (i.e., which advertising form has the highest revenue elasticity), we perform in the fourth step (Table 3) a restricted policy simulation (Pauwels, 2004). This analysis enables us to interpret the effects of an impulse on a marketing action by itself (i.e., without it changing the other endogenous advertising variables in the model). The restricted impulse response function of this simulation provides us with advertising elasticities as traditionally interpreted (i.e., the effect of a 1% increase in advertising, while keeping everything else constant). A key benefit of restricted impulse response functions is that they do not require estimation of a new, restricted model (which would bias the coefficient estimates); instead, they only restrict the simulation of performance effects, based on the estimates from the full model's estimation (Pauwels, 2004).

For ease of illustration, after estimating the effects per advertising form for the five product categories, we translate the advertising elasticities back to the overall firm level. To do this, we take the weighted mean of the elasticities using the precision of the parameter estimates (i.e., 1 divided by the parameter's standard error). This calculation assumes that a 1% increase in advertising variable *i* is allocated over different product categories, similar to how the budget is normally allocated over those categories. We translate the effect of advertising on each stage in the website funnel in the same way. To reach an optimal online advertising budget allocation, we follow the procedure of Danaher and Dagger (2013) and Naik and Raman (2003). For the budget allocation we can only use variables which are measured in terms of monetary expenditure. The variable 'referrals', which does not directly involve monetary expenditures but is measured in the amount of referrals (i.e., the daily number of website sessions started by clicking on non-paid for links on third-party websites), is thus not included in the budget allocation. Further details on the budget allocation procedure appears in Online Appendix B.

5. Results

5.1. Model selection

The Granger-causality tests show that all nine advertising variables, organic searches, affiliates, four stages of the website funnel, and revenues are significantly Granger-caused by at least one other variable. This finding indicates that we must treat them all as endogenous. The augmented Dickey–Fuller and Kwiatkowski–Phillips–Schmidt–Shin tests indicate unit roots in organic searches for all product categories and in expenditures on comparison websites and SEA for the S&L category. Because there are three variables with a unit root for S&L, we perform a cointegration test (Johansen, 1991), which shows that the variables are not cointegrated. Consequently, we estimate all five SVAR models with the variables with a unit root in first differences. For all five models, the BIC underscores the use of one lag, which gives us 43 parameters to estimate in the revenue equation for each product category. Because each product category has 385 observations, the resulting 8.95 observation-to-parameter ratio is considered sufficient (Hanssens, Parsons, & Schultz, 2003, p. 185).

Table 4 presents the parameter restriction tests. The BIC and MAPE indicate that both current and lagged backward and skipping paths can be excluded. This empirical evidence is consistent with technical limitations on funnel progression and managerial intuition. Therefore, we proceed with a much simpler SVAR model—namely, the model in which all four blocks in Eq. (2) are set to zero.

Table 4
Restricted model tests.

	Block 1	Block 2	Block 3	Block 4	MAPE	BIC
Model 1	✓	✓	✓	✓	4.24%	113,168.2
Model 2	X	✓	✓	✓	4.10%	113,098.0
Model 3	X	✓	X	✓	4.10%	113,055.6
Model 4	X	X	✓	✓	3.97%	113,046.0
Model 5	X	X	X	X	3.89%	112,971.1

On the product-category level, this final model explains 99% of the variation in revenues for fashion and S&L, 96% for EE&H, 94% for H&G, and 92% for B&W. The explained variance is lower for earlier funnel stages in the purchase funnel (e.g. 87%, 90%, 89%, 77%, and 84% for product page stage for the respective product categories). Key drivers of the high explanatory power of revenues are the advertising variables, which together explain 78% of the variance in revenues for fashion, 85% for S&L, 77% for EE&H, 63% for H&G, and 61% for B&W. This result reflects the richness of our data for both online and offline forms of advertising. The sales promotion dummies, lagged sales revenue, current and lagged checkout sessions, and day-of-week dummies capture much of the remaining variance in sales revenues. The previous stage in the purchase funnel, as well as the lag variable of the current funnel stage, explains much of the variance in the current stage, as can also be derived from the correlations table (Table 2). However, the advertising variables significantly add explanatory power in funnel stage progression on top of the other variables for each funnel equation, as we verified in a series of F-tests ($p < .01$ in all cases). This indicates that our data are appropriate to investigate funnel progression and, thus, to answer our *where* question.

5.2. Revenue elasticities and the purchase funnel effects of each ad form

From the selected SVAR, we calculate the restricted impulse responses of revenue to a 1% change in the advertising variables. Table 5 captures the key results with respect to long-term advertising elasticity (the cumulative effect) and wear-in and wear-out for each advertising activity. A wear-in (wear-out) of 1 means that the first (last) significant effect of advertising on revenue occurs on the day of the impulse, a wear-in (wear-out) of 2 means that the first (last) significant effect occurs on the second day, and so forth. As Table 5 shows, the wear-out of most channels is within 1 or 2 days. This finding is in line with the results of other studies (e.g. Li & Kannan, 2014).

