NET Institute*

www.NETinst.org

Working Paper #15-15

September 2015

**Towards a Digital Attribution Model:
Measuring the Impact of Display Advertising on Online Consumer Behavior**

Anindya Ghose and Vilma Todri
Stern School of Business, New York University

* The Networks, Electronic Commerce, and Telecommunications (“NET”) Institute, <http://www.NETinst.org>, is a non-profit institution devoted to research on network industries, electronic commerce, telecommunications, the Internet, “virtual networks” comprised of computers that share the same technical standard or operating system, and on network issues in general.

Towards a Digital Attribution Model: Measuring the Impact of Display Advertising on Online Consumer Behavior

Anindya Ghose and Vilma Todri*

Abstract

The increasing availability of individual-level data has raised the standards for measurability and accountability in digital advertising. Using a massive individual-level data set, our paper captures the effectiveness of display advertising across a wide range of consumer behaviors. Two unique features of our data set that distinguish this paper from prior work are: (i) the information on the actual viewability of impressions and (ii) the duration of exposure to the display advertisements, both at the individual-user level. Employing a natural experiment enabled by our setting, we use difference-in-differences and corresponding matching methods as well as instrumental variable techniques to control for unobservable and observable confounders. We empirically demonstrate that mere exposure to display advertising can increase users' propensity to search for the brand and the corresponding product; consumers engage both in *active search* exerting effort to gather information through search engines as well as through direct visits to the advertiser's website, and in *passive search* using information sources that arrive exogenously, such as future display ads. We also find statistically and economically significant effect of display advertising on increasing consumers' propensity to make a purchase. Furthermore, we find that the advertising performance is amplified up to four times when consumers are targeted earlier in the purchase funnel path and that the longer the duration of exposure to display advertising, the more likely the consumers are to engage in direct search behaviors (e.g., direct visits) rather than indirect ones (e.g., search engine inquiries). We also study the effects of various types of display advertising (e.g., prospecting, retargeting, affiliate targeting, video advertising, etc.) and the different goals they achieve. Our framework for evaluating display advertising effectiveness constitutes a stepping stone towards causally addressing the digital attribution problem.

Keywords: Online Advertising, Big Data, Analytics, Display Advertising, Advertising Effectiveness, Digital Attribution, Natural Experiment.

JEL Classification: L86, M37

* Anindya Ghose is a Professor of Information Systems and Marketing, Leonard N. Stern School of Business, New York University. E-mail: aghose@stern.nyu.edu. Vilma Todri is a Ph.D. candidate in Information Systems, Leonard N. Stern School of Business, New York University. E-mail: vtodri@stern.nyu.edu. We gratefully acknowledge the financial support from the NET Institute (www.NETinst.org).

1. Introduction

In 2013, for the first time in the history of the advertising industry in the U.S., digital ad spending surpassed TV broadcast advertising, which has traditionally been considered the most effective mass-marketing medium (IAB 2013). Thanks to the prevalence of digital advertising and various technological advancements, advertisers are nowadays able to track consumers' digital footprints at a more granular level and they can gain deeper insights into online consumer behavior as well as how advertising exposures might affect consumers' online behavior.

However, the consumer journey to an online purchase has become more complicated as consumers are likely to become exposed to multiple types of advertising (i.e., also known as 'multi-channel' exposures) during the purchase funnel. Besides, the usual problem of endogeneity that arises in the context of advertising is further exacerbated in the case of digital advertising due to the unprecedented possibilities for refined targeting. Hence, although digital advertising generates opportunities for greater measurability and accountability compared to traditional advertising media, there still exist significant challenges that hinder disentangling the different effects and imputing the true effectiveness of advertising. Given these challenges, it is not surprising that ad-hoc and naïve advertising performance evaluation models, such as the last-click attribution model, or simplistic performance metrics, such as the click-through rates (CTR), have prevailed in the advertising industry so far.

In order to accurately and fairly determine the effectiveness of specific advertising channels in spurring desired consumer actions, we need to understand what actually causes the desired changes in consumer behaviors. Towards achieving this goal, we first need to be able to draw causal inferences from observational data as randomized controlled experiments impose significant opportunity costs for advertisers who are often not keen on relinquishing the opportunity to advertise. Additionally, we need to understand the broad impact of advertising exposures on consumers' behaviors. This is especially important since some advertising exposures are consumer-initiated and, therefore, advertisers who would like to control the

presence and frequency of these advertising exposures on consumers' purchase funnel paths need to understand how they can increase consumers' propensity to initiate these exposures. For instance, a search advertising exposure is only triggered in a consumer's funnel path after the consumer initiates a search session expressing interest for a brand or a product. Therefore, advertisers who would like to control the presence and the frequency of search advertising exposures should examine what triggers consumers to initiate a search session.

This paper resolves some of the challenges that hinder measurability and accountability in digital advertising by drawing on economic methods that allow us to make causal inferences with observational data. Harnessing the value of big data and the ad-tech advancements that allow for more precise measurements, such as the viewability of impressions, we exploit an exogenous shock to the firm's targeting mechanism that simulates a natural experiment. The natural experiment framework allows us to compare the online behavior of two groups of users: those who view the display advertisements and those who eventually do not view the display advertisements while both groups are automatically targeted by the same marketing campaign. Therefore, this paper contributes to the nascent digital attribution literature with the proposed framework serving as the cornerstone of a digital display channel attribution model; it demonstrates how advertisers can attribute credit across the various types of display advertising (e.g., prospecting, retargeting, affiliate targeting, etc.), as well as how display advertising triggers other paid and non-paid advertising exposures in the funnel path. Second, this paper evaluates the impact of display advertising exposures across a wide range of consumer behaviors, rather than just a single proxy or metric. In particular, we examine whether display advertising exposures increase the users' likelihood to engage in *active search*, exerting effort to gather information, and *passive search*, responding to information sources that arrive exogenously, as well as the likelihood to increase the users' propensity to make an online purchase. Therefore, this study also contributes to the advertising effectiveness literature by investigating the effects of display advertising on a wide range of online consumer behaviors and the relative magnitudes

of the corresponding advertising effects. *Our data set indicates that on average 55% of the display ads are not rendered viewable*; other digital platforms report similar statistics (Horf 2014). Besides, to provide deeper insights into display advertising effectiveness, we move beyond the binary treatment of subjects that prior literature typically adopts; to our knowledge, we are the first to study and quantify the effect of the duration of consumer exposure to the display advertisement on online consumer behavior in a real-world setting. Additionally, we study the dynamics of the display advertising effects on online consumer behaviors taking into consideration how far down each consumers is in the funnel path, as indicated by the relative position of the touchpoint in the consumer funnel path.

We find that mere exposure to display advertising can significantly increase users' propensity to search for the brand and the corresponding product. In particular, display advertising leads consumers to engage both in *active search*, exerting effort to gather information, as well as *passive search*, using information sources that arrive exogenously. The most prominent effect of display advertising is the increased propensity of consumers to engage in active search; consumers are up to 36.07% more likely to make a direct visit on the advertiser's website after a display advertising exposure and up to 25.7% more likely to initiate and click on a brand-related search engine session, relative to the mean of these activities. However, the effect of display advertising in passive search behavior is not negligible; consumers can be up to 28.46% more likely to click on a future display advertisement when targeted early in the funnel paths. In general, studying the dynamics of these effects, we find that advertising effects are amplified up to four times when consumers are targeted earlier in the purchase funnel path, and the longer the duration of an exposure to display advertising, the more likely the consumer is to engage in direct search behaviors (e.g., direct visits) rather than indirect ones (e.g., search engine inquiries). Furthermore, display advertising directly increases consumers' propensity to make a purchase. Our framework for evaluating display advertising effectiveness also constitutes a stepping stone towards causally addressing the digital attribution problem.

2. Related Work

This paper is related to several streams of research in the fields of Information Systems and Marketing pertaining to online advertising. Regarding the effectiveness of display advertising, its performance was initially assessed with simple proxies –such as the click-through-rate (CTR)– in an effort to capture consumers’ active response to advertising, not just a probable exposure to it (Hollis 2005). However, such proxies demonstrated decreasing performance over time, generating concerns whether display advertising is indeed effective. Nevertheless, the CTR performance metric is only an intermediate proxy for other quantities of interest to the advertisers (Lewis et al. 2013). Moving beyond CTRs, researchers evaluated the impact of display advertising exposures on consumers’ purchase intentions (Lewis and Reiley 2010; Manchanda et al. 2006) and other consumer behaviors, such as brand recall (Dreze and Hussherr 2003) or brand-specific page views (Rutz and Bucklin 2012). Moreover, various researchers have attempted to identify contextual factors that enhance display advertising effectiveness, such as targeting or obtrusiveness (Goldfarb and Tucker 2011). This paper adds to the literature of display advertising by evaluating its effectiveness across a wider spectrum of consumer behaviors –such as active and passive search behaviors as well as the propensity to convert–, by studying the dynamics of these effects on the funnel path, and evaluating their relative magnitudes.

