



Do Display Ads Influence Search? Attribution and Dynamics in Online Advertising

**Pavel Kireyev
Koen Pauwels
Sunil Gupta**

Working Paper

13-070

February 9, 2013

Copyright © 2013 by Pavel Kireyev, Koen Pauwels, and Sunil Gupta

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Do Display Ads Influence Search?

Attribution and Dynamics in Online Advertising

Pavel Kireyev

Koen Pauwels

Sunil Gupta¹

February 9, 2013

¹ Pavel Kireyev is a Ph.D. student and Sunil Gupta is the Edward Carter Professor of Business Administration at the Harvard Business School, and Koen Pauwels is Professor at Ozyegin University, Istanbul, Turkey.

Do Display Ads Influence Search?

Attribution and Dynamics in Online Advertising

Abstract

As firms increasingly rely on online media to acquire consumers, marketing managers feel comfortable justifying higher online marketing spend by referring to online metrics such as click-through rate (CTR) and cost per acquisition (CPA). However, these standard online advertising metrics are plagued with attribution problems and do not account for dynamics. These issues can easily lead firms to overspend on some actions and thus waste money, and/or underspend in others, leaving money on the table.

We develop a multivariate time series model to investigate the interaction between paid search and display ads, and calibrate the model using data from a large commercial bank that uses online ads to acquire new checking account customers. We find that display ads significantly increase search conversion. Both search and display ads also exhibit significant dynamics that improve their effectiveness and ROI over time. Finally, in addition to increasing search conversion, display ad exposure also increases search clicks, thereby increasing search advertising costs. After accounting for these three effects, we find that each \$1 invested in display and search leads to a return of \$1.24 for display and \$1.75 for search ads, which contrasts sharply with the estimated returns based on standard metrics. We use these results to show how optimal budget allocation may shift dramatically after accounting for attribution and dynamics. Although display benefits from attribution, the strong dynamic effects of search call for an increase in search advertising budget share by up to 36% in our empirical context.

Introduction

Firms are motivated to spend more of their marketing budget online as consumers increasingly use online media to find information. Worldwide digital advertising spending in 2012 was \$103 billion, or about 20% of total money spent on advertising, and is expected to increase to \$163 billion, or 25% of total advertising spend, by the end of 2016 (eMarketer 2013). In 2012, almost half of all digital ad dollars worldwide were spent on paid search, and 38% were used for display ads (ZenithOptimedia 2012).

The introduction of online metrics such as click through rate (CTR) and cost per acquisition (CPA) by Google and other online advertisers has made it easy for marketing managers to justify their online ad spend in comparison to the budgets used for television and other media. However, these metrics suffer from the fundamental problem of attribution, since they give credit to the last click and ignore the impact of other ad formats that may have helped a consumer move down the conversion funnel. Consider, for example, a consumer searching online for a bank to open a new checking account. During this search, the consumer sees a paid search ad for a particular bank, clicks on it, and converts, recalling that she saw display ads of the same bank a few weeks earlier. How should search and display ads be credited for the conversion, and to what extent?

Most managers recognize the attribution problem, and intuitively believe that display and search ads interact to influence consumers. Recently, analytical firms and ad agencies have started addressing this problem, but most of their solutions tend to be ad-hoc. For example, some industry models give equal weight or credit to all ad exposures received by a consumer in, say, a two week period; others give more weight to recent ad exposures and exponentially lower weight to past ads (Havas Digital 2010).

As firms spend more of their ad dollars on online search and display, managers and researchers alike recognize a need for more careful attribution adjustment that takes into account the journey consumers follow before conversion. Unfortunately, the effects different advertising media have on consumers at different stages of consideration are not yet well

understood (Marketing Science Institute 2012). In this research we use time series models to infer the interaction between search and display ads. Specifically, we address the following questions:

- Do display ads influence paid search and vice versa?
- If so, how large are these effects and what dynamic patterns do they follow?
- What are the implications for online marketing metrics and optimal budget allocation?

This research draws broadly on two streams of literature – online advertising effectiveness (specifically to our context, display and search) and the spillover effects of online advertising.

In the context of display ads, researchers have studied the impact of ad exposure on click-through behavior (Chatterjee et al. 2003), long-term brand awareness (Drèze and Hussherr 2003), and repurchase decisions (Manchanda et al. 2006). Research has also explored the potential of targeted display advertising (Sherman and Deighton 2001, Shamdasani et al. 2001, Moore et al. 2005) and the consequences of its intrusiveness (Edwards et al. 2002, Goldfarb and Tucker 2011a, 2011b). Lewis and Riley (2011) use a randomized experiment to measure the causal effect of online display advertising on offline retail sales.

In the context of paid search, researchers have focused on understanding optimal advertising strategy in complex search engine environments. Ghose and Yang (2009) and Rutz et al (2012) adopt a keyword-specific approach to understand the performance of individual keywords and guide optimal keyword investment decisions. Further work examined spillover within search. Yang and Ghose (2010) identify complementarities across organic and paid search listings, and Rutz and Bucklin (2011) find spillover effects from generic search to branded search. Wiesel et al (2011) model consumer progression through the purchase funnel, and explain how online advertising may drive sales in the offline channel. In contrast to the above study, which mainly examines advertising effectiveness within a particular online channel, we study ad effectiveness taking into account the interaction and feedback between both search and display channels. Furthermore, we focus on the role of search and display advertising in

customer acquisition in the commercial banking industry, where consumer decision process tends to be longer and more involved, and the attribution problem is more severe.

Several studies consider spillovers and synergies in online and offline consumer behavior. Naik and Peters (2009) propose a hierarchical model to capture synergies within the offline channel and across online and offline channels. Their model builds on earlier work (Naik and Raman 2003), and argues that investing in offline and online advertising simultaneously generates greater revenues than investing in each channel individually.

A small number of studies examine the interaction between paid search and display. Papadimitrou et al (2012) conduct a field experiment to explore the impact of display exposure on search queries. They find that exposure to a display ad increases the number of relevant search queries submitted by 5%-25%. Lewis and Nguyen (2011) conduct a field experiment to explore the impact of display advertising on advertiser- and competitor-branded search queries. Using a very short time window (10 minutes) they find a 27%-45% lift in searches attributable to the display advertising exposure. A number of studies conducted by industry researchers also explore the impact of display advertising on paid search. A survey conducted by iProspect (2009) finds that about 50% of all internet users react to a display ad by conducting a search related to the brand or product described in the ad. The study finds that 14% of users make a purchase after conducting the search. A study by comScore (Fulgoni and Mom, 2008) finds a 38% lift in branded search activity for consumers exposed to a display ad. The study tracked individual consumers exposed to a display ad, and compared their behavior with a group of similar consumers not exposed to display advertising. An iCrossing study (Malm and Hamman 2009) finds a 14% change in search visits after a company activated its display advertising campaign. In our application for the bank, the bank's ad agency also conducted an experiment to find that display ads improve search ad conversion by 15-20%.