Of the 45 (9×5) long-term advertising elasticities, 29 (64.4%) are significantly different from zero. This result compares favorably with the 57% significant advertising elasticities that Sethuraman et al. (2011) find in their meta-analyses. Furthermore, 55.2% of the significant advertising elasticities (16 of 29) are (also) significant after the first day. This finding indicates that examining only the same-day impact would lead to biased results.

In terms of sign and size, all significant advertising elasticities are positive and vary between .003 and .215. Of the 29 significant advertising elasticities, one (i.e., product-related SEA for S&L) is from an advertising variable included in first differences. This elasticity represents the performance effect of permanently increasing marketing spending by 1%. Given the difference in interpretation, we focus the remainder of our analysis on the 28 elasticities of marketing variables included in levels and exclude variables in first differences.

5.3. Product category differences

Table 5 shows consistency across product categories in the relative size and ordering of the elasticities. We do, however, observe some differences. First, EE&H revenues do not increase with any FIC but do increase with price comparison site advertising. Such products are rather expensive search goods, and thus induce price comparison (Nelson, 1970). Second, fashion and S&L see no impact of price comparison – but retargeting works in these categories. A likely reason is that consumers have well-defined preferences in the early stages of their journey, which makes retargeting more effective (Lambrecht & Tucker, 2013). Finally, product-related SEA increases revenues for fashion and EE&H, which are two of the most popular product categories purchased online (Nanji, 2013). This limits the research-shopper phenomenon (Verhoef et al., 2007) of consumers who search for products online but wind up buying them in bricks-and-mortar stores—a likely scenario for H&G (home and gardening), for example.

Table 5
Long-term advertising elasticities and timing in days per product category.

		Fashion			EE&H			H&G			S&L			B&W		
		Advert. elasticity	Wear-in	Wear-out												
FIC	Television	n.s.	–	–	n.s.	–	–	.006	1	2	.004	2	2	n.s.	–	–
	Radio	.003	2	2	n.s.	–	–	.007	1	1	.003	1	1	n.s.	–	–
	E-mail	.005	2	2	n.s.	–	–	n.s.	–	–	.008	2	3	.014	1	1
CIC separated	SEA product	.084	2	2	.048	1	1	n.s.	–	–	.032*	2	2	n.s.	–	–
	SEA branded	.032	1	2	.036	1	2	.042	1	1	n.s.	–	–	.051	1	1
CIC integrated	Retargeting	.051	1	1	n.s.	–	–	n.s.	–	–	.040	1	1	n.s.	–	–
	Referrals	.129	1	1	.184	1	2	.175	1	2	.204	1	3	.129	1	2
	Portals	.182	1	2	.173	1	2	.119	1	1	.123	1	1	.075	1	1
	Comparison	n.s.	–	–	.154	1	1	n.s.	–	–	n.s.*	–	–	.215	1	2

n.s. = not significant.

* The variable is in first differences; therefore, the elasticity must be interpreted as follows: “a 1% permanent change in variable x will give a temporally increase of y in sales revenues,” where x is the advertising variable and y is the elasticity.

Table 6
Advertising elasticities split by activity type and product category.

Activity type	Example	Mean elasticity	Weighted elasticity	Fashion elasticity	EE&H elasticity	H&G elasticity	S&L elasticity	B&W elasticity
FICs	Television, radio, e-mails	.003 ^b	.003 ^b	.003	.000	.004	.005	.005
CICs (content-separated)	SEA, retargeting	.027 ^b	.030 ^b	.056	.028	.014	.020	.017
CICs (content-integrated)	Portals, price comparison, referrals	.133 ^a	.130 ^a	.104	.171	.098	.163	.139

^a Significantly ($p < .01$) different from the other two activities.

^b Significantly ($p < .05$) different from the other two activities.

5.4. Answering the “which, when, and where” questions

The key aim of our study is to provide guidance to researchers and managers on strategic decision making with respect to budget allocation across multiple forms of advertising. To do so, we investigate the differential effectiveness of different advertising forms by answering three questions: *which* advertising forms are most effective, *when* does the effect take place, and *where* does this effect occur in the purchase funnel? In addition, we illustrate the biases of simple rules, such as last-click methods, for budget allocation decisions. Next, we examine the implications of our results, guided by the three main questions and the discussion of biases when using last-click methods.

5.4.1. Answering the “which” question

To answer the *which* question (i.e., which advertising form is the most effective), we use the two recent classifications of advertising forms introduced previously to provide further insights into our results. First, we classify forms of advertising as CICs and FICs, following Wiesel et al. (2011) and Li and Kannan (2014). Significant elasticities occur 53.3% of the time for FICs (8 of 15), consistent with John Wanamaker's century-old notion that half the money spent on advertising is wasted. For CICs, 70.0% of elasticities (21 of 30) are significant, which is in line with previous studies that show that CICs are more effective than FICs (e.g. Li & Kannan, 2014, Sarner & Herschel, 2008, Wiesel et al., 2011). When splitting CICs into content-integrated and content-separated categories, as indicated in Table 5, we find that 60% (9 of 15) of the elasticities for content-separated CICs are significant. In contrast, 80% (12 of 15) of the elasticities for content-integrated CICs are significant.