This paper is also related to the stream of literature studying spillover effects and synergies in multi-channel advertising. As multi-channel advertising became prominent, researchers attempted to evaluate the complementarity effects employing aggregate-level data. For example, Naik and Raman (2003) adopt an integrated marketing communications perspective to emphasize the importance of synergy in multimedia activities. In addition, prior research has also studied spillover effects between specific advertising channels. For instance, Joo et al. (2013) demonstrate that television advertising increases consumers’ search behavior, as reflected in the number of product category-relevant searches and the brand’s share of keywords searched.

Ghose (2014) report a field experiment to examine cross-media synergies between web and mobile advertising, and demonstrated that for many brands a mix of web and mobile display advertising triggers, more clicks and conversions and higher sales amounts than web or mobile ads alone. In the same vein, Xu et al. (2015) show in the context of cross-device browsing behavior that the tablet channel acts as a substitute for the PC channel while it acts as a complement for the smartphone channel. That is, an increase in ecommerce sales on smartphones can be attributed to the introduction of tablets. Focusing on display and search advertising, Kireyev et al. (2013) consider the interaction and dynamic effects of search and display advertising with aggregate-level data while Papadimitriou et al. (2011) study the impact of display advertising only on search engine queries employing a campaign-level analysis. In contrast, in this paper we employ individual-level data to investigate multiple interactions between display ads and online search behavior. Individual-level data does not entail the loss of information that aggregate-level data does and provides the ability to control for unobserved heterogeneity at the consumer and advertisement level. Compared to prior literature, in this study, we also examine the impact of display advertising on *both active search behaviors*, requiring users to actively gather information, *and passive search behaviors*, under which users are employing information sources that arrive exogenously. Besides, we also study higher engagement levels of active search behavior beyond search engine queries, such as direct visits to the advertiser's website. Finally, to the best of our knowledge, this is the first paper to study the effect of display advertising beyond the usual binary treatment that prior literature adopts; we evaluate the impact of the duration of exposure to display advertisement on online consumer behavior in a real-world setting.

Another closely related stream of literature is the emerging work on digital attribution. Digital attribution is concerned with allocating credit for a consumer's purchase across the marketing channels to which the consumer was exposed during her purchase funnel path. Given the plethora of advertising channels, media, and targeting techniques, it is crucial for advertisers and

marketers to disentangle the influence of each channel in order to better optimize their return of investment (ROI), among other advertising goals. Most of the related works in this stream of literature conduct an empirical analysis for the channel attribution problem with a few exceptions that propose a game-theoretic approach (Berman 2013). The existing empirical digital attribution models capitalize on the variation of advertising exposures at the individual level across consumers in order to estimate the real effect of advertising. For instance, Abhishek et al. (2012) develop a hidden Markov model of individual consumer behavior where they define as conversion various actions, such as requesting a quote and looking up a dealer's address, rather than making an purchase. Similarly, Li and Kannan (2013) develop a nested logit model to examine the nature of carryover and spillover effects of prior visits to a firm's website through different channels while Zantedeschi et al. (2013) develop a Bayesian Tobit model drawing on marketing mix models to measure the effectiveness of advertising exposures. Likewise, Shao and Li (2011) develop a bagged logistic regression model to quantify the attribution of different advertising channels. Capturing interdependencies among advertisement clicks, Xu et al. (2014) develop a stochastic model for online purchasing and advertisement clicking that incorporates mutually exciting point processes.

However, Lewis and Reiley (2010) have demonstrated that without an experimental framework, using methods based on endogenous cross-sectional variation in advertising exposure, one could obtain a very inaccurate estimate of advertising effectiveness. Furthermore, when making an online purchase, consumers are influenced by multiple other factors. For instance, Bronnenberg et al. (2012) find strong evidence that past experiences are an important driver of current consumption. Thus, we propose an incremental lift approach to the digital attribution problem. In particular, recognizing that the consumers who are targeted through an advertising channel usually have a non-zero propensity of conversion, as formulated by prior experiences and the reservoir of brand equity, we propose that the digital attribution model should be based on the incremental lift that advertising exposures have towards key marketing objectives.

3. Data

Our unique data set was collected in collaboration with a large online media analytics and optimization platform company that manages the entire campaign of an online U.S.-based retailer.¹ The advertiser runs paid keyword advertising campaigns across multiple popular search engines as well as display advertising campaigns across various websites employing multiple targeting techniques (e.g., retargeting, prospecting targeting, affiliate advertising, etc.). The data spans all online advertisements that were run by the company during a period of six months, from May to October 2013; during this period, the company did not engage in other advertising activities apart from the ones we observe in our data set.

In particular, we have an individual-level data set consisting of advertising exposures and user-initiated actions, with users tracked across different advertising channels and media. The data set contains advertising exposures and user-initiated actions corresponding to all the display and search advertising exposures as well as consumers' direct visits to the advertiser's website. More specifically, an observation in our dataset refers to a display impression or click, a search click, or a direct visit to the advertiser's website. Furthermore, in our data set we are able to distinguish between brand and non-brand related search queries based on the campaign that was triggered to serve the search advertisement. In particular, the corresponding campaigns have very refined targeting criteria that prevent a non-brand search advertising campaign from serving an advertisement to a brand related search query. Additionally, we have information regarding the advertising network, the type of targeting employed and the specific advertising campaign id through which every display impression was served.

Given that our corporate partner advertiser sells experience goods, the consumers usually spend some time conducting research on the product before they decide to make an online purchase, if any. Therefore, as one would expect, the consumers are exposed to multiple touchpoints across

¹ Due to the nature of the data-sharing agreement, we are unable to reveal the name of the firm.

their funnel paths related to various advertising channels and formats. Finally, the website of the retailer does not allow consumer reviews to be posted online; thus, we can study advertising effects in the absence of user-generated content interactions.

3.1. Consumer Funnel Paths and Viewability of Impressions

We construct the consumer funnel paths by connecting the touchpoints that are related to the same unique customer identifier. If a funnel path is successful, we do observe the corresponding conversion, which means that the consumer made an actual purchase on the advertiser's website. Each time a consumer exits a funnel path successfully, she enters a new funnel path as a returning customer (i.e., the customer does not cease to exist). In total, the data set contains observations corresponding to all the display and search advertising exposures as well as direct visits for all the consumers who had at least one advertising exposure in their corresponding funnel paths. Table A1 of the online appendix reports the summary statistics of our data set at the observational level.

Our data set also includes two unique features that provide granular information regarding the advertising exposures. First, we have information on the actual viewability of the display impressions. In particular, we know when an impression was visible on a consumer's screen area for more than one second and, hence, whether it is rendered viewable. Second, we have information on the duration of exposure to the display advertisements once they were rendered viewable. The viewability of advertising exposures is a new metric in the field of online advertising; it was introduced in 2012 but nowadays has gained significant popularity redefining the measurement standards of the field (e.g., see Fig.1 of online appendix). To our knowledge, this is the first paper to exploit the newly established metric of viewability in order to impute the effectiveness of advertising from observational data and tackle the channel attribution problem.

Regarding the viewability of the advertising exposures, there are many circumstances under which an impression of a display ad may be rendered non-viewable. For instance, when the

display is loaded on the website but the consumer does not happen to scroll down to the area that would render the impression viewable. Additionally, several contextual factors, such as the browser window size, the screen resolution, and the screen orientation of the user might determine whether an advertising impression will be rendered viewable. Other possible scenarios include circumstances under which the viewer does not have the appropriate plug-ins for an interactive ad to be displayed, or when the viewer utilizes some type of ad blocker software – so that even if the browser loads the impression, the user never views the ad.. An impression is also rendered non-viewable when the publisher places an image or another layer overlapping the ad. For instance, *The New York Times* publisher displays such a layer when a non-subscriber attempts to view more articles than their paywall limit allows.² Finally, based on this information, we filter out impressions that are loaded by non-human technologies, such as crawlers and web proxies.

As far as the collection of information on viewability is concerned, currently there exist two main ways to measure ad viewability. The first way to measure ad viewability is to utilize the geometric method and it typically involves comparing the position of the four corners of the ad relative to the host webpage and then comparing the four corners of the browser's viewport relative to the host webpage. Comparing these two, advertisers can make inferences about whether the ad is within the viewport. Another variant of the geometric method approach involves comparison of the screen rather than the host page. Comparisons are also made between the ad and the mouse cursor, and the mouse cursor and the viewport. The second way to measure ad viewability is to monitor browser optimization functions. By monitoring how a browser allocates resources to render an ad, one can determine whether an ad was rendered viewable. The

² Additionally, an impression might be rendered non-viewable due to ad frauding techniques that an advertising network might employ. For example, advertising networks might serve multiple ads on top of each other in a legitimate ad slot where only the ad on the top can be rendered viewable (i.e., ad stacking). Similarly, advertising networks might employ pixel stuffing where many pixel frames are placed all over a webpage and trigger multiple ad impressions that are invisible in the naked eye.

browser-optimizations approach may be used to measure the viewability of ad impressions across all major browsers even when the ads are embedded in (nested) unfriendly iframes.