As these studies suggest, display ads appear to influence the effectiveness of search ads. However they lack two important elements that we consider in our study. First, almost all of these studies ignore the dynamic effects of advertising, whereby display ads may impact consumers' search behavior over time. Studies that attempt to incorporate dynamics do so in

an ad-hoc fashion. For example, the ad agency for our bank decided to use a two-week period (an ad-hoc assumption) to examine its effect. In our application, we show that these dynamic effects are very strong and may last several weeks. Ignoring them can lead to significant underestimation of the effectiveness of online ads. Second, most of the previous studies used click-through rates or similar metrics to measure the impact of display ads on search. In contrast, we examine how display ads influence search clicks, conversion and ultimately the profitability of the firm. This allows us to determine appropriate budget allocation between search and display.

The remainder of this article is organized as follows. First, we present our conceptual framework by explaining the relation between attribution and consumer funnel progression. Then, we present the data, modeling methodology, and empirical analysis. To conclude, we provide a set of attribution and dynamics adjusted marketing metrics, and discuss managerial implications.


Conceptual Framework

The consumer journey can be conceptualized as a conversion funnel. A consumer may be exposed to a brand through display ads, she may click on these ads to get more information, and may eventually convert. This is the direct impact of display ads on conversion that most studies find to be very small. Alternatively, a consumer could be actively searching for a product online, where she encounters a search ad, clicks on it, and converts. This is the direct effect of search ads, which is usually bigger than the direct effect of display ads. It is common to measure these direct effects of display and search using online metrics such as CTR, CPC, and CPA.

Besides direct conversion, passive forms of advertising exposure may also influence consumers' consideration sets, and subsequent active engagement with the firm moves consumers down the funnel towards conversion. In this scenario, display ads may influence consumers at the top or middle of the purchase funnel while search ads may have more impact at the bottom of the funnel.

Figure 1 shows how the firm's online advertising strategy may influence consumers' purchase behavior and the firm's budget allocation. In this framework, the firm allocates a budget between search and display ads that determines the number of ad impressions to consumers. These in turn affect display or search clicks, and eventually, conversion. Two important aspects of our framework should be noted. First, we expect strong interaction between search and display ad impressions and clicks. The empirical results will show if display ads indeed influence search ad effectiveness, and if so, by how much. Second, the system explicitly recognizes endogeneity, whereby the firm's advertising budget influences consumers' exposure and purchase behavior, which in turn affects how much the firm spends on advertising.

Figure 1: Impact of Online Advertising on Consumer Behavior



Data Description

We use data from a large commercial bank that operates mainly in the southern U.S. After the financial crisis, advertising to acquire new consumers became increasingly important given reduced margins and declining consumer confidence. The bank invests heavily in both paid search and display ads to acquire customers for its checking account.

For the calendar year 2010, the bank and its advertising agency provided us weekly data on the bank's online marketing expenditure, search and display impressions and clicks, and the number of online applications completed by consumers for a new checking account. The bank invested about \$1 million in online advertising, almost equally split between search and display.² There are two limitations of our data set. First, the bank does not track if online advertising influences consumers to open a checking account in its retail branch. This means that we cannot investigate the impact of online advertising on offline behavior and vice versa. Second, our dataset consists of only aggregate levels of consumer behavior. Although the lack of individual-level data is a limitation for our study, managers routinely use aggregate data to assess the performance of their online campaigns.

Paid search data capture weekly spend, clicks, impressions, and the number of applications completed through the paid search ad's landing page. The display data also contain information on weekly spend, clicks, and impressions. Using internet cookies, applications completed were attributed to display advertising if a consumer had seen a display ad at least one month before converting through the display ad network's landing page using organic search or a direct link.³ However, the display advertising data excluded display-driven paid search conversions, as the paid search campaigns are overseen by platforms maintained by search engines unrelated to the display ad networks.

² To maintain the confidentiality of the client bank, we have disguised some of the data while maintaining the relationship between the variables of interest.

³ This is another form of ad-hoc attribution between display ads and other online media. However, we do not investigate this in our study due to the lack of individual-level data available to us.

The bank invested in five search engines and eleven ad networks. We aggregate over search engines and ad networks to the week level to avoid over-parameterization as our primary interest lies in the interplay of search and display advertising, as opposed to the performance of individual search engines or display ad networks. Furthermore, the bank's limited investment over a number of smaller ad networks and search engines makes it difficult to estimate the impact at the level of a search engine or ad network.

Table 1 presents a correlation matrix of the variables in our data. The notations used for the variables are indicated below:

- SA_t : Checking account applications completed through paid search in week t .
- DA_t : Checking account applications completed after exposure to a display ad.
- SI_t : Paid search ad impressions.
- SC_t : Paid search ad clicks.
- SE_t : Weekly expenditure on paid search advertising.
- DI_t : Display ad impressions.
- DC_t : Display ad clicks.
- DE_t : Weekly expenditure on display advertising.

Table 1: Correlation Matrix

	SA	DA	SI	SC	SE	DI	DC	DE
SA								
DA	0.76							
SI	0.80	0.60						
SC	0.66	0.56	0.65					
SE	0.80	0.75	0.88	0.68				
DI	0.71	0.89	0.59	0.44	0.75			
DC	0.67	0.88	0.58	0.41	0.74	0.98		
DE	0.56	0.84	0.54	0.34	0.69	0.94	0.94	

Table 1 shows that many variables are highly correlated, especially those related to the same marketing instrument. Display impressions exhibit a 0.98 correlation with clicks and a 0.94 correlation with spend. Therefore, we excluded display clicks and spend from the analysis.

In the case of paid search, spend exhibits a high correlation (0.88) with impressions, so we exclude search advertising spend from the analysis to minimize possible collinearity.

Table 2: Summary statistics (per week)

Variable	SA	DA	SI	SC	DI
Mean	143	139	230,927	9,778	6,023,264
Median	125.5	150.5	194,013	9,478	6,176,496
Maximum	278	304	745,911	21,073	14,885,122
Minimum	43	0	45,565	2,685	242
Std Dev	66	97	134,968	4,060	5,008,892

Figure 1: Weekly Trend




Table 2 provides summary statistics of our data and Figure 2 illustrates the weekly trend of the resulting series. The decreasing trend present in all variables arises as a consequence of the bank exhausting its advertising budget, and hence decreasing its investments over time to avoid overspending. We incorporate this trend as a non-deterministic component of the model, allowing for the other endogenous variables to explain it.