We obtain a clear effectiveness distinction along the dimension of firm- versus customer-initiated forms of advertising. The mean advertising elasticity is .003 for FICs and .080 for CICs (.003 and .074 respectively if we weigh the elasticities by precision). In other words, in our dataset, CICs are 26.7 times (or 24.7 times if we look at the weighted mean) more effective than FICs. This finding is consistent with previous studies that report much higher revenue elasticity for CICs than for FICs (Sarner & Herschel, 2008; Wiesel et al., 2011) and in line with the theoretic expectations in Section 1. The likely reason is that prospective customers are more attentive to information that is directly relevant to what they seek; however, firm-initiated advertising often reaches them at the wrong time and with a sub-optimal message. Our CIC–FIC effectiveness ratio of 26.7 and 24.7 is somewhat higher than Sarner and Herschel's (2008) ratio of 15 but far lower than Wiesel et al.'s (2011) ratio of 400 between SEA and FICs. A reason for the lower effectiveness of CICs could be that the company Wiesel et al. (2011) examined is relatively small, whereas ours is already established and therefore has less to gain from paying for SEA (Blake et al., 2015).

When we split CICs into content-separated and content-integrated, as in Table 6, we find that content-separated activities are on average 9 (unweighted) or 10 (weighted) times more effective than FICs, while content-integrated activities are 44.3 (unweighted) or 43.3 (weighted) times more effective in generating revenues than FICs. Our research is the first to show this higher revenue effectiveness of content-integrated than content-separated CICs. However, this finding is consistent with previous studies on the effect of traditional content-integrated advertising on purchase intentions (e.g. D'Astous & Chartier, 2000; Russell, 2002) and the positive perceptions customers have of these messages (Tutaj & van Reijmersdal, 2012). It is also in line with the theoretic expectations in Section 1. The differences among the three groups – FICs, content-separated CICs, and content-integrated CICs – are all statistically significant, as indicated in Table 6. Furthermore, the findings hold for all five product categories. This gives us confidence that our key substantive findings are not product-category specific, although there are too few observations per category to test for significance – a key avenue for future research.

Regarding face validity, we can compare our model's results with the model-free indications in Section 3. As to the main message of this paper, content-integrated comparison site advertising indeed has a larger revenue elasticity than content-separated product SEA; consistent with Fig. 4. Could this be an artifact of the methodology? Unlikely, because our model reveals that comparison site advertising has no significant revenue impact in 3 out of 5 categories. In fashion, the largest category, the model attributes the revenue increase to product SEA instead of to comparison site advertising (recall from Section 3 that both advertising forms increased prior to the revenue hike). In the second largest category, EE&H, both product SEA and comparison site advertising increase revenue, but the elasticity of comparison site advertising is three times as large, consistent with our observations in Section 3. And in both H&G and S&L categories, the model results confirm our model-free intuition that neither advertising form substantially increases revenue.⁷ Finally, the B&W category results show a substantially larger revenue elasticity for comparison site versus product SEA expenditures, consistent with our observations in Fig. 4.

⁷ Note from Table 5 that product-related SEA does have an impact on revenue in the S&L category, but since this variable is measured in first differences it means that a 1% permanent change in product-related SEA expenditures will give a temporally increase of .032% in product category revenues, which is extremely limited.

5.4.2. Answering the “when” question

Immediate (same-day) revenue effects are significant for FICs in only 37.5% (3 of 8) of the cases. For content-separated CICs, this is true in 77.8% (7 of 9) of the cases, whereas it is always the case for content-integrated CICs (12 of 12 cases). Therefore, the effect of content-integrated CICs occurs quickly, whereas it is lagged for FICs. This observation makes sense because CICs are triggered by a consumer action indicating (at least some) interest in the category, and these differences are in line with the expectations in Section 1 of this article. For wear-out, we do not observe much difference between the groups. In approximately half the cases, there is (still) a revenue effect after the first day.

5.4.3. Answering the “where” question

Finally, after investigating which advertising form is most effective and when this effect takes place, we investigate *where* in the purchase funnel the effect takes place (e.g. whether an advertising form brings more visitors to a website, or whether it increases visitors' conversion probability). By grouping according to advertising category, we obtain the mean funnel elasticities for FICs, content-separated CICs, and content-integrated CICs. Fig. 5 illustrates these funnel elasticities and shows that both content-integrated and content-separated CICs outperform FICs in generating home page traffic. However, the two types of CICs do not differ significantly from each other in this respect. The likely reason is that many content-integrated CICs link to a specific product page that is of interest to both the third party and the customer who clicked on the link. This interest should translate into higher conversion through the funnel. We find that all later website funnel metrics show a significantly greater elasticity for content-integrated activities. Content-integrated CICs tend to distinguish themselves from other forms of advertising in the later stages of the funnel. They are more effective in attracting people to product-specific pages; from there, visitors move more smoothly through the website funnel, thus producing higher total revenues. These findings stay the same when we weight the elasticities from all stages based on the precision with which they are estimated.