3.2. Big Data Infrastructure

In order to estimate the empirical models presented in Section 4, we used a high-performance computing cluster. The employed cluster³ constitutes a powerful and reliable high-performance computing infrastructure that provides advanced computing technologies and allows us to efficiently manage the data of complex, high-volume computational processes. Tapping into the unique characteristics of our big data set and taking full advantage of our high-performance infrastructure, we parallelize the execution of our estimation procedures to make feasible the estimation of our empirical models in real time. For instance, the proposed individual-level difference-in-differences model takes on average 0.87 seconds to be estimated for our data set using a single computing node with 20 processors. Furthermore, the total time complexity grows linearly with the number of observations; for N training examples and F features the total time complexity asymptotically is $O(F^2 N)$. Additionally, apart from using in parallel multiple computing nodes, we also parallelize around 80% of the most computationally expensive procedures achieving a significant increase in performance; at least more than four times faster compared to using the same hardware infrastructure without using this second level of parallelization.

4. Empirical Methods

Our identification strategy for estimating the effects of display advertising is based on the unique feature of our massive data set, pertaining to the viewability of the display ad impressions and the duration of exposure to the display advertisement. Previous academic research on display advertising assumed that an impression was always viewable to the user whenever loaded.

³ The specific cluster consists of 160 nodes with more than 3200 CPU cores featuring the latest Intel Xeon-based technology for supercomputers and more than 16 terabytes of memory running a high-performing enterprise class UNIX-based operating system. The theoretical peak performance limit of our parallel supercomputing systems is measured to more than 75 teraFLOPS.

However, *our data set indicates that on average 55% of the display ads are not rendered viewable*; other digital platforms report similar statistics (Horf 2014). Such measurement issues are very important and might have led to biased estimates of the effects of display advertising in previous research studies. Hence, tracking the viewability of impressions greatly enhances the veracity of advertising data and alleviates the usual attenuation bias of measurement error which is further magnified in the case of multivariate regression models and panel data models (Griliches and Hausman 1986).

The discussed circumstances that affect the viewability of a display ad impression serve as an exogenous shock to the firm's targeting and simulate a natural experiment creating two groups of users: those who view the display ad and those who do not; both groups are automatically targeted by the same marketing campaign fulfilling certain targeting criteria. Hence, only some of the customers who are targeted through the campaigns are indeed 'treated' (i.e., their impression is rendered viewable). Our natural experiment framework avoids the self-selection and other treatment selection biases, as the viewability of the display ad mimics the exogeneity of a randomized experiment. The treatment and control groups are similar in every other way as they fulfill the advertiser's targeting criteria and do not exhibit systematic differences.

Nevertheless, we also employ alternative methods and use various falsification tests to examine and rule out alternative explanations. Another significant advantage of the proposed estimation methodology is that it can be parallelized and scaled to big data sets, as shown in this study.

4.1. Individual-Level Difference-in-Differences

Let i be a consumer who is targeted with a display advertisement. At display ad occasion j , individual i belongs to group $G_{ij} \in \{0,1\}$, where group $g = 1$ indicates the treatment group (i.e., each individual who was targeted with a display ad *and* her impression was rendered viewable). For each occasion, individuals belonging to both the treatment and control groups are observed before and after the treatment time period, $T_{ij,t} \in \{0,1\}$ where $t = 1$ indicates the post-treatment

period. For $i = 1, \dots, N$, a random sample from the population, the group identifier of individual i at occasion j , G_{ij} , and the time period T_{ij} can be treated as random variables. We denote by Y_{ijt} the outcome that is observed for individual i at occasion j in the time period t . Let Y_{ij1}^0 denote the outcome for individual i for the after period, if she does not receive the treatment at occasion j , and Y_{ij1}^1 the outcome for the same individual for the after period, if she does receive the treatment at occasion j . Only one of these outcomes is realized as for each display ad impression occasion one individual either belongs to the treatment or the control group (i.e., for individual i at each occasion j during the after period, only Y_{ij1}^0 or Y_{ij1}^1 is observed). Then, the parameter of interest in this setting is the mean impact of treatment on the treated. Thus, the DID estimate is:

$$\begin{aligned} \tau^{DID} = & \left[\mathbb{E}[Y_{ij1}^1 | G_{ij} = 1, T_{ij} = 1] - \mathbb{E}[Y_{ij0}^1 | G_{ij} = 1, T_{ij} = 0] \right] \\ & - \left[\mathbb{E}[Y_{ij1}^0 | G_{ij} = 0, T_{ij} = 1] - \mathbb{E}[Y_{ij0}^0 | G_{ij} = 0, T_{ij} = 0] \right]. \end{aligned}$$

In particular, in the DID estimator, the average gain over time in the non-exposed (control) group is subtracted from the gain over time in the exposed (treatment) group. This double differencing removes biases in second period comparisons between the treatment and control groups, which could be the result of permanent differences between these groups, as well as biases from the comparison over time in the treatment group, which in turn could be the result of time trends unrelated to the treatment (Heckman et al. 1998a; Imbens and Wooldridge 2009).⁴ In this set-up, we observe the outcomes for the treated and untreated groups before and after the treatment (Y_{ij0}, Y_{ij1}). For instance, when evaluating the brand search behavior of the user, the outcome Y_{ij0} indicates whether the user had or had not searched before the treatment and the outcome Y_{ij1} indicates whether she initiated a search session during the second period (i.e., after the treatment, if treated), as captured by the next touchpoint in the funnel path.

⁴ When the treatment effect is not constant across individuals, the standard DID estimate can be interpreted as the average effect of the intervention on the treatment group.

We generalize this framework to allow for general forms of heteroscedasticity by relaxing the full independence assumption to only mean independence (Athey and Imbens 2006).

Additionally, we control for unobserved heterogeneity at the individual level, by allowing for time-invariant, individual-specific fixed effects ξ_i , potentially correlated with G_{ij} . Furthermore, we also allow for ad-level heterogeneity with fixed effects φ_j . Denoting the treatment effects across individuals by $\tau (= Y_{ijt}^1 - Y_{ijt}^0)$, the realized outcome satisfies:

$$Y_{ijt} = a + \beta T_{ij} + \gamma G_{ij} + \tau T_{ij} * G_{ij} + \varphi_j + \xi_i + \varepsilon_{ijt}.$$

The β coefficient represents the time effect, the γ coefficient the targeting effect, and the error term is normalized to have a zero mean. In order to make the interpretation of our results more straightforward, we employ a linear probability model (LPM) that yields results in terms of probability changes and allows for the coefficients to be comparable across models and groups. LPMs are unbiased and consistent estimates of a variable's average effect, specifically the average marginal effect, on $P(y = 1)$ (Wooldridge 2010). Therefore, if one is interested in the average effect estimate, it is entirely appropriate and reasonable to choose LPM over other methods, such as logistic regression (Mood 2010).

4.2. Difference-in-Differences Matching

The aforementioned DID estimation methodology constitutes a popular method for estimating average treatment effects (ATE) that controls for unobservables. The key underlying assumption of DID that differences between treatment and control groups would have remained constant in the absence of the treatment does not always hold though. In fact, one cannot evaluate the plausibility of this assumption with two time periods. Thus, we use additional estimation methods, such as matching on observables, to corroborate our findings. Matching estimators constitute an alternative methodology for estimating ATEs. In our context, matching estimators allow us to evaluate the ATE, even if the generating process of viewable and non-viewable display ads had not been completely random.

Heckman et al. (1997) developed the generalized *difference-in-differences matching estimator* that allows for temporally time-invariant differences in outcomes between the treated and non-treated individuals. They find that the generalized difference-in-differences matching estimator is more effective than conventional matching methods, such as propensity score matching, in removing biases from the data, especially when it is contaminated by temporally invariant components of bias, such as unobserved effects. The nonparametric conditional DID matching estimator is a two-stage evaluation methodology, which first estimates the probability of an individual receiving the treatment and then uses the estimated probability. The DID matching estimator is then:

$$\tau^{DDM} = \frac{1}{n_1} \sum_{i \in I_1 \cap S_p} \left\{ (Y_{ij0}^1 - Y_{ij1}^0) - \sum_{k \in I_0 \cap S_p} w(i, k) (Y_{kj0}^0 - Y_{kj1}^0) \right\},$$

where the weights w are provided by the cross-sectional matching estimator employed, I_1 denotes the set of treated individuals, I_0 the set of non-treated individuals, and S_p the region of common support. Essentially, in the first step, we initially estimate the propensity score and then compare observations that have identical (or very similar) scores. Such matching methods allow the mechanism of allocation to treated and untreated groups to be not completely random (e.g., users who have visited the website more frequently in the past might be more likely to have a viewable impression). We also control for such differences by calculating the probability of treatment based on consumers' past exposures and online behavior.