Methodology and Analysis


We use persistence modeling techniques to capture the complex dynamic interdependencies in online advertising (Dekimpe and Hanssens 1999). Persistence modeling extends multivariate time series methods into the domain of marketing, thereby enabling researchers to model the effects of spillover and feedback dynamics through a system of equations involving marketing actions and consumer response. Persistence modeling is particularly relevant in the context of online advertising as the associated multivariate time series techniques require no stringent a priori restrictions on model structure and allow all variables of interest to affect each other.

Persistence modeling involves several steps. A series of tests are used to determine the correct model specification. Granger causality tests are used to identify which variables enter the system endogenously. Unit root tests are done to determine which of the endogenous variables exhibit non-stationary behavior and should enter the model in differences. Next, cointegration tests are used to identify stationary linear combinations of non-stationary endogenous variables that must be considered in the specification to correct for temporary deviations away from the implied long-run equilibria.

Granger causality tests, conducted pair-wise for variable lag-lengths ranging from 1 to 20, suggest that all variables should enter the system endogenously. Figure 3 presents a schematic of the Granger causality results. For example, an arrow from *DI* to *SC* indicates that *DI* is found to Granger-cause *SC* for at least one of the lag-lengths considered. Interestingly, no arrow exists from display impressions to search applications, implying that if display does affect search, the effect travels through search impressions and search clicks. The complex nature of

interdependencies depicted in Figure 3 points to the appropriateness of using a flexible approach, such as persistence modeling, to capture cross-ad spillovers and online advertising dynamics.

Figure 3: Granger-causality graph



We conduct Augmented Dickey-Fuller (ADF) and Kwiatkowski-Phillips-Schmidt-Shin (KPSS) unit root tests to determine if the endogenous variables are evolving or stationary. Table 3 summarizes the resulting statistics of the unit root tests. The KPSS test identifies all series as evolving, whereas the ADF test identifies all but *SI* as evolving. We choose to include *SI* as an evolving variable following the outcome of the KPSS test to allow for richer cointegration possibilities.

Table 3: Summary of unit root test results

Test\Variable	SA	DA	SI	SC	DI
ADF	-2.006	-0.428	-4.444	-2.657	-1.984
KPSS	0.872	0.874	0.813	0.587	0.850

Note: Bold numbers indicate significant evidence of non-stationarity

The Johansen cointegration trace test identifies three cointegrating relations. These relations can be interpreted as long-run equilibrium conditions which may arise as a result of firm budgeting rules or consumer decision processes. Based on the outcomes of the Granger causality, unit root and cointegration tests, we specify a vector error correction model (VEC) with all variables as endogenous. The interpretation of VEC models is particularly interesting from a substantive perspective. Both managers and researchers may expect a long-term equilibrium linking search and display applications to a combination of firm control variables (e.g. search and display ad impressions) and consumer actions (e.g. clicks on these ads). Economist Walter Enders states that economic theory abounds with equilibrium theories which, if they involve non-stationary variables, “require the existence of a combination of the variables that is stationary” (2010, p.356). Within marketing, researchers often voice opinions about the necessary intricate relation between firm activity, consumer activity, and purchase action. Error correction models have been used to study the long-run impact of a product harm crisis (Van Heerde et al 2007), market share cannibalization by new innovations (Van Heerde et al 2010), and long run sales sensitivity to price changes (Fok et al 2006).

The general form of the VEC model with K lags is given by equation 1,

$$\Delta Y_t = \Gamma_0 D_{1t} + \sum_{k=1}^K \Gamma_k \Delta Y_{t-k} + \alpha e_{t-1} + u_t,$$

where $e_{t-1} = \beta \begin{bmatrix} D_{2t} \\ Y_{t-1} \end{bmatrix}$, (1)

and $u_t \sim MVN(0, \Sigma)$,

In equation (1), Y_t is the vector of endogenous variables at time t , D_{1t} and D_{2t} are vectors of deterministic components (e.g. intercept, trend), e_t is a matrix of cointegrating relations, $\Gamma_0, \Gamma_1, \dots, \Gamma_k, \alpha$ and β are parameter matrices to be estimated, and Σ is the covariance matrix of the multivariate-normally distributed error terms u_t . The coefficients in $\Gamma_1, \dots, \Gamma_k$ capture the effects of past changes in the endogenous variables on their current deviations. The coefficients in α reflect the speed of adjustment of the endogenous variables towards the equilibrium cointegrating relations defined in e_{t-1} . We refine the model further by

allowing for an intercept in both D_{1t} and D_{2t} . The intercept in the model specification allows for the possibility of a deterministic time trend to exist concurrently with the stochastic one implied by the error correction model. The intercept term in the cointegrating vector is included to account for the initial values of the endogenous variables. The Bayesian Information Criterion identifies a lag-length of 1 as optimal. The resulting model specification is indicated in equation (2):

$$\begin{aligned}
 \begin{bmatrix} \Delta SA_t \\ \Delta DA_t \\ \Delta SI_t \\ \Delta SC_t \\ \Delta DI_t \end{bmatrix} &= \begin{bmatrix} \gamma_{10} \\ \gamma_{20} \\ \gamma_{30} \\ \gamma_{40} \\ \gamma_{50} \end{bmatrix} + \begin{bmatrix} \gamma_{11} & \gamma_{12} & \gamma_{13} & \gamma_{14} & \gamma_{15} \\ \gamma_{21} & \gamma_{22} & \gamma_{23} & \gamma_{24} & \gamma_{25} \\ \gamma_{31} & \gamma_{32} & \gamma_{33} & \gamma_{34} & \gamma_{35} \\ \gamma_{41} & \gamma_{42} & \gamma_{43} & \gamma_{44} & \gamma_{45} \\ \gamma_{51} & \gamma_{52} & \gamma_{53} & \gamma_{54} & \gamma_{55} \end{bmatrix} \begin{bmatrix} \Delta SA_{t-1} \\ \Delta DA_{t-1} \\ \Delta SI_{t-1} \\ \Delta SC_{t-1} \\ \Delta DI_{t-1} \end{bmatrix} \\
 &+ \begin{bmatrix} \alpha_{11} & \alpha_{12} & \alpha_{13} \\ \alpha_{21} & \alpha_{22} & \alpha_{23} \\ \alpha_{31} & \alpha_{32} & \alpha_{33} \\ \alpha_{41} & \alpha_{42} & \alpha_{43} \\ \alpha_{51} & \alpha_{52} & \alpha_{53} \end{bmatrix} \begin{bmatrix} \beta_{10} & 1 & 0 & 0 & \beta_{11} & \beta_{12} \\ \beta_{20} & 0 & 1 & 0 & \beta_{21} & \beta_{22} \\ \beta_{30} & 0 & 0 & 1 & \beta_{31} & \beta_{32} \end{bmatrix} \begin{bmatrix} 1 \\ SA_{t-1} \\ DA_{t-1} \\ SI_{t-1} \\ SC_{t-1} \\ DI_{t-1} \end{bmatrix} + \begin{bmatrix} u_{SA,t} \\ u_{DA,t} \\ u_{SI,t} \\ u_{SC,t} \\ u_{DI,t} \end{bmatrix} \quad (2)
 \end{aligned}$$

The parameters are recovered in two steps. First, Johansen's procedure is used to estimate the cointegrating vectors. Then, the first differences of the endogenous variables are regressed on an intercept, their lags and the cointegrating vectors to recover the remainder of the coefficients. Not all the coefficients in this model are identified. In particular standard errors cannot be recovered for β_{10} , β_{20} and β_{30} . Furthermore, an arbitrary normalization is required to identify the remaining coefficients of the β matrix.