6. Managerial implications for budget allocation and the last-click bias

To allocate budgets across advertising forms, managers need to know how different advertising forms contribute to revenues. In this section we therefore optimize the budget allocation in order to (1) compare our results with simple rules like last-click (online allocation only) and (2) to make budget allocation recommendations (online and offline). One of the most popular methods in practice, used in more than half the cases, is the last-click attribution method (Econsultancy, 2012). We compare our SVAR with two last-click methods. First, the “last-click/same-session” method examines which advertisement form started the session. If the customer started the session without clicking on an advertisement link (i.e., by typing in the URL him- or herself), the session is called a “direct load.” Direct load can be interpreted as goodwill or (brand) awareness developed by the retailer; it can also come from traffic from offline advertising forms (e.g. television, radio) for which no “click” occurs. Second, the “last-click/7-days” method includes additional historical information. When a session is started by typing in a retailer's URL, the customer's sessions for up to the previous 7 days are investigated. When a previous session during those 7 days was started by clicking on an advertisement, that advertisement receives the credit for the current session. If there was no previous session during the previous 7 days or if all the sessions during the previous 7 days were also direct load, the current session is also labeled “direct load” (i.e., not driven by advertising). With both of these last-click methods, no credit is given to earlier advertisements, the effect of seasonality, or baseline revenues. Furthermore, the impact of offline advertising and other forms of advertising that generate visitors but not (direct) clicks is ignored. For these reasons, studies have found that last-click attribution is insufficient and delivers biased insights (Jordan et al., 2010; Li & Kannan, 2014).

How does the data provider allocate its online advertising budget? Table 7 compares the current online budget allocation with the allocation derived from the two last-click attribution methods.⁸ Note that the current allocation does not strictly follow any of the last-click attribution methods. When confronted with this difference, the retailer's managers indicated that they adjust the last-click recommendations using their own judgment and use trial-and-error to search for better ways to allocate their budgets. Specifically, management believed that last-click methods over-estimated the power of e-mails and branded SEA, while they underestimated the power of product-related SEA, retargeting, and price comparison sites. As a result, management spends less on e-mails (1.6% instead of the recommended 48%–50%) and branded SEA (10% instead of 16%–20%). However, managers are unsure about whether they have gone far enough or even too far in these adjustments. This practice underscores managers' need for guidance in allocating firm advertising budgets.

To offer guidance, we calculate the online budget allocation suggested by our SVAR model using the procedure of Danaher and Dagger (2013) and Naik and Raman (2003). The last column of Table 7 shows our suggested online budget allocation. For both e-mails and branded SEA, our analysis agrees with management intuition and recommends allocating budgets substantially lower than those offered by last-click methods. Especially the power of e-mail is over-estimated by last-click attribution. Our rationale is that most customers only react to e-mail when they have decided what and where to buy based on other sources (including other advertising forms) and are reminded about their purchase intentions by the e-mail, which provides easy click access to the product. Indeed, Li and Kannan (2014) find that customers use some channels more frequently to navigate to a website even if they have previously visited the website through other channels. However, in their study, the use of a last-click approach leads to under-estimation of the e-mail effect. The difference between our findings and theirs might be due to the nature of e-mail

⁸ The retailer directly provided us with the data and outcomes of last-click/same-session and last-click/7-days. Last-click methods are only available for online advertising forms; thus, we concentrate on those in our analysis. Because referrals do not involve any costs, we do not include referrals in the budget allocation.

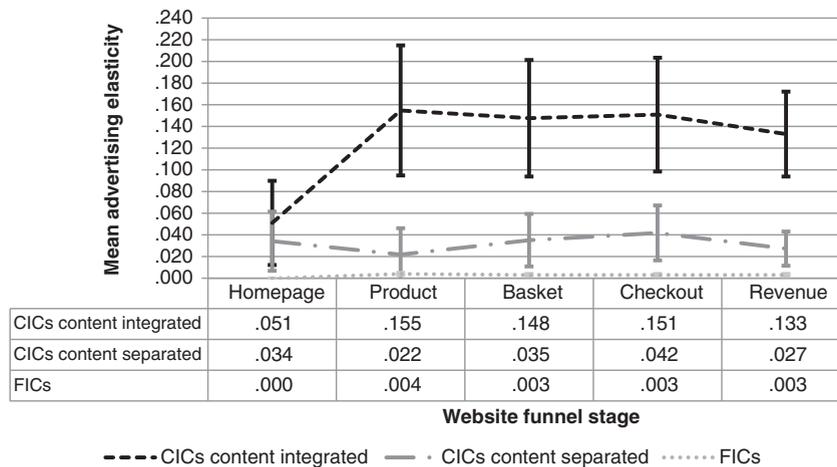


Fig. 5. Advertising elasticities for funnel stages.

(i.e., permission-based) targeting and content focus (e.g. customer acquisition vs. cross-selling). Our online data provider increased the amount of e-mails during various periods to obtain more clicks but did not observe a substantial change in website visits or revenues. This trial-and-error approach to increasing the number of e-mails, which the retailer already considered high, is closer to the optimal method identified by the SVAR than it is to last-click methods.