In order to ensure that the proposed matching method reliably estimates the ATE, we assess the quality of the matching method by evaluating the area of the common support (i.e., overlap) between the estimated probability densities of the propensity score for the treated and non-treated individuals. The overlap describes the extent to which the range of the estimated propensity score is the same across the two treatment groups. As shown in Figure 1 and Table 1, there exists substantial overlap of the propensity score densities for the two groups, which

assures us that we can reliably estimate the ATE with our data set. In particular, the significant overlap in the propensity score densities enables the matching method to find control individuals that are similar to treated individuals and treated individuals that are similar to control individuals, respectively. Therefore, the proposed matching method can reliably use the observed counterfactuals that exist in our data set and generate causal inferences without the need to depend on model-based extrapolations beyond the support of the observed data (Rubin 1997). Furthermore, conditioning on the propensity score should result in distributional balance of the observed covariates between the treatment groups. Hence, we further assess the quality of the matching method by evaluating the balance of the observed covariates. In particular, following Austin (2011), we estimate the standardized differences in means and the variance ratio for each observed covariate between treatment and control groups. As shown in Table 2 and Figure 2, matching on the estimated propensity score has balanced the covariates since the standardized differences are all close to zero, and the variance ratios are all close to one.

4.3. Identification with Potentially Endogenous Treatment

The aforementioned empirical models control for possible sources of endogeneity in various ways. For instance, the individual-level panel data difference-in-differences estimator controls for unobserved time-invariant individual-level effects and, therefore, accounts for individual-specific sources of endogenous selection into the treatment group. Similarly, the difference-in-difference matching estimator allows selection to treatment as a function of past browsing behavior, past advertising exposures, and consumers' responses to these exposures. In addition, it controls for temporally invariant differences in outcomes between treated and untreated individuals (Heckman et al. 1998b; Heckman et al. 1997). Nevertheless, despite these various controls for potential endogeneity, there might still exist potential unobserved time-varying confounders that could potentially result in inconsistent estimates of the effect of an advertising exposure. For example, users who tend to browse webpages with content related to the product itself might be more likely to engage with the content of the webpage for longer periods of time

and, thus, the corresponding advertising impression might be more likely to be rendered viewable. Hence, under such a scenario, it would be challenging to identify the effect of the advertisement itself, disentangled from the inherent interest of the users who choose to browse online content relevant to the advertised product.

4.3.1. Controlling for time-varying unobserved confounders

In order to address any concerns regarding potential time-varying unobserved confounders, we exploit the characteristics of variety and volume of big data. The variety and volume of big data can contribute in alleviating various biases from previously unobservable confounders. More specifically, in our context, we extend the individual-level panel data difference-in-differences model specification controls for targeting criteria and mechanisms at three different levels; unobserved factors related to targeting mechanisms are time varying as every attempt by the advertiser for an advertising exposure is characterized by different targeting criteria and is served through different advertising networks and campaigns. *First*, we control for the specific type of targeting of the specific advertising exposure attempt. Controlling for the type of targeting, we control for cases under which the content of the webpage is relevant to the ad being shown (i.e., contextual targeting) and we are able to purge the coefficient of interest from potential unobserved targeting confounders. *Second*, we control for the advertising network and affiliate partner that serve the advertisement. Extending our specification by controlling for the network that served the advertisement, we are able to control for different mechanisms and incentives that might come into play when networks position the advertiser's ads. For instance, various advertising networks utilize different technologies to estimate the match of the ad with certain webpages. *Third*, we also control for the specific campaign that triggered the advertising exposure. The targeting criteria for serving advertisements are defined at the campaign level; each campaign corresponds to tightly defined targeting criteria as such categorization allows the evaluation of the effectiveness of different type of targeting methods and specifications at an aggregate level. Hence, controlling for the campaign through which the advertisement was

triggered, we are able to control for cases under which the webpage content is relevant to the ad being shown as well as for more granular targeting factors, such as different ways to employ contextual targeting (e.g., topic versus keyword contextual targeting).

Furthermore, the potential problem of the omitted variable bias (e.g., bias introduced by omitting how relevant the website content is with the advertisement) can be addressed by using a proxy variable for the unobservable confounder of relevant webpage content. In our case, we use the viewability ratio of the campaign through which the specific advertising exposure was attempted as a proxy variable that measures the relevance of the content of the website the consumer was browsing. More specifically, under the scenario that users who browse online content that is more relevant to the advertised product are more likely to have their advertising impressions rendered viewable, advertising campaigns that served impressions in webpages with relevant content should be characterized by higher viewability ratios. Hence, in order to address this potential limitation, we extend our main specification by introducing the viewability ratio of the campaign as a proxy for serving ads to websites that match the advertisement itself.

4.3.2. Instrumental Variables Fixed Effects

In order to further investigate concerns for potential time varying unobserved effects as well as to investigate more general concerns about any type of unobserved confounders that could potentially introduce biases in the estimates of interest, we additionally employ the instrumental variables fixed effects estimator. As instrument, we use information on hyper-local weather for very granular time intervals (i.e., 20 minute time intervals) matched with individual-level advertising exposures. Weather data is correlated with the (potentially endogenous) explanatory variable of a viewable impression because browsing the Internet is an activity that competes with other outdoors activities the users might be enjoying. In particular, bad weather conditions (e.g., level of precipitation) reduce the opportunity costs of outdoor activities, such as sports, and, thus, increase the attractiveness of browsing the Internet. Therefore, *ceteris paribus*, as consumers spend more time browsing the Internet, an advertising impression is more likely to be rendered

viewable. On the other side, bad weather conditions are not correlated with the dependent variables as the demand for the product is not driven or affected in any way by weather conditions. In order to employ weather information as instruments, we collected weather data from the National Weather Service (NOAA) by mapping the latitude and longitude of each individual to the closest weather station in order to gather granular level information on precipitation, temperature, and humidity levels.

5. Empirical Results

In this section, we present and discuss the results of our empirical models. Each of the following subsections analyzes a set of outcome variables that are properly categorized to capture similar effects on consumer behavior changes. The corresponding results for each outcome are presented and discussed individually, focusing on the average treatment effect (ATE) of display advertising on various consumer behaviors, as shown in Table 3; the dynamics of the effects of display advertising in the funnel path, as shown in Table 4; and the effect of the duration of the treatment shown in Tables 5 and 6, respectively. Finally, the last subsection discusses how display advertising effectiveness varies by the type of targeting deployed by the advertiser, as shown in Tables 9 and A8. Noting a few things here will make the interpretation of the results more straightforward. First, the display advertising effectiveness across all the tables is captured by the ‘ATE’ variables. The average treatment effect is measured as the change in a desired consumer activity, such as purchase, which is the result of a randomly drawn individual from the population of users targeted by the advertisers’ marketing campaigns being successfully exposed to a display advertisement. Second, the ‘treated’ variable controls for any systematic differences in consumers that are targeted; therefore, the ‘ATE’ variables are not contaminated by such effects. Third, the advertising exposure variables in Tables 4 and 6 (i.e., ATE exposure 1-5) indicate the relative position of the display advertising exposure in the consumer funnel path, rather than the absolute positions. As far as the goodness of fit of the models is concerned, given the particular econometric specifications, the R-squared indicates a good fit of the model to our

data in alignment with previous studies in this area (Forman et al. 2009; Tucker 2014); fitting a mean deviated model, the effects for individuals are simply subtracted out of the model and, thus, their overall effect is not quantified on the fit of the model leading to seemingly low within R^2 measures (Baltagi 2008). Finally, for each specification, additional robustness tests have been conducted and presented in Section 6.

5.1. Spurring Interest in the Brand

The coefficients of the impact of display advertising (‘ATE’ on Table 3) are positive and significant for all the consumer behavior outcome variables, indicating that display advertising generates and stimulates the consumers’ interest for the brand. In particular, we empirically show that consumers engage both in *active search*, exerting effort to gather information and *passive search*, using information sources that arrive exogenously. More specifically, an exposure to a display advertisement significantly increases both the chances of a consumer to visit the website directly and the chances to initiate a search session on any search engine using brand-related search queries.

(Insert Table 3 About Here)

Among the variables that capture an increased interest in the advertiser’s brand (i.e., direct visit, search engine inquiries, click on future display ads), we show that direct visits to the website is the most prevalent effect of display advertising. Specifically, an additional exposure to display advertising increases direct visits website activity (after the exposure) by an average of 1.01%, as shown in Column 2 of Table 3. This estimate is also of important economic significance as the coefficient of interest suggests that exposure to a single advertising impression increases consumers’ probability to make a direct visit to the advertiser’s website by 36.07%, relative to the mean of direct visit probability. Additionally, targeting the consumers earlier in the funnel path can increase the aforementioned effect by up to 2.82%, as shown in Column 2 of Table 4.