Table 4 summarizes the full set of parameter estimates and asymptotic standard errors. The model exhibits good fit for a model in differences, with individual equation R^2 statistics ranging from 0.27 to 0.45. Portmaneau tests fail to find significant evidence of residual autocorrelation and normality tests fail to reject normality of the residuals. Furthermore, generalized fluctuation tests for structural change fail to find significant evidence of parameter instability.

Table 4: VEC parameter estimates (asymptotic t-statistics in parentheses)

Cointegrating Eq:					
	$e_{SA,t-1}$	$e_{DA,t-1}$	$e_{SI,t-1}$		
SA_{t-1}	1.000000	0.000000	0.000000		
DA_{t-1}	0.000000	1.000000	0.000000		
SI_{t-1}	0.000000	0.000000	1.000000		
SC_{t-1}	-0.026943 [-6.19270]	0.001345 [0.58142]	-17.58090 [-5.00585]		
DI_{t-1}	1.77E-06 [0.51097]	-2.14E-05 [-11.6546]	-0.009488 [-3.39954]		
<i>constant</i>	109.1041	-26.13430	5200.822		
Error Correction:					
	ΔSA_t	ΔDA_t	ΔSI_t	ΔSC_t	ΔDI_t
$e_{SA,t-1}$	-0.091170 [-1.52498]	-0.175613 [-2.74302]	79.30134 [0.66615]	5.523423 [0.89771]	7495.433 [1.21460]
$e_{DA,t-1}$	-0.107651 [-1.01232]	-0.139443 [-1.22448]	-429.4406 [-2.02803]	-29.77596 [-2.72068]	32198.84 [2.93331]
$e_{SI,t-1}$	-0.000157 [-2.85132]	0.000195 [3.31444]	-0.402342 [-3.67005]	0.002561 [0.45207]	3.103370 [0.54608]
ΔSA_{t-1}	-0.443194 [-2.78235]	-0.138795 [-0.81368]	-721.8211 [-2.27574]	-3.728723 [-0.22745]	-7399.273 [-0.45002]
ΔDA_{t-1}	0.234133 [1.90658]	-0.433164 [-3.29385]	87.36913 [0.35729]	-2.568378 [-0.20322]	-11868.29 [-0.93627]
ΔSI_{t-1}	0.000189 [2.63937]	-4.02E-05 [-0.52406]	0.327024 [2.29089]	-0.014157 [-1.91883]	1.069115 [0.14448]
ΔSC_{t-1}	-0.002645 [-1.62591]	-0.000102 [-0.05828]	-4.195442 [-1.29520]	-0.127549 [-0.76186]	249.0512 [1.48318]
ΔDI_{t-1}	-3.72E-06 [-2.05150]	5.22E-06 [2.68708]	-0.009798 [-2.71173]	-0.000447 [-2.39148]	0.144877 [0.77352]
<i>constant</i>	-1.594333 [-0.45454]	-6.117405 [-1.62862]	-7426.962 [-1.06336]	-356.0680 [-0.98637]	-249869.3 [-0.69013]
R^2	0.396648	0.381397	0.453215	0.370934	0.266994

It is difficult to directly interpret the parameters of persistence models, so we proceed to derive implications by impulse response analysis.

The Effects of Search and Display Ads

As recommended for multivariate time series models (Sims 1980), we use impulse response functions to analyze the impact of search and display advertising, and assess significance by applying a one standard error band to the impulse response coefficients⁴ (Sims and Zha 1999, Dekimpe and Hanssens 1999). Pesaran and Shin (1998) provide a derivation of the generalized impulse response function, which captures the impact of an unexpected shock to the endogenous variables in a VEC model by constructing two forecasts and taking their difference. One forecast takes the shock into consideration, while the other does not. The difference of the two forecasts provides the incremental impact of the shock. Impulse response functions trace the impact of a shock to one endogenous variable through other endogenous variables, thereby providing a cumulative view of all dynamic interactions that take place.

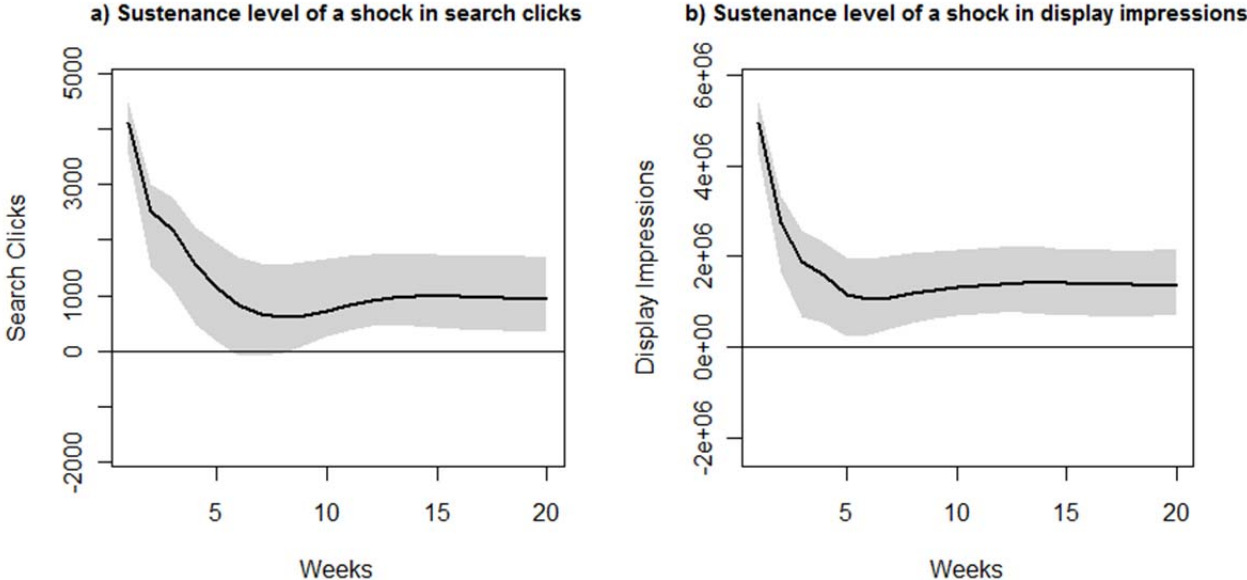
For search clicks and display impressions, an unexpected shock represents an investment injection by the firm. In the case of search and display applications, and search impressions, it would imply a scenario in which we would observe an unexpected increase in applications or search viewership, holding display exposure and search clicks unchanged. It is common in practice to make budgeting decisions based on search clicks and CPC, and display impressions and cost per thousand (CPM) impressions. Therefore, we use search clicks and display impressions as the marketing variables of interest and interpret the forecasts that result from their shocks as the effects of increases in marketing investment. We apply one standard deviation shocks to the marketing variables and study their sustenance, implications for performance, and interaction between search and display.