The last-click/7-days method is close to the SVAR results for product-related SEA and retargeting, two advertising forms for which our analysis opposes the management's current allocation. Both advertising forms are content-separated: their message is not part of the website's editorial content. Each of the three analytic methods consistently shows that such advertising forms are less effective in terms of sales than the data provider believes them to be. In contrast, only our SVAR analysis shows the power of content-integrated advertising: we recommend that the data provider allocate 43% to portals and 29% to price comparison sites. Last-click methods underestimate these effects, plausibly because these forms of advertising generate few direct visits but a large amount of brand and product exposure, which is an important part of content-integrated CICs. Our analysis supports the data provider's decision to upwardly adjust the last-click method allocation but also shows that this adjustment does not go far enough for the content-integrated advertising forms. This is consistent with research on anchoring and adjustment (Wilson, Houston, Etling, & Brekke, 1996) and suggests that last-click methods can severely bias allocations even when management adjusts the anchor in the right direction.

What happens to retailer revenues under each of the four allocations in Table 7? Using our model estimates, the retailer would increase revenues by 21.01% by switching from its current allocation to our recommended allocation. In contrast, switching to the last-click/same-session (last-click/7-day) method would decrease revenues by 9.68% (12.05%) from the current allocation. Accordingly, the allocation from last-click/same-session (last-click/7-day) yields 25.36% (27.32%) less revenues than our recommended allocation. The retailer is unlikely to achieve our recommended allocation immediately because CICs (e.g. SEA expenditures) can only be increased when enough additional customers use these channels (Wiesel et al., 2011). Thus, in practice, we recommend that managers *move toward* the SVAR's suggested optimal allocation, e.g. in a field experiment.

Finally, we run the allocation procedure including the two offline advertising forms, television and radio. Table 8 reports the results of this procedure. One of the downsides of last-click procedures, which the table shows, is that these methods cannot garner direct traffic through these channels and therefore cannot allocate sales to these advertising forms. As a result, with last-click methods the current budget for offline advertising forms is redistributed over the online advertising forms.

An important caveat to Table 8 relates to interpreting the very low allocation to offline advertising forms, based on their low revenue elasticity in our model estimated on daily data. Such low elasticities can be expected if the retailer uses offline advertising mainly to build and retain brand awareness, perception, and goodwill rather than to drive immediate sales (e.g. Hanssens, 2009; Sethuraman et al.,

Table 7

Comparison of online budget allocation and sales revenues: current, SVAR, and last-click methods.

	Current allocation	Last-click/same-session	Last- click/7-days	SVAR
E-mail	1.58%	47.71%	49.81%	1.71%
SEA product	53.12%	10.06%	12.45%	10.49%
SEA branded	10.38%	20.11%	15.81%	10.14%
Retargeting	11.21%	5.18%	7.31%	5.74%
Portals	11.56%	10.92%	8.89%	42.64%
Comparison	12.14%	6.03%	5.73%	29.28%
Revenue forecast (relative to current)	100.00%	90.32%	87.95%	121.01%
Revenue forecast (relative to SVAR)	82.64%	74.64%	72.68%	100.00%

Table 8

Comparison of total budget allocation and sales revenues: current, SVAR, and last-click methods.

	Current allocation	Last-click/same-session	Last-click/7-days	SVAR
E-mail	.96%	47.71%	49.81%	1.68%
SEA product	32.03%	10.06%	12.45%	10.33%
SEA branded	6.26%	20.11%	15.81%	9.99%
Retargeting	6.76%	5.18%	7.31%	5.66%
Portals	6.97%	10.92%	8.89%	42.02%
Comparison	7.32%	6.03%	5.73%	28.86%
Television	23.23%	n/a	n/a	.65%
Radio	16.47%	n/a	n/a	.81%
Revenue forecast (relative to current)	100.00%	101.78%	99.11%	139.14%
Revenue forecast (relative to SVAR)	71.87%	73.14%	71.23%	100.00%

2011). In the online world, such goodwill may translate into the direct load; i.e., visits that were started by typing in the website's URL instead of using (online) advertising.⁹ This online retailer enjoys an average direct load of 54.9% of all online visits (ranging from 37.8% for EE&H to 63.0% for fashion), the remaining 45.1% of the visits come to the website via one of the online advertising channels. This direct load is higher than the 39.7% share in offline advertising out of the total advertising spending. Because we do not have data from the start of the online retailer, we cannot establish how much of this direct load can be attributed to past online advertising, past offline advertising or other factors such as word-of-mouth or a personal positive experience.