Going beyond the binary treatment, we can see that an additional minute of exposure to a display advertisement can direct visit activity by 0.26%, as shown in Table 5. This suggests that the longer the duration of the consumer's exposure to the display advertisement, the more likely the consumer is to make a direct visit, rather than visit the website through the rest of the channels. Also, the duration of the exposure has a positive and significant effect on consumers' propensity to make direct visits to the advertiser's website throughout the funnel path, as shown in Table 6. These findings have important implications for marketers because, contrary to other display advertising effects, the search behavior exhibited through direct visits requires a higher level of engagement from the users; consumers need to recall the brand and remember the retailer's website in order to make a direct visit sometime after being exposed to a display advertisement. Moreover, as a result of display advertising, consumers engage in active search behavior collecting information through other channels as well. In particular, one display advertising exposure increases brand search activity by an average of 0.72%, as shown in Column 1 of Table 3. The coefficient is also of important economic significance as the exposure to a single advertising impression increases consumers' brand search intent by 25.7%, relative to the mean of brand search activity. Interestingly, this effect becomes up to four times larger (3.02%) when the advertiser targets the consumer earlier in the funnel path, as shown in Column 1 of Table 4. Hence, display advertising can have a positive impact on consumers' propensity to search, even when they do not directly engage with the display advertisement itself. Additionally, as shown in Table 6, an additional minute of exposure to the display advertiser can affect consumers' search engine inquiries when this exposure takes place very early in the funnel path.

(Insert Tables 4, 5 & 6 About Here)

The results detailed above suggest that if advertisers do not account for spillovers from display advertising, they might overestimate the effectiveness of search advertising towards increasing consumers' propensity to convert. Accordingly, advertisers might underestimate the effects of display advertising if they do not account for spillovers to other types of search behaviors, such

as direct visits and search engine inquiries. Many single-source ad-hoc attribution models, such as the last-click channel attribution model, would give the whole credit to search paid advertising, if that was the last touchpoint of a successful funnel path that concluded with a conversion. Similarly, such models would not give any credit at all to the display advertising that lead to a direct website visit, if direct visit last touchpoint of a successful funnel path that concluded with a conversion. However, as our results indicate, a consumer might choose to perform a branded search or make a direct visit as a result of a display advertising exposure. Finally, display advertising leads consumers to engage in passive search using information sources that arrive exogenously to them. Specifically, display advertising makes consumers more responsive to future display advertisements as it increases display advertisement clicks activity by an average of 0.37%, as shown in Column 5 of Table 3. Apart from statistically significant, the estimate is of important economic significance as well given that the coefficient of interest suggests that exposure to a single advertising impression increases consumers' intention to click on a future display advertisement by 28.46%, relative to the mean of display click advertisement probability. Similar effects are observed for the duration of the exposure to the display advertisement, as shown in Column 5 of Tables 5 and 6.

5.2. Spurring Interest in the Product

The coefficients of the impact of display advertising on the consumer behavior outcome variables that capture interest for the product category (i.e., non-brand paid advertising, organic search results) are all positive and significant, indicating that display advertising generates and stimulates the consumers' interest for the product category in general and beyond the brand itself. However, compared to the effects of the brand interest, those of the product interest are smaller in magnitude and exhibit decreasing marginal effects.

In particular, a display advertising exposure increases visits through the organic results of search engines, by 0.54% on average, as shown in Column 3 of Table 3. This coefficient suggests that

exposure to a single advertising impression increases consumers' intent to visit through an organic search result by 20.77%, relative to the mean of organic search visits. Similarly, a display advertising exposure increases the generic keyword sessions and clicks on the advertiser's paid-search advertisement, by 0.31% on average, as shown in Column 4 of Table 3. Apart from statistically significant, the estimates are also of important economic significance as the coefficient of interest suggests that exposure to a single advertising impression increases consumers' intent to visit through generic keywords search sessions by 17.76%, relative to the mean of generic search visits. Such effects decrease later in the funnel path, as shown in Table 4, because as consumers move further down the funnel path, they shape their preferences. Finally, on average, increasing the duration of the exposure to display advertising does not have an overall significant impact on the aforementioned outcomes, as shown in Table 5. However, studying the dynamics of these effects, we see that an additional minute of exposure to display advertising can increase consumers' propensity to initiate and click on generic keywords sessions very early in the funnel path, as shown in Table 6.

We show that display advertising can have positive spillover effects to other brands in the same product category, based on the increase in non-brand search inquiries. In other words, display advertising might not only benefit the advertiser, but also some competitors, since a user is more likely to search for the product using a generic keyword and might end up choosing a competitor's product offering.

5.3. Increasing consumers' propensity to convert

Up to this point, we have discussed the effects of display advertising on spurring the interest in the brand and the product, as exhibited by various active and passive search consumer behaviors. Next, we discuss the impact of display advertising on increasing consumers' propensity to convert and successfully complete their current funnel paths. In prior literature, the findings regarding the ability of display advertisements to drive conversions are conflicting and there is a debate regarding whether display advertising is an effective means of generating direct sales.

We evaluate the likelihood of purchase after a successful display advertising exposure at the completion of the funnel path. Evaluating the impact of display advertising on increasing consumers' propensity for conversions, the effect is positive and significant across the whole funnel path, as shown in Column 6 of Table 4. Apart from statistically significant, the estimates are also of important economic significance as the coefficient of interest suggests that exposure to a single advertising impression increases consumers' purchase intent by 7.1%, relative to the mean of conversion probability. Similar effects are observed regarding the duration of the exposure to the display advertisement, as shown in Column 6 of Table 6. These results suggest that display advertising has a positive impact on the decision of consumers to convert.

Overall, evaluating a wide range of consumer behaviors, we show that an advertising exposure is able to create and sustain an increased interest for the brand and product.

5.4. Results under Potentially Endogenous Treatment

In this section, we discuss the results of various alternative methods we employ which allow for potentially endogenous allocation mechanisms into the treatment group. First, Table 7 shows the results of the difference-in-differences matching estimator; we find that our results remain qualitatively very similar. Furthermore, the results of the additional specifications that control for potentially time-varying unobserved confounders with various controls for targeting mechanisms are shown in Tables A3-A5 of the online appendix. Again, we find that the estimates of interest remains robust after controlling for potential time-varying unobserved effects related to targeting mechanisms and advertising networks. Similarly, Table A6 of the online appendix, presents the results of the specification employing a proxy variable for the potential unobserved confounder and shows that the estimates remain very robust.

Moreover, the results of the instrumental variable (IV) fixed effects methods are presented in Table 8. In particular, Table 8 presents the results of the panel data IV method deploying hyper-local granular weather information at the 20-minute interval of the advertising exposure as an

instrument. Furthermore, Table A7 of the online appendix presents the results of the panel data instrumental variable method employing just the precipitation level at the 20-minute interval of the advertising exposure as an instrument. Overall, tapping into the variety of big data sets by augmenting our data with hyper-local weather information during the advertising exposures, we are able to employ the IV method and find that our results remain qualitatively very consistent, which ensures us that the coefficients of interest regarding the effectiveness of a successful advertising exposure are not driven by time-varying unobserved effects or other unobserved confounders. Besides, while each of the employed strategies for assessing causal effects is based on different assumptions, yet all produce convergent results and, hence, enhance the credibility of our results.

5.5. Effectiveness by Type of Display Advertising

Digital attribution pertains to attributing credit not only across the different channels an advertiser employs to promote and attenuate the marketing message in the online world but also across the different mediums and platforms. Hence, an additional step of our analysis is the evaluation of the effectiveness of display advertising across the different types of targeting and mediums.

During the six-month period of our data set, five different types of display targeting were deployed by the advertiser. The first type of display advertising is related to the affiliate advertising; a type of performance-based targeting, according to which the affiliate channels are getting rewarded by the advertiser for bringing in consumers through their own marketing efforts. Based on the results presented in Table 9, affiliate advertising increases consumers' propensity to search for the brand through search advertising, by an average of 0.38%, and through direct visits to the advertiser's website, by an average of 0.70%. Furthermore, this type of advertising is effective in the long run for generating sales, as a display advertising exposure by this type of targeting is associated with a 0.40% increase in the non-immediate conversions, as shown in Table 9.

(Insert Table 9 About Here)

Moreover, branding pre-roll advertising, which entails a limited 15 to 30 seconds video advertisement auto-playing before a user-selected video content, has a very strong positive and significant effect on the consumers' propensity to visit the advertiser's website through organic search results (i.e., increases the chances by 102%); while it does not seem to increase consumers' propensity to make a conversion, as shown in Table 9.

Retargeting, a form of online advertising that allows marketers to target consumers based on their previous actions, decreases the chances that the user will visit the advertiser's website after the exposure by 1.28%, as shown in Table 9. A possible explanation for this finding is that retargeting might be annoying for many consumers, causing them to actively avoid searching for the brand immediately after an exposure, but in the long run it seems effective as it is increasing consumers' propensity to convert by an average of 0.81%.

Lastly, prospecting display advertising, which deals with finding new customers early in their funnel path, and CPA advertising, which is a more performance-based type of advertising that minimizes the cost per acquisition (i.e., an online conversion), are more effective towards spurring consumers' interest in the brand and the product, across all the different outcome variables. For instance, both increase the chances that the consumer will visit the advertiser's website after the exposure, by roughly more than 4%, as shown in Column 1 of Table 9. Our results demonstrate how effective this type of targeting is towards spurring consumers' interest in the product and brand. Qualitatively similar results are shown for the duration of the exposure of the consumers to the display advertisements in the Table A8 online appendix.