Figure 4 presents the sustenance levels of search clicks and display impressions. Sustenance measures the response of a variable to a one standard deviation shock to itself. The

⁴ We calculate confidence bands for the impulse response functions by simulating 1000 random draws from a multivariate normal distribution with mean zero and covariance matrix equal to the residual covariance matrix of the model, using these draws to perturb the data, and estimating the impulse response functions 1000 times on the resulting simulated datasets. Quantiles of the distributions of coefficients provide an indication of the accuracy of the impulse response functions. We take the 16th and 84th percentiles of the empirical distribution to approximate a one standard error band.

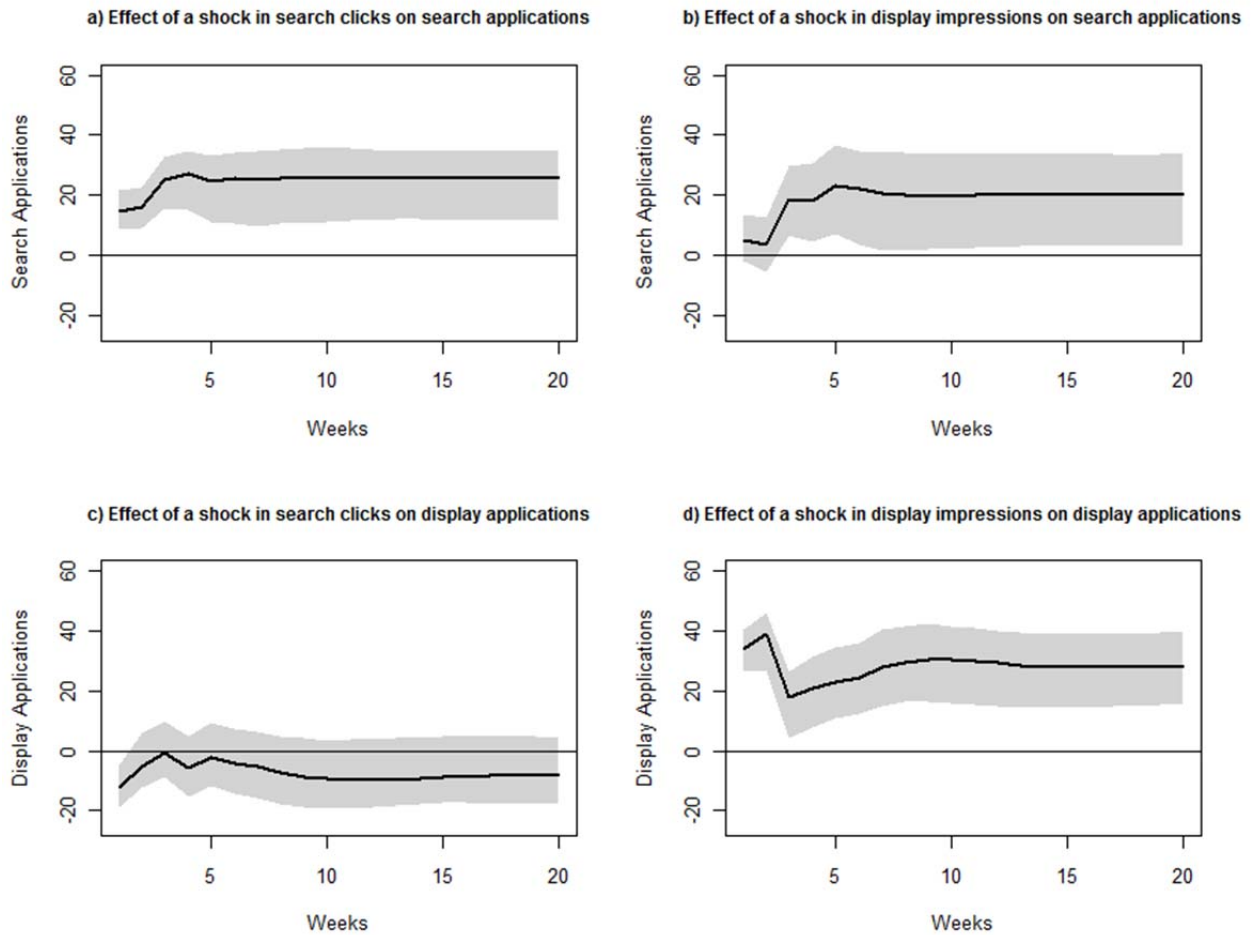
plots suggest that persistent investment and complex consumer transitions between different channels of the conversion funnel lead to sustained levels of long-run exposure to marketing. Panel 4a shows that a shock of 4,000 search clicks wears-in after 8-10 weeks and stabilizes at about 900 clicks per week in the long run. Display impressions follow a similar pattern according to panel 4b. A shock of 5 million impressions wears-in over a period of 7-8 weeks and stabilizes at a sustained level of 1.4 million impressions per week.

Figure 4: Sustenance levels of marketing variables



The plots in Figure 5 show the performance impact of marketing. The top row captures the impact of initial shocks and persistence in marketing exposure on search applications. Panel 5a shows the impact of search clicks on search applications. A shock of 4,000 clicks generates 15 search applications initially. After a wear-in period of 4 weeks, 900 clicks (Figure 4a) generate 26 applications per week (Figure 5a). A smaller number of search clicks is required to maintain a higher level of search applications in the long run, suggesting that the effectiveness of an injection to search advertising increases as it persists over time.