A second caveat regarding budget allocation (in both Tables 7 and 8) is our assumption that prospects lost by cutting one advertising form are not (fully) recovered by finding their way directly to the retailer. To fully analyze this 'recovery' issue, we would need individual-level data in controlled field experiments, as executed for eBay by Blake et al. (2015). They find that cutting branded keyword SEA has no sales impact for eBay: prospects that type in the retailer's name in a search engine are likely determined to check out that retailer and thus will use other ways (such as natural search clicks) to go to its page. Cutting non-branded keyword SEA does lower sales, but only for new and infrequent customers. Note that several of Blake et al.'s (2015) findings may help understand ours. Content-separated CICs such as SEA may be mostly spent on the regular customers (as shown for eBay in Blake et al., 2015), while content-integrated CICs such as referrals, portals and price comparison sites may feature more new and infrequent customers. Likewise, retargeting focuses on customers that have already looked at the page of the specific product advertised. Following our and Blake et al.'s (2015) rationale, customers who are really still interested in purchasing that product would find a way to return – the others are unlikely to be convinced by an ad at a time they are checking out an unrelated website. Would Blake et al.'s (2015) zero sales impact on branded keyword SEA for eBay apply for all CIC advertising at our retailer? Only a field experiment would answer that question, but we find it unlikely for at least two reasons: First, previous field experiments at companies other than eBay have shown a substantial impact of CICs (Montaguti, Neslin, & Valentini, 2015; Wiesel et al., 2011). Second, even looking at just two forms of SEA (branded and non-branded), Blake et al. (2015) find different sales effectiveness. They also state that SEA is "unlikely to play a persuasive role" in contrast to "display ads" to which their findings may not apply. Thus, the recovery issue appears to differ on a continuum by advertising form – consistent with our rationale. In sum, distinguishing these effects in field experiments for different advertising forms and different companies represents a promising avenue for further research.

7. Limitations and further research

This study has several limitations that provide fruitful avenues for further research. Although we examine multiple product categories, the data come from one online retailer. Our substantive results are similar across all five categories, indicating broader insight into the effectiveness of the advertising forms considered. However, further studies are necessary to generalize our findings to other contexts. For example, research could compare more and less established retailers to analyze whether the effectiveness of a form of advertising is dependent on how well the firm is known, as suggested by Blake et al. (2015) and Li and Kannan (2014) and confirmed by Demirci, Pauwels, Srinivasan, and Yildirim (2014). Field experiments are important to verify the causality and policy implications (Blake et al., 2015; Li & Kannan, 2014). Research could also investigate advertising effectiveness in other sectors. In addition, we do not know the advertising expenditures of manufacturers and the retailer's competitors. To the extent that these third-party activities are correlated with the focal retailer's advertising, this omission could affect our effect size estimates for the retailer's own advertising activities. This lack of data also constitutes a limitation that our data provider and several other (online) retailers must address in practice when making decisions about budget allocation. As previous studies (e.g. Pauwels, 2004, Srinivasan et al., 2010) have demonstrated, it is straightforward to include such third-party marketing activity in (structural) VAR models and to generate insights into dynamic competition patterns. Another limitation is that we examine a relatively short period. We checked for but did not find evidence of time-varying parameters. Other studies (e.g. Sethuraman et al., 2011) have shown that the effectiveness of advertising decreases over time. Furthermore, although we use daily data and divide these into product categories, the data are still at an aggregate level. Within the same day, we do not know which stages precede

⁹ We thank an anonymous reviewer for this suggestion.

which other stages. Another limitation of aggregate data is their inability to account for heterogeneity in response in advertising (Manchanda, Dubé, Goh, & Chintagunta, 2006). Further research could extend our findings by distinguishing advertising form effectiveness for different customers.

8. Conclusion

In summary, the data presented here are rich with respect to use of online advertising forms, purchase funnel metrics, the number of product categories, and offline advertising forms. Thus, we are able to provide a more nuanced understanding of the differential effectiveness of different forms of online advertising, which, to the best of our knowledge, has previously been unavailable. In doing so, we answer three questions about effectiveness: *Which* advertising forms are most effective? *When* does the effect take place? and *where* in the purchase funnel does the effect take place? Such an understanding of advertising form effectiveness is necessary to provide guidance on strategic decision making with respect to budget allocation across multiple forms of advertising. This is particularly important because managers tend to rely on trial-and-error or simple rules, such as last-click methods, because they lack knowledge and/or better approaches. Our study shows the substantial biases of this method. We hope that our research not only stimulates additional research but also encourages the future use of marketing science approaches to budget allocation decisions rather than relying on gut feelings, trial-and-error, or overly simplistic heuristics.

Acknowledgment

The authors thank the anonymous online retailer who has provided the data and other relevant insights to conduct this research. The authors thank Dominique M. Hanssens, Raoul Kübler, Bernd Skiera, Shuba Srinivasan, and Peter C. Verhoef for their valuable suggestions on this manuscript. The authors also thank the participants of the 2012 INFORMS Marketing Science Conference in Boston, the participants of the 2014 Theory + Practice in Marketing Conference in Evanston, the participants of the 2015 EMAC Conference in Leuven, the two anonymous reviewers, and the guest editor, Werner Reinartz, for their valuable comments.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <http://dx.doi.org/10.1016/j.ijresmar.2015.12.001>.