6. Additional Robustness Checks

We conducted several additional robustness tests with supplementary data and alternative model specifications to examine whether the key results remain consistent. In the following paragraphs, we describe in detail these additional robustness checks.

6.1. Falsification Checks

One might think that it is plausible that the previous set of models is simply picking up spurious advertising effects as a result of pure coincidence, a general increase in the corresponding metrics for our dependent variables, or other unobserved factors. To assess this possibility that the advertising exposure variable is capturing significant effects by chance or because of other confounding factors, we run various falsification tests using the same models (in order to maintain consistency) but randomly indicating which subjects have been treated (i.e., consumers who have been targeted by the advertising mechanism and the display advertisement was rendered visible) and not treated (i.e., consumers who have been targeted by the advertising mechanism but the display advertisement was *not* rendered *viewable*). In particular, we use a standard (pseudo) random number generator in order to create a dummy variable that indicates which of the targeted consumers are treated. Under this falsification test, since the treatment variable does not bear the real information of whether the display advertisement was viewable, the corresponding coefficient should not pick any effect in the falsification models and show that there is no impact.

The results of these falsification tests for the binary treatment are shown in Tables A9 of online appendix. We see that, under this check, the corresponding coefficients ‘ATE’ are not statistically significant, indicating that our findings are not a statistical artifact of our specification, but we indeed discovered the actual effects. Similar results were obtained for *all* the dependent variables in our study. Thus, the falsification tests show that the relationship between our dependent variables and advertising exposures did not arise spuriously but online display advertising significantly affects consumer behavior. The results of additional falsification tests for the duration of the exposure to the display advertisement are shown in Table A10; we see that coefficients are non-statistically significant across all the outcome variables.

6.2. DID Robustness to Alternative Specifications

Imbens and Wooldridge (2009) discuss an alternative specification for the DID model, with panel data that assumes unconfoundedness given lagged outcomes, in order to make treated and control individuals comparable on lagged outcomes. In order to check the sensitivity of our results to these different assumptions and specifications, we run our models with an additional covariate of the lagged outcome for each different consumer behavior outcome. As the results in Tables A12 and A13 in the online appendix indicate, our findings remain qualitatively the same. Thus, although these two models make fundamentally different assumptions, the ATE of display advertising still remains positive and significant.

Furthermore, in order to exploit the natural ordering of the advertising exposures, we test an alternative specification instead of using dummy (categorical) variables. In particular, we capture the non-linearity of the ATE of display advertising across the exposures, by using the inverse of the logarithm of the relative position of advertising exposure. As shown in Tables A14 and A15, this alternative specification corroborates our findings. Using such a specification, marketers can better understand the dynamic effects as a function of the relative position of the display advertising exposures in consumers' funnel paths.

6.3. Controlling for Demographics in the DID Matching

In order to also control for unobserved consumer demographic characteristics, which might drive the firm's targeting in the DID matching model, we enhance our data set by collecting demographics information at the zip code level. For each consumer identifier in our database, we map the associated IPv4 addresses of the customer into the zip code of the area in which the consumer is located.⁵ Then, we collect various demographics, such as the median household income, median age, male to female ratio, educational attainment, etc., at the zip code level.⁶ In order to control for the sensitivity of the DID matching estimator to unobserved demographics, we first predict the propensity score (i.e., the probability that a consumer will be successfully

⁵ Maxmind GeoIP2 Web Service was used for this mapping.

⁶ Source: U.S. Census Bureau, 2010 Census of Population.

treated with a display advertisement that is rendered viewable), introducing such demographic variables in the model in addition to consumer's advertising exposures and past browsing behavior. Then, we rerun the DID matching models and evaluate whether our results remain qualitatively the same. As shown in Table A16 of the online appendix, the effects of display advertising on spurring consumers' interest in the brand and the product remain qualitatively the same, after controlling for consumer demographics characteristics at the zip code level.

6.4. Controlling for seasonality and other time-related effects

In order to also control for seasonality trends and other time-related effects, we introduce in the employed individual-level panel data difference-in-differences model specification time dummies and time trends at the monthly, weekly, and daily level, respectively. Time dummies allow us to capture specific effects related to that month or week of the year that might affect the outcome and potentially bias the ATE. Alternatively, time trends allow us to model how the overall direction of the respective outcome moves over time and purge such effects from the coefficient of interest.

Tables A17 and A18 present the results when we introduce in our model time dummies and time trends at the month level, respectively. Similarly, Tables A19 and A20 present the results when we introduce time dummies and time trends at the weekly level, respectively. Furthermore, Table A21 presents the results when we introduce non-linear weekly time trends. Finally, Table A22 presents the results when we introduce weekly time dummies and daily time trends, at the same time. The results remain consistent under all specifications. Furthermore, our results also remain consistent when we control for the day of the week time dummies. Overall, we find that the estimate of interest remains highly robust when we control for average time-varying unobserved effects at the monthly, weekly, and daily level, which further enhances our confidence that the effect of a successful display exposure is not driven by time-varying trends or seasonality effects.

6.5. Controlling for Skewed Duration of Exposure

Furthermore, in our data set, the distribution of the duration of consumers' exposure to display advertisements is skewed right, with a skewness coefficient of 1.42. In order to control for extreme values of the durations that might drive the results of the ATE of display advertising for non-binary exposures, we rerun the non-binary treatment models after normalizing the duration of the exposure. Table A23 reports the results for this specification where it is shown that our coefficients remain qualitatively the same.

6.6. Additional Control Variables

Finally, various other robustness checks were employed. For instance, we controlled for whether the display advertisement has a dynamic content and, hence, whether the obtrusiveness of the ad affects the nature of our results (see Tables A24 and A25 in the online appendix). We also control for additional variables, including the actual position of the user in the funnel path (see Tables A26 and A27). As shown in the aforementioned tables, the results remain qualitatively the same.

7. Conclusions, Managerial Implications, and Limitations

In this study, we analyze the relationship between display advertising and several online consumer behaviors, including but not limited to the consumers' propensity to convert. We argue that the effectiveness of display advertising and any other advertising channel should be evaluated across a wide range of consumer behaviors, apart from increasing consumers' propensity to convert. Furthermore, we employ a natural experiment framework that allows us to evaluate advertising effectiveness as well as build causal attribution models with individual-level data. Finally, we move beyond the standard binary treatment of subjects that the prior literature typically adopts; to the best of our knowledge, this is the first to study and causally identify the effect of the duration of consumer's exposure to the display advertisement in a real-world setting.

7.1. Key findings

The results presented in this paper demonstrate that mere exposure to display advertising (i.e., without requiring the user to engage with the ad) spurs interest towards the advertiser's brand and product. We empirically show that consumers, as a result of display advertising, engage both in *active search*, exerting effort to gather information through search engines and direct visits to the advertiser's website, as well as in *passive search*, using information sources that arrive exogenously, such as future display ads. Interestingly, the longer the duration of an exposure to display advertising, the more likely the consumer is to engage in direct search behaviors (e.g., direct visits) rather than indirect ones (e.g., search engine inquiries). Furthermore, the effect of display advertising in passive search behavior is not negligible; consumers can be up to 1.47% more likely to click on a future display advertisement early in the funnel path. This study also provides deep insights into the dynamics of the effects of display advertising; we find that advertising effects are amplified up to four times when consumers are targeted earlier in the purchase funnel path, rather than later. Furthermore, display advertising also increases consumers' propensity to make a purchase.

Furthermore, our incremental lift approach demonstrates how advertisers can attribute credit across the various types of display advertising (e.g., prospecting, retargeting, affiliate targeting, pre-roll video, etc.) as well as how display advertising triggers other paid and non-paid advertising exposures in the funnel path. We argue that instead of attributing the whole credit for a conversion across all the channels, advertisers should attempt to identify the incremental lift that each channel has towards changing the consumers' behavior in desired ways. An incremental lift approach is appropriate as consumers almost never enter the funnel path in a vacuum since there exists a reservoir of past experiences, brand equities, and other factors, such as social influences, that will affect a consumer's decision to make a purchase.

7.2. Managerial Implications

The deployment of the proposed multi-channel attribution model itself from advertisers can also have significant managerial implications. In particular, implementing the proposed display channel attribution method, advertisers will be able to use big data analytics in order to substitute away marketing budget from less effective and cost efficient channels and media and enhance the overall effectiveness of their digital marketing strategy (Tucker 2012). At the same time, through the implementation of the proposed model, advertisers will be able to discern the most effective advertising channels and further capitalize on their targeting features to improve their return on investment. For instance, examining the various effects across the different types of display advertising targeting types in our dataset, it is evident that some targeting methods (e.g., retargeting advertising) perform better than others (e.g., branding pre-roll video advertising) towards increasing consumers' propensity to make a conversion. Conversely, other targeting methods are better suited for driving consumer awareness. In addition, the proposed channel attribution model can facilitate performance-based advertising by allowing advertisers to reward their digital display networks partners based on their true effectiveness towards driving key marketing objectives. Besides, the proposed channel attribution model can also facilitate real-time bidding while allowing the advertisers to perform real-time analysis of their advertising campaigns. Finally, the channel attribution model also enables better optimization of more granular advertising choices for advertisers, such as the effectiveness of display creative strategies or the identification of the optimal time intervals between advertising exposures. In particular, the proposed framework can be effectively used to identify subgroups of treated and untreated individuals that have been exposed to different tactical advertising strategies in order to identify the true causal effect of these advertising strategies towards achieving various key marketing objectives.