Figure 5: Performance impact of marketing variables



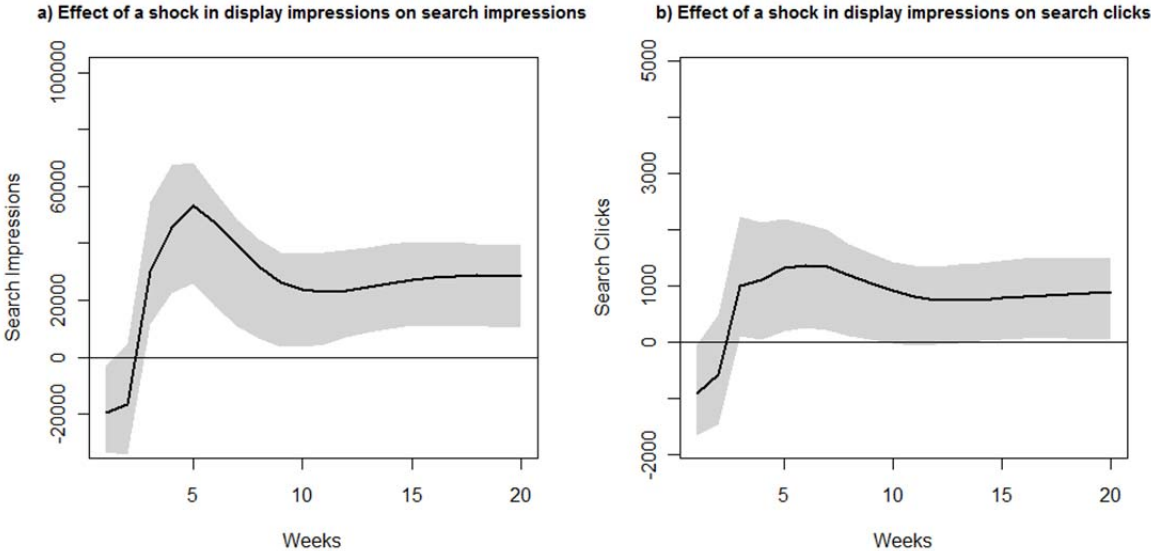
Panel 5b shows the impact of display impressions on search applications. As expected, an initial increase in display impressions does not generate any search applications. However, after a period of two weeks, display impressions positively impact search applications. A sustained level of 1.4 million impressions (Figure 4b) generates about 20 search applications per week (Figure 5b). Consistent with our conceptual framework, display exposure appears to drive consumers to paid search over time.

The bottom row of Figure 5 captures the impact of online ads on display applications. A shock to search clicks does not affect display applications (Figure 5c), except for the initial period, which may point to consumers who would have applied through display substituting into the search channel. Panel 5d shows that the effect of display impressions on display

applications is powerful and immediate. A shock of 5 million impressions (Figure 4b) generates 34 applications immediately. After one week, display applications dip and then stabilize at 28 applications per 1.4 million impressions (Figure 4b) in the long run.

To further understand the impact of advertising, we consider the interaction between search and display ads to see how increased levels of display ads may drive search impressions and clicks. Figure 6 plots the impact of a shock in display impressions on search impressions and search clicks. In the short-run, we observe a decrease in both search impressions and clicks, which may be driven by consumer substitution across channels. In the long-run, a sustained increase in display impressions drives a significant increase in search impressions and clicks, suggesting that display exposure not only increases conversion through search, but also drives search visitation and search clicks. This finding, together with the lack of direct Granger-causality between display impressions and search applications, suggests that display advertising drives search applications through search impressions and clicks. Hence, in calculating the overall impact of display advertising, we must take into account the potential associated increases in costs from search advertising.

Figure 6: Impact of display advertising on search funnel progression



We explored the sensitivity of the impulse response analysis to our modeling assumptions. In particular, we investigated how strongly the results depend on the non-stationarity and cointegration of the data series. We estimated a basic vector-autoregressive simultaneous equations model, with the endogenous variables entered in levels. Although such a specification is known to yield biased estimates, it is informative to see how significantly these biases affect our results. The VAR specification yielded qualitatively similar findings for the short-term and wear-in period. As implied by this specification, the long-run effects were not persistent. Hence, non-stationarity and cointegration only drive the long-run behavior of the impulse response functions. This shows that only the part of our model that is attributable to non-stationarity depends on it, and hence invokes confidence in the short-run and wear-in impulse response estimates.

We also estimated variance decompositions of forecast errors to confirm that display advertising indeed drives search behavior as implied by the Granger-causality tests and impulse response analysis. Variance decompositions showed that an exogenous shock to display impressions explained 40% of the forecast error variance in search impressions, 17% in search clicks, and 16% in search applications, suggesting that display impressions indeed move consumers through search media.

Managerial Implications

Standard online metrics such as CPA and ROI are static measures that ignore attribution or dynamic effects. As implied by the impulse response analysis, shocks to display advertising increase both exposures to search marketing and search applications. Moreover, the performance effects are non-stationary and stabilize only 2-4 weeks after the initial marketing shock, implying that marketing metrics should take into account not only attribution, but also the dynamic effects of marketing.

CPA and ROI of Search Ads

We begin with the implications for search metrics, where only dynamics need to be taken into account, since search clicks do not impact display applications (Figure 5c). Although a

complicated pricing and bidding system drives cost per click (CPC), as a simplification we assume that CPC remains constant over the impulse response forecast.

Table 5 contrasts CPA and ROI calculated in a standard fashion with their dynamic counterparts as implied by the impulse response analysis. CPA is calculated as the total search expenditure divided by the sum of all search applications. The standard approach is a static measure of cost per acquisition (or application in our context) that is commonly used by most marketing managers and search engines like Google. In contrast, dynamic CPA incorporates the long-run effects implied by the impulse response functions.

Table 5: CPA and ROI for search ads

	Standard	Dynamic (immediate)	Dynamic (wear-in)	Dynamic (long-run)	% Change: Standard vs. Dynamic
CPA	\$73	\$296	\$52	\$38	-47.5%
ROI	\$1.27	\$0.46	\$1.52	\$1.75	+37.7%

Numbers in this table have been rounded.

Using the bank’s annual expenditure on search ads and the total number of search applications, we find the standard CPA to be \$73, a number that the bank and its ad agency used to assess the performance of its search ads. For estimating the dynamic effect of search, recall that a shock of 4,000 search clicks generates about 15 applications in the short run. Using the average CPC rate of \$1.07 from our data, the immediate dynamic CPA is then \$296. However, as discussed earlier, search clicks increase by 900 per week and generate about 26 search applications per week in the long run, implying a long run CPA of \$38, 48% lower than the standard CPA.

Table 5 also provides a measure of ROI that shows the return for every \$1 invested in search ads. According to the bank, about 80% of the customers who complete online applications are approved for a checking account, and two-third of the approved customers actually fund the account (i.e., put some money in their account within a month). In other words, $80\% \times 67\% = 53.6\%$ of the applications become active customers. In addition, both search

and display ads were generally accompanied by a promotional offer that included a free iPod Nano, iPod Touch or \$100-\$150 cash. The bank estimates that on average the effective cost of these promotions is about \$100 for each new active customer acquired through the online channel. The bank further estimates the average customer lifetime value (CLV) to be \$300 for every active account.

Using this information we calculated the ROI for the standard and the dynamic approaches. For example, the standard CPA is \$73, but the effective cost of getting an active account is $\$[(73/0.536)+100] = \236 , and the benefit of this account in the long run is its CLV of \$300. ROI is then simply the benefit (\$300) divided by the effective cost of an active account (\$236), or 1.27 for the standard approach.