References

- Agarwal, A., Hosanagar, K., & Smith, M. D. (2011). Location, location, location: An analysis of profitability of position in online advertising markets. *Journal of Marketing Research*, 48(6), 1057–1073.
- Alba, J., Lynch, J., Weitz, B., Janiszewski, C., Lutz, R., Sawyer, A., & Wood, S. (1997). Interactive home shopping: Consumer, retailer, and manufacturer incentive to participate in electronic marketplaces. *Journal of Marketing*, 61(3), 38–53.
- Anderl, E. (2015). Mapping the customer journey: A graph-based framework for attribution modeling. *McKinsey Dissertation Award presentation, EMAC 2015, Leuven, Belgium*.
- Blake, T., Nosko, C., & Tadelis, S. (2015). Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica*, 83(1), 155–174 01.
- Blattberg, R. C., Kim, B.-D., & Neslin, S. A. (2008). *Database marketing: analyzing and managing customers*. Heidelberg, Germany: Springer.
- Bluefly (2011). Annual report. December <http://www.annualreports.com/Company/6201> (accessed March 16, 2015).
- Braun, M., & Moe, W. W. (2013). Online display advertising: Modeling the effects of multiple creatives and individual impression histories. *Marketing Science*, 32(5), 753–767.
- Bronto (2012). From abandon to conversion: Why shoppers abandon carts and what merchants can do about it. http://mkto.brnto.com/AbandonedCartWP_AbandonedCart.html (accessed March 18, 2015).
- Court, D., Elzinga, D., Mulder, S., & Vetvik, O. J. (2009). The consumer decision journey. *The McKinsey Quarterly*, 3, 96–107.
- Danaher, P. J., & Dagger, T. S. (2013). Comparing the relative effectiveness of advertising channels: A case study of a multimedia blitz campaign. *Journal of Marketing Research*, 50(4), 517–534.
- D'Astous, A., & Chartier, F. (2000). A study of the factors affecting consumer evaluations and memory of product placement in movies. *Journal of Current Issues and Research in Advertising*, 22(2), 31–40.
- Dekimpe, M. G., & Hanssens, D. M. (1999). Sustained spending and persistent response: A new look at long-term marketing profitability. *Journal of Marketing Research*, 36(4), 397–412.
- Dekimpe, M. G., & Hanssens, D. M. (2000). Time-series models in marketing: Past, present and future. *International Journal of Research in Marketing*, 17(2–3), 183–193.
- Dekimpe, M. G., & Hanssens, D. M. (2007). Advertising response models. In G. J. Tellis, & T. Ambler (Eds.), *Handbook of advertising* (pp. 247–263). London: Sage.
- Demirci, C., Pauwels, K., Srinivasan, S., & Yildirim, G. (2014). "Conditions for owned, paid and earned media impact and synergy". *Marketing Science Institute, report 14* (pp. 101).
- Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366), 427–431.
- Econsultancy (2012). Quarterly digital intelligence briefing: Making sense of marketing attribution. November <http://econsultancy.com/nl/reports/quarterly-digital-intelligence-briefing-making-sense-of-marketing-attribution> (accessed March 17, 2015).
- Frambach, R. T., Roest, H. C. A., & Krishnan, T. V. (2007). The impact of consumer Internet experience on channel preference and usage intentions across the different stages of the buying process. *Journal of Interactive Marketing*, 21(2), 26–41.
- Franses, P. H. (2005). On the use of econometric models for policy simulation in marketing. *Journal of Marketing Research*, 42(1), 4–14.
- Gensler, S., Verhoef, P. C., & Böhm, M. (2012). Understanding consumers' multichannel choices across the different stages of the buying process. *Marketing Letters*, 23(4), 987–1003.
- Ghose, A., & Yang, S. (2009). An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Science*, 55(10), 1605–1622.
- Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3), 424–438.
- Hanssens, D. M. (2009). *Empirical generalizations about marketing impact: What we have learned from academic research*. Cambridge, MA: Marketing Science Institute.
- Hanssens, D. M., Parsons, L. J., & Schultz, R. L. (2003). *Market response models: Econometric and time series analysis*. 12. Springer Science & Business: Media.
- IAB (2013). 2012 Internet advertising revenue full-year report. April <http://www.iab.net/media/file/IABInternetAdvertisingRevenueReportFY2012POSTED.pdf>