Our paper has limitations, which primarily arise from the data. Although we have data on all the search sessions that resulted either in a click on a paid advertisement, or in a click on the organic results for the specific advertiser, we do not have information about search sessions that led to

competitors' websites. However, this would only imply that our results could underestimate the effect of display advertising on spurring consumers' interest in the brand and the product, in which case we should interpret the corresponding estimates as a lower bound of the corresponding effects. Future work could also examine whether the effects we identify are persistent across various industries and business models and further enhance the external validity of the results. Additionally, in the current paper, we focus on examining consumers' behavior as captured by their actions at the next touchpoint of the funnel path. Future work could further examine long-term behavioral changes on outcomes other than the consumer' propensity to convert and evaluate the role of online advertising in the establishment of positive brand equity. Notwithstanding these limitations, we hope our paper paves the way for more research in this increasingly important and emerging stream of work in digital attribution.

References

- Abhishek, V., Fader, P., and Hosanagar, K. 2012. "The Long Road to Online Conversion: A Model of Multi-Channel Attribution," *Available at SSRN 2158421*).
- Athey, S., and Imbens, G.W. 2006. "Identification and Inference in Nonlinear Difference - in - Differences Models," *Econometrica* (74:2), pp. 431-497.
- Austin, P.C. 2011. "An Introduction to Propensity Score Methods for Reducing the Effects of Confounding in Observational Studies," *Multivariate behavioral research* (46:3), pp. 399-424.
- Baltagi, B. 2008. *Econometric Analysis of Panel Data*. John Wiley & Sons.
- Berman, R. 2013. "Beyond the Last Touch: Attribution in Online Advertising," *Available at SSRN 2384211*).
- Bronnenberg, B.J., Dubé, J.-P.H., and Gentzkow, M. 2012. "The Evolution of Brand Preferences: Evidence from Consumer Migration," *American Economic Review* (102:6), pp. 2472-2508.
- Dreze, X., and Hussherr, F.X. 2003. "Internet Advertising: Is Anybody Watching?," *Journal of interactive marketing* (17:4), pp. 8-23.
- Forman, C., Ghose, A., and Goldfarb, A. 2009. "Competition between Local and Electronic Markets: How the Benefit of Buying Online Depends on Where You Live," *Management Science* (55:1), pp. 47-57.
- Ghose, A.H., Sang Pil; Park, Sung-Hyuk. 2014. "Analyzing the Interdependence between Web and Mobile Advertising: A Randomized Field Experiment," *Working Paper*).
- Goldfarb, A., and Tucker, C. 2011. "Online Display Advertising: Targeting and Obtrusiveness," *Marketing Science* (30:3), pp. 389-404.

- Griliches, Z., and Hausman, J.A. 1986. "Errors in Variables in Panel Data," *Journal of econometrics* (31:1), pp. 93-118.
- Heckman, J., Ichimura, H., Smith, J., and Todd, P. 1998a. "Characterizing Selection Bias Using Experimental Data," National bureau of economic research.
- Heckman, J.J., Ichimura, H., and Todd, P. 1998b. "Matching as an Econometric Evaluation Estimator," *The Review of Economic Studies* (65:2), pp. 261-294.
- Heckman, J.J., Ichimura, H., and Todd, P.E. 1997. "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme," *The review of economic studies* (64:4), pp. 605-654.
- Hollis, N. 2005. "Ten Years of Learning on How Online Advertising Builds Brands," *Journal of advertising research* (45:02), pp. 255-268.
- Horf, R. 2014. "Digital Ad Fraud Is Improving - but Many Ads Still Aren't Seen by Real People."
- IAB, P.W.H.C. 2013. "Internet Advertising Revenue Report 2013 Full-Year Results." IAB.
- Imbens, G.W., and Wooldridge, J.M. 2009. "Recent Developments in the Econometrics of Program Evaluation," *Journal of Economic Literature* (47:1), pp. 5-86.
- Joo, M., Wilbur, K.C., Cowgill, B., and Zhu, Y. 2013. "Television Advertising and Online Search," *Management Science* (60:1), pp. 56-73.
- Kireyev, P., Pauwels, K., and Gupta, S. 2013. "Do Display Ads Influence Search? Attribution and Dynamics in Online Advertising," *Harvard Business School, Boston, MA*).
- Lewis, R., Rao, J.M., and Reiley, D.H. 2013. "Measuring the Effects of Advertising: The Digital Frontier," National Bureau of Economic Research.
- Lewis, R., and Reiley, D. 2010. "Does Retail Advertising Work: Measuring the Effects of Advertising on Sales Via a Controlled Experiment on Yahoo," Working paper.
- Li, H., and Kannan, P. 2013. "Attributing Conversions in a Multichannel Online Marketing Environment: An Empirical Model and a Field Experiment," *Journal of Marketing Research*).
- Manchanda, P., Dubé, J.-P., Goh, K.Y., and Chintagunta, P.K. 2006. "The Effect of Banner Advertising on Internet Purchasing," *Journal of Marketing Research* (43:1), pp. 98-108.
- Mood, C. 2010. "Logistic Regression: Why We Cannot Do What We Think We Can Do, and What We Can Do About It," *European Sociological Review* (26:1), pp. 67-82.
- Naik, P.A., and Raman, K. 2003. "Understanding the Impact of Synergy in Multimedia Communications," *Journal of Marketing Research* (40:4), pp. 375-388.
- Papadimitriou, P., Garcia-Molina, H., Krishnamurthy, P., Lewis, R.A., and Reiley, D.H. 2011. "Display Advertising Impact: Search Lift and Social Influence," *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining: ACM*, pp. 1019-1027.
- Rubin, D.B. 1997. "Estimating Causal Effects from Large Data Sets Using Propensity Scores," *Annals of internal medicine* (127:8_Part_2), pp. 757-763.
- Rutz, O.J., and Bucklin, R.E. 2012. "Does Banner Advertising Affect Browsing for Brands? Clickstream Choice Model Says Yes, for Some," *Quantitative Marketing and Economics* (10:2), pp. 231-257.
- Shao, X., and Li, L. 2011. "Data-Driven Multi-Touch Attribution Models," *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining: ACM*, pp. 258-264.

- Tucker, C. 2012. "The Implications of Improved Attribution and Measurability for Online Advertising Markets," *Competition in the Online Environment*).
- Tucker, C. 2014. "The Reach and Persuasiveness of Viral Video Ads," *Marketing Science* (Forthcoming (Articles in Advance)), September 29, 2014.
- Wooldridge, J.M. 2010. *Econometric Analysis of Cross Section and Panel Data*. MIT press.
- Xu, K., Chan, J., Ghose, A., and Han, S.P. 2015. "Battle of the Channels: The Impact of Tablets on Digital Commerce," *Working Paper*).
- Xu, L., Duan, J.A., and Whinston, A. 2014. "Path to Purchase: A Mutually Exciting Point Process Model for Online Advertising and Conversion," *Management Science* (60:6), pp. 1392-1412.
- Zantedeschi, D., Feit, E.M., and Bradlow, E.T. 2013. "Measuring Multi-Channel Advertising Effectiveness Using Consumer-Level Advertising Response Data." Citeseer.

Tables and Figures

Notes: Significance levels defined as * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; Robust standard errors are in parentheses.