The results show that accounting for long run dynamic effects reduces search CPA by 48% and increases its ROI by 38% compared to the standard metrics that ignore these dynamic effects. In other words, the firm may be underinvesting in search by relying on standard metrics.

CPA and ROI of Display Ads

Results for impulse response functions show that search ads do not affect display applications, but display impressions influence both search and display applications. Therefore, in addition to dynamics, display metrics should incorporate attribution. Moreover, shocks to display advertising not only increases search applications, but also increases search clicks, which may lead to greater search cost. Display attribution must take into account not only the benefit, but also this additional cost of spillover into the search channel.

To make this point salient, Table 6 presents a comparison of three methods for calculating display CPA and ROI – without attribution to search, with attribution to search applications, but without accounting for additional search cost, and with additional search applications and search cost both considered. For these calculations we use the average cost-per-thousand (CPM) impressions of \$2.05 from our data. The calculations for CPA and ROI follow the same logic as before, except that now we also include the impact of display impressions on search applications and search cost.

Table 6: CPA and ROI for display ads

	Standard		Dynamic (immediate)	Dynamic (wear-in)	Dynamic (long-run)	% Change
CPA	\$88	No attribution	\$298	\$120	\$99	+12.7%
		Attribution to search applications only	\$258	\$71	\$57	-35.2%
		Attribution to search applications and clicks	\$233	\$94	\$76	-14.1%
ROI	\$1.14	No attribution	\$0.46	\$0.92	\$1.05	-7.3%
		Attribution to search applications only	\$0.52	\$1.29	\$1.45	+28.0%
		Attribution to search applications and clicks	\$0.56	\$1.09	\$1.24	+9.6%

The long-run CPA for display is 35% lower than in the standard approach when we account for its impact on search applications but ignore the additional cost it may drive. However, even after accounting for the additional cost, long-run CPA for display is 14% lower than in the standard approach. ROI of display impressions exhibits a similar pattern – it is 28% higher than the standard ROI when only attribution to search applications is considered, and is about 10% higher when both additional search applications and extra search costs due to display ads are included.

Budget Allocation

It is clear from the previous analysis that search ads are more effective than display ads even when we account for attribution effects of display ads. This is due to the fact that search ads show a significant dynamic effect, which is perhaps reasonable in the context of a bank checking account, where consumers are likely to take several weeks before making a decision. These results have direct implications for budget allocation. How should the firm allocate its online advertising budget between search and display and how does this allocation compare to the firm's current allocation?

In a non-stationary scenario, the firm should allocate budget according to the long-run effectiveness of marketing instruments (Dekimpe and Hanssens 1999). Optimal budgeting would then allocate shares according to the ratio of display and search advertising elasticities. Table 7 presents the advertising elasticities⁵ of marketing actions after taking into account attribution and dynamics, and Table 8 shows the actual and proposed budget allocation.

Table 7: Advertising elasticities

	Ad Elasticity (immediate)	Ad Elasticity (wear-in)	Ad Elasticity (long-run)
Search	0.12	0.71	0.96
Display	0.19	0.46	0.57

Consistent with our previous results, we find that search elasticities are significantly higher than the display elasticities, suggesting that the firm would be better off spending more on search than its current 50% budget allocation. Display advertising yields a lower elasticity even after accounting for attribution. Wiesel et al (2011) similarly find a high search advertising

⁵ The long-run elasticity reflect the percent change in the total number of display and search applications from a 1% change in investment for a particular marketing instrument, taking into account any additional costs it may drive. We use sample means instead of the last observation in the series as our data exhibit a decreasing tendency, whereas impulse response functions are calculated as averages over the data range. Therefore, using the last observation in the series may yield inconsistent elasticity estimates.

elasticity of 4.35 in the context of a multichannel furniture retailer. Dinner et al (2011) find a long-run search elasticity of 0.49 and a display elasticity of 0.15. Manchanda et al (2006) find a display elasticity of around 0.02 with respect to repeat purchase behavior. While our elasticity estimates are within the broad range of the estimates found in the previous studies, it is important to note that our context of bank applications and the \$100 incentive offered by the bank makes direct comparisons across studies somewhat difficult.

Given the current advertising budget of the firm, the optimal allocation between search and display ads is given by the ratio of their elasticities. Table 8 shows the actual and proposed budget allocation.

Table 8: Actual and proposed budget allocations

	Actual Budget	Proposed Budget	% Change
Search	\$542,000	\$739,000	+36%
Display	\$639,000	\$442,000	-31%

The firm is currently allocating 54% of its online ad budget on display advertising even though the standard metrics used by the firm show the search CPA (\$73) to be about 20% lower than the display CPA (\$88). The bank and its ad agency made this allocation recognizing that the standard metrics do not account for attribution. Given the nature of the product category they expected display to have significant impact on search applications. To test this hypothesis the ad agency conducted a field experiment and found that search effectiveness improved by about 20% when it was followed by display ads. This factored into their budget allocation.

However, our model suggests that search should have 63% of the total budget, or almost 36% higher than the budget currently allocated by the firm, and display budget should be reduced. It may seem counterintuitive to reduce the budget for display ads after accounting for its attribution (something that the firm is also trying to do through its experiment), but a simple attribution analysis ignores two important aspects. First, it ignores the additional cost of search clicks that are accompanied by the search applications generated by display. Second,

and perhaps more important in our application, the firm is ignoring dynamic effects that are particularly strong for search ads.

Conclusions

Our goal in this study was to find out if online display ads influence search (attribution problem), if online advertising, more generally, exhibits dynamic effects, and if so, what implications this has for the firm's budget allocation. We used persistence modeling on data from a bank that used online advertising to acquire new customers for its checking account. We found that display ads have a significant impact on search applications, as well as clicks. The majority of this spillover was not instant, but took effect only after two weeks. On the other hand, search advertising did not lead to an increase in display applications.

Our findings suggest that simple static metrics, commonly used in the industry, may not accurately measure the effectiveness of online advertising. We propose dynamic versions of the classic metrics and find that search CPA is 48% lower than the static CPA, while search ROI is 38% higher than the static ROI. Similar pattern emerges for display advertising, where we also account for attribution. This made display CPA 14% lower and ROI 10% higher than their standard counterparts. Finally, we show that these revised measures of ad effectiveness lead to a very different budget allocation than the one used currently by the firm. Specifically, we find that even though our proposed allocation gives credit to display due to its effect on search applications, search ad budget should be increased by 36% from its current level due to its strong dynamic effects, and display ad budget should be decreased by 31%.