- IAB (2014). 2013 Internet advertising revenue full-year report. April http://www.iab.net/media/file/IAB_Internet_Advertising_Revenue_Report_FY_2013.pdf
- Ilfield, J. S., & Winer, R. S. (2002). Generating web site traffic: An empirical analysis of web site visitation behavior. *Journal of Advertising Research*, 42(5), 49–61.
- Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. *Econometrica*, 59(6), 1551–1580.
- Jordan, P., Mahdian, M., Vassilivtskii, S., & Vee, E. (2010). The multiple attribution problem in pay-per-conversion advertising. *Lecture Notes in Computer Science*, 6982, 31–43.
- Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics*, 54(1–3), 159–178.
- Lambrecht, A., & Tucker, C. (2013). When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research*, 50(5), 561–576.
- Lecinski, J. (2011). Winning the zero moment of truth. <http://www.thinkwithgoogle.com/research-studies/2011-winning-zmot-ebook.html>
- Leeflang, P. S. H., Wieringa, J. E., Bijmolt, T. H. A., & Pauwels, K. H. (2015). *Modeling markets: Analyzing marketing phenomena and improving marketing decision making*. New York: Springer-Verlag.
- Lehmann, D. R. (2004). Metrics for making marketing matter. *Journal of Marketing*, 68(4), 73–75.
- Li, H., & Kannan, P. K. (2014). Attributing conversions in a multichannel online marketing environment: An empirical model and a field experiment. *Journal of Marketing Research*, 51(1), 40–56.
- Litterman, R. B. (1983). A random walk, Markov model for the distribution of time series. *Journal of Business & Economic Statistics*, 1(2), 169–173.
- Lütkepohl, H., & Krätzig, M. (2009). *Applied time series econometrics*. Cambridge, UK: Cambridge University Press.
- Manchanda, P., Dubé, J.-P., Goh, K. Y., & Chintagunta, P. K. (2006). The effect of banner advertising on Internet purchasing. *Journal of Marketing Research*, 43(1), 98–108.
- Manchanda, P., Rossi, P. E., & Chintagunta, P. K. (2004). Response modeling with nonrandom marketing-mix variables. *Journal of Marketing Research*, 41(4), 467–478.
- Montaguti, E., Neslin, S. A., & Valentini, S. (2015). “Can marketing campaigns induce multichannel buying and more profitable customers? A field experiment.” *Marketing science* in press.
- Naik, P. A., & Raman, K. (2003). Understanding the impact of synergy in multimedia communications. *Journal of Marketing Research*, 40(4), 375–388.
- Nanji, A. (2013). Online shopping trends 2013: Most popular categories, top purchase drivers. <http://www.marketingprofs.com/charts/2013/12195/online-shopping-trends-most-popular-categories-top-purchase-drivers#ixzz3IN1ie5Fe> (accessed March 17, 2015).
- Nelson, P. (1970). Information and consumer behavior. *Journal of Political Economy*, 78(2), 311–329.
- Pauwels, K. (2004). How dynamic consumer response, competitor response, company support, and company inertia shape long-term marketing effectiveness. *Marketing Science*, 23(4), 596–610.
- Pauwels, K., Currim, I., Dekimpe, M. G., Ghysels, E., Hanssens, D. M., Mizik, N., & Naik, P. (2005). Modeling marketing dynamics by time series econometrics. *Marketing Letters*, 15(4), 167–183.
- Pauwels, K., Srinivasan, S., & Franses, P. H. (2007). When do price thresholds matter in retail categories? *Marketing Science*, 26(1), 83–100.
- Pauwels, K., Stacey, C., & Lackman, A. (2013). Beyond likes and tweets: Marketing, social media content and store performance. *Marketing Science Institute, working paper 13* (pp. 125).
- Pesaran, M. H., & Smith, R. P. (1998). Structural analysis of cointegrating VARs. *Journal of Economic Surveys*, 12(5), 471–505.
- Russell, C. A. (2002). Investigating the effectiveness of product placements in television shows: The role of modality and plot connection congruence on brand memory and attitude. *Journal of Consumer Research*, 29(3), 306–318.
- Russell, C. A., & Belch, M. (2005). A managerial investigation into the product placement industry. *Journal of Advertising Research*, 45(1), 73–92.
- Sarner, A., & Herschel, G. (2008). A checklist for evaluating an inbound and outbound multichannel campaign management application. *Report G00160776*. Gartner: Stamford, CT.
- Sethuraman, R., Tellis, G. J., & Briesch, R. A. (2011). How well does advertising work? Generalizations from meta-analysis of brand advertising elasticities. *Journal of Marketing Research*, 48(3), 457–471.
- Shankar, V., & Malhotra, E. C. (2007). The growth of interactions and dialogs in interactive marketing. *Journal of Interactive Marketing*, 21(2), 2–4.
- Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, 48(6), 1–48.
- Skiera, B., & Nabout, N. A. (2013). PROSAD: A bidding decision support system for profit optimizing search engine advertising. *Marketing Science*, 32(2), 213–220.
- Srinivasan, S., Rutz, O. J., & Pauwels, K. (2015). Paths to and off purchase: Quantifying the impact of traditional marketing and online consumer activity. *Journal of the Academy of Marketing Science*, 1–14.
- Srinivasan, S., Vanhuele, M., & Pauwels, K. (2010). Mind-set metrics in market response models: An integrative approach. *Journal of Marketing Research*, 47(4), 672–684.
- Tellis, G. J., & Franses, P. H. (2006). Optimal data interval for estimating advertising response. *Marketing Science*, 25(3), 217–229.
- Trusov, M., Bucklin, R. E., & Pauwels, K. (2009). Effects of word-of-mouth versus traditional marketing: Findings from an Internet social networking site. *Journal of Marketing*, 73(5), 90–102.
- Tutaj, K., & van Reijmersdal, E. A. (2012). Effects of online advertising format and persuasion knowledge on audience reactions. *Journal of Marketing Communications*, 18(1), 5–18.
- Verhoef, P. C., Neslin, S. A., & Vroomen, B. (2007). Multichannel customer management: Understanding the research-shopper phenomenon. *International Journal of Research in Marketing*, 24(2), 129–148.
- Wiesel, T., Pauwels, K., & Arts, J. (2011). Marketing's profit impact: Quantifying online and off-line funnel progression. *Marketing Science*, 30(4), 604–611.
- Wilson, T. D., Houston, C. E., Etling, K. M., & Brekke, N. (1996). A new look at anchoring effects: Basic anchoring and its antecedents. *Journal of Experimental Psychology: General*, 125(4), 387–403.
- Zellner, A. (1962). An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American Statistical Association*, 57(298), 348–368.