Our study has several limitations that can provide avenues for future research. We do not consider spillovers effects of search and display into other channels. Future research may examine the effects of online ads on conversions and funnel progression in mobile and offline channels. We use aggregate data that does not allow us to untangle the mechanism that may be driving consumer decisions. Using disaggregate data, future research could provide richer insights into the consumer journey and progression, and the differential impact of various

marketing instruments at various stages of the conversion funnel. Future studies may also wish to generalize our findings by examining multiple products and contexts.

Overall, our research suggests that managers should carefully consider the interaction and dynamic effects of search and display advertising. Our results show that classic metrics used in practice are highly biased since they do not account for these effects. As a result firms may be making suboptimal budget allocation decisions.

References

- Chatterjee, Patrali, Donna L. Hoffman and Thomas P. Novak (2003), "Modeling the Clickstream: Implications for Web-Based Advertising Efforts," *Marketing Science*, 22 (4), 520-541.
- 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.
- Dinner, Isaac M., Harald J. van Heerde and Scott A. Neslin (2011), "Driving Online and Offline Sales: The Cross-Channel Effects of Digital versus Traditional Advertising," Working Paper.
- Drèze, Xavier, and François-Xavier Hussherr (2003), "Internet Advertising: Is Anybody Watching?" *Journal of Interactive Marketing*, 17 (4), 8-23.
- Edwards, S. M., Li, H., and Less, J.-H. (2002), "Forced Exposure and Psychological Reactance: Antecedents and Consequences of the Perceived Intrusiveness of Pop-Up Ads," *Journal of Advertising*, 31 (3), 83-95.
- eMarketer (2013), "Digital to Account for One in Five Ad Dollars," Retrieved from <http://www.emarketer.com/Article/Digital-Account-One-Five-Ad-Dollars/1009592#kPpaVi4w0A2DpGPA.99>
- Enders, W. (2008). *Applied econometric time series*. John Wiley & Sons.
- Fok, D., Horváth, C., Paap, R., & Franses, P. H. (2006), "A Hierarchical Bayes Error Correction Model to Explain Dynamic Effects of Price Changes," *Journal of Marketing Research*, 43(3), 443-461.
- Fulgoni, G. M., & Morn, M. (2008), "How online advertising works: Whither the click," *comScore.com Whitepaper*.
- Ghose, A., S. Yang (2009), "An empirical analysis of search engine advertising: Sponsored search in electronic markets," *Management Science*, 55(10) 1605–1620.

- Goldfarb, Avi and Catherine Tucker (2011a), "Online Display Advertising: Targeting and Obtrusiveness," *Marketing Science*, 30 (3), 389-404.
- Goldfarb, Avi and Catherine Tucker (2011b), "Search Engine Advertising: Channel Substitution When Pricing Ads to Context," *Management Science*, 57 (3), 458–70.
- Havas Digital (2010), "Artemis Attribution Weighting," [White paper]. Retrieved from http://www.havasdigital.com/wp-content/uploads/2011/02/HD_Insight_AttributionDM.pdf
- iProspect (2009), "Search Engine Marketing and Online Display Advertising Integration Study," [White paper]. Retrieved from http://www.iprospect.com/wp-content/uploads/2011/11/iProspectStudy_May2009_Search-Engine-Marketing-and-Online-Display-Advertising-Integration-Study.pdf
- Lewis, R., & Nguyen, D. (2012), "Wasn't That Ad for an iPad? Display Advertising's Impact on Advertiser-and Competitor-Branded Search," Working Paper.
- Lewis, R., & Reiley, D. (2011), "Does retail advertising work: Measuring the effects of advertising on sales via a controlled experiment on Yahoo," Working paper.
- Malm, K., & Hamman, D. (2009), "The effects of display media on search traffic," *White paper*. Retrieved from <http://www.icrossing.com/sites/default/files/effects-of-display-on-search-traffic.pdf>
- Manchanda, P; Dube, JP; Goh, KY; et al. (2006), "The effect of banner advertising on Internet purchasing," *Journal of Marketing Research*, 43 (1): 98-108.
- Marketing Science Institute (2012). 2012-2014 Research Priorities. http://www.msi.org/MSI_RP12-14.pdf
- Moore, Robert S., Claire Allison Stammerjohan, Robin A. Coulter (2005), "Banner Advertiser - Web Site Context Congruity and Color Effects on Attention and Attitudes," *Journal of Advertising*, 34 (2), 71-84.

- Naik, P. A., & Peters, K. (2009), "A hierarchical marketing communications model of online and offline media synergies," *Journal of Interactive Marketing*, 23(4), 288-299.
- Naik, P. A., & Raman, K. (2003), "Understanding the impact of synergy in multimedia communications," *Journal of Marketing Research*, 375-388.
- Papadimitriou, P., Garcia-Molina, H., Krishnamurthy, P., Lewis, R. A., & Reiley, D. H. (2011), "Display advertising impact: Search lift and social influence," In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 1019-1027). ACM.
- Pesaran, H. H., & Shin, Y. (1998), "Generalized impulse response analysis in linear multivariate models," *Economics Letters*, 58(1), 17-29.
- Rutz O. J., Bucklin R. E., Sonnier G. P. (2012), "A latent instrumental variables approach to modeling keyword conversion in paid search advertising," *Journal of Marketing Research*, 49(3): 306–319.
- Rutz O. J., Bucklin R. E. (2011), "From generic to branded: A model of spillover in paid search advertising," *Journal of Marketing Research*, 48(1), 87–102.
- Shamdasani, P. N., Andrea J. S. Stanaland and Juliana Tan (2001), "Location, Location, Location: Insights for Advertising Placement on the Web," *Journal of Advertising Research*, 41(4), 7-21.
- Sherman, L., & Deighton, J. (2001), "Banner advertising: Measuring effectiveness and optimizing placement," *Journal of Interactive Marketing*, 15(2), 60-64.
- Sims, C. A. (1980), "Macroeconomics and Reality," *Econometrica*, 48(1), 1-48.
- Sims, C. A., & Zha, T. (1999), "Error bands for impulse responses," *Econometrica*, 67(5), 1113-1155.
- 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, 90-102.

- Van Heerde, H. J., Helsen, K., & Dekimpe, M. G. (2007), "The impact of a product-harm crisis on marketing effectiveness," *Marketing Science*, 26(2), 230-245.
- Van Heerde, H. J., Srinivasan, S., & Dekimpe, M. G. (2010), "Estimating cannibalization rates for pioneering innovations," *Marketing Science*, 29(6), 1024-1039.
- Wiesel, T., Pauwels, K., & Arts, J. (2011), "Practice Prize Paper—Marketing's Profit Impact: Quantifying Online and Off-line Funnel Progression," *Marketing Science*, 30(4), 604-611.
- Yang, S., A. Ghose (2010), "Analyzing the relationship between organic and sponsored search advertising: Positive, negative, or zero interdependence?" *Marketing Science*, 29(4) 602–623.
- ZenithOptimedia (2012), "Advertising Expenditure Forecasts December 2012," [White paper].