Figure 1: Overlap of estimated propensity score matching probability distributions of treated and untreated individuals illustrates the good quality of matching

	Min	Median	Mean	Max
Propensity Score Treated (with demographics)	0.00049	0.47180	0.49940	0.99890
Propensity Score Non Treated (with demographics)	0.00001	0.39750	0.40050	0.99360
Propensity Score Treated (without demographics)	0.00044	0.47220	0.49950	0.99880
Propensity Score Non Treated (without demographics)	0.00001	0.39740	0.40040	0.99340

Table 1: Descriptive statistics for estimated propensity score for treated and non-treated individuals

Figure 2: Covariate balance summary for the difference-in-differences matching estimator

Covariate balance summary for matching estimator (1)				
	Standardized differences		Variance ratio	
	Raw	Matched	Raw	Matched
Lagged Frequency of Display Clicks	-0.0158	-0.0001	0.9110	0.9889
Lagged Frequency of Display Impressions	0.0095	0.0004	0.4504	0.9874
Lagged Frequency of Direct Website Visits	-0.0095	-0.0003	1.0434	1.0106
Lagged Frequency of Organic Searches	-0.0207	-0.0001	0.9442	0.9988
Lagged Frequency of Brand Paid Searches	-0.0253	-0.0010	1.0174	0.9931
Lagged Frequency of Brand Non-Paid Searches	-0.0209	-0.0004	0.8523	0.9695
Number of days since Last Purchase	-0.0106	0.0001	1.0379	1.0092
Lagged Ratio of Successful Viewable Treatments	0.0477	0.0007	1.3381	1.0277

Table 3: Assessing the quality of matching with covariate balance summary for the difference-in-differences matching estimator

	Search Brand	Direct Visit	Organic Search	Search Non-Brand	Display Click	Non-Immediate Conversion
Time	-0.0551***	-0.0547***	-0.0501***	-0.0418***	-0.0246***	0.0131***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0001)
Treated	-0.0020***	-0.0018***	-0.0011***	-0.0008***	-0.0010***	-0.0011***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)
Average Treatment Effect	0.0072***	0.0101***	0.0054***	0.0031***	0.0037***	0.0022***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0002)	(0.0001)
Constant	0.0550***	0.0552***	0.0502***	0.0418***	0.0251***	0.0243***
	(0.0005)	(0.0005)	(0.0005)	(0.0004)	(0.0003)	(0.0002)
Individual-level Heterogeneity	✓	✓	✓	✓	✓	✓
Ad-level Heterogeneity	✓	✓	✓	✓	✓	✓
R-squared	0.0474	0.0443	0.0436	0.0384	0.0192	0.0141
F-statistic	3094.3290	2880.0827	2835.8551	2484.2280	1215.6310	891.4947
N. of observ.	3000000	3000000	3000000	3000000	3000000	3000000

Table 4: ATE of Display Advertising Exposure

	Search Brand	Direct Visit	Organic Search	Search Non-Brand	Display Click	Non-Immediate Conversion
Time	-0.0551***	-0.0547***	-0.0501***	-0.0418***	-0.0246***	0.0131***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0001)
Treated	-0.0010***	-0.0011***	-0.0002	0.0001	-0.0006***	-0.0008***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)
ATE Exposure 1	0.0302***	0.0282***	0.0265***	0.0227***	0.0147***	0.0084***
	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0004)	(0.0003)
ATE Exposure 2	0.0193***	0.0167***	0.0156***	0.0136***	0.0092***	0.0045***
	(0.0006)	(0.0006)	(0.0005)	(0.0005)	(0.0004)	(0.0003)
ATE Exposure 3	0.0127***	0.0120***	0.0099***	0.0076***	0.0059***	0.0032***
	(0.0006)	(0.0006)	(0.0005)	(0.0005)	(0.0004)	(0.0003)
ATE Exposure 4	0.0071***	0.0095***	0.0061***	0.0041***	0.0038***	0.0023***
	(0.0006)	(0.0006)	(0.0005)	(0.0005)	(0.0004)	(0.0003)
ATE Exposure 5	0.0048***	0.0078***	0.0035***	0.0007	0.0025***	0.0015***

	(0.0006)	(0.0006)	(0.0005)	(0.0005)	(0.0004)	(0.0003)
ATE Other Exposures	-0.0008*	0.0052***	-0.0017***	-0.0039***	0.0001	0.0003*
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0002)	(0.0002)
Constant	0.0544***	0.0548***	0.0496***	0.0412***	0.0248***	0.0241***
	(0.0005)	(0.0005)	(0.0005)	(0.0004)	(0.0003)	(0.0002)
Individual-level Heterogeneity	✓	✓	✓	✓	✓	✓
Ad-level Heterogeneity	✓	✓	✓	✓	✓	✓
R-squared	0.0486	0.0448	0.0447	0.0396	0.0197	0.0144
F-statistic	2836.7881	2606.5979	2596.0580	2290.1442	1115.9656	812.7444
N. of observ.	3000000	3000000	3000000	3000000	3000000	3000000

Table 5: ATE of Display Advertising Exposure (Funnel Path)

	Search Brand	Direct Visit	Organic Search	Search Non-Brand	Display Click	Non-Immediate Conversion
Time	-0.0526***	-0.0506***	-0.0486***	-0.0411***	-0.0236***	0.0140***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0001)
Treated	0.0018***	0.0000	0.0016***	0.0010***	0.0005**	-0.0004**
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)
ATE Duration	-0.0001	0.0026***	0.0001	-0.0002	0.0003***	0.0003***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Constant	0.0542***	0.0520***	0.0498***	0.0418***	0.0247***	0.0244***
	(0.0005)	(0.0005)	(0.0005)	(0.0004)	(0.0003)	(0.0002)
Individual-level Heterogeneity	✓	✓	✓	✓	✓	✓
Ad-level Heterogeneity	✓	✓	✓	✓	✓	✓
R-squared	0.0480	0.0418	0.0444	0.0392	0.0195	0.0143
F-statistic	2937.7869	2541.6991	2702.6603	2376.6326	1157.6243	845.0611
N. of observ.	2820934	2820934	2820934	2820934	2820934	2820934

Table 6: ATE of Display Advertising Exposure (Exposure in mins)

	Search Brand	Direct Visit	Organic Search	Search Non-Brand	Display Click	Non-Immediate Conversion
Time	-0.0527***	-0.0507***	-0.0486***	-0.0411***	-0.0236***	0.0140***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0001)
Treated	0.0028***	0.0007**	0.0025***	0.0018***	0.0009***	-0.0002
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)
ATE Exposure 1	0.0085***	0.0091***	0.0076***	0.0070***	0.0043***	0.0019***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)
ATE Exposure 2	0.0043***	0.0054***	0.0037***	0.0036***	0.0022***	0.0010***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0001)
ATE Exposure 3	0.0018***	0.0036***	0.0018***	0.0014***	0.0012***	0.0006***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0001)
ATE Exposure 4	-0.0001	0.0027***	0.0003	0.0001	0.0004**	0.0004***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0001)
ATE Exposure 5	-0.0012***	0.0018***	-0.0006**	-0.0010***	-0.0002	0.0001
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0001)

ATE Other Exposures	-0.0029***	0.0007***	-0.0025***	-0.0026***	-0.0010***	-0.0002***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Constant	0.0536***	0.0516***	0.0493***	0.0413***	0.0244***	0.0243***
	(0.0005)	(0.0005)	(0.0005)	(0.0004)	(0.0003)	(0.0002)
Individual-level Heterogeneity	✓	✓	✓	✓	✓	✓
Ad-level Heterogeneity	✓	✓	✓	✓	✓	✓
R-squared	0.0492	0.0424	0.0453	0.0403	0.0200	0.0145
F-statistic	2690.2477	2305.6820	2469.6012	2184.9995	1062.6883	763.5120
N. of observ.	2820934	2820934	2820934	2820934	2820934	2820934

Table 7: ATE of Display Advertising Exposure (Exposure in mins)

	Search Brand	Direct Visits	Organic Search	Search Non-Brand	Display Clicks	Non-Immediate Conversion
ATE Expos. 1	0.0087***	0.0118***	0.0071***	0.0019***	0.0029***	0.0068***
	(0.0004)	(0.0004)	(0.0003)	(0.0003)	(0.0004)	(0.0013)
ATE Expos. 2	0.0067***	0.0084***	0.0055***	0.0031***	0.0033***	0.0012
	(0.0004)	(.0004)	(0.0004)	(0.0003)	(0.0003)	(0.0007)
ATE Expos. 3	0.0054***	0.0087***	0.0049***	0.0022***	0.0035***	0.0013**
	(0.0004)	(0.0005)	(0.0004)	(0.0003)	(0.0004)	(0.0007)
ATE Expos. 4	0.0028***	0.0051***	0.0036***	0.0010***	0.0023***	-0.0012
	(0.0003)	(0.0004)	(0.0003)	(0.0002)	(0.0003)	(0.0007)
ATE Expos. 5	0.0034***	0.0062***	0.0027***	0.0012***	0.0025***	-0.0011
	(0.0004)	(0.0005)	(0.0003)	(0.0003)	(0.0003)	(0.0006)
ATE Expos. 6	0.0042***	0.0066***	0.0029***	0.0016***	0.0025***	-0.0001
	(0.0004)	(0.0005)	(0.0004)	(0.0003)	(0.0004)	(0.0006)
ATE Expos. 7	0.0036***	0.0065***	0.0035***	0.0008**	0.0026***	0.0001
	(0.0004)	(0.0005)	(0.0004)	(0.0003)	(0.0004)	(0.0006)
ATE Expos. 8	0.0035***	0.0064***	0.0027***	0.0012***	0.0019***	0.0010
	(0.0004)	(0.0005)	(0.0004)	(0.0002)	(0.0004)	(0.0006)
ATE Expos. 9	0.0036***	0.0055***	0.0035***	0.0009**	0.0026***	0.0001
	(0.0005)	(0.0005)	(0.0004)	(0.0003)	(0.0004)	(0.0006)
ATE Expos. 10	0.0020***	0.0072***	0.0031***	0.0010***	0.0021***	0.0010
	(0.0004)	(0.0005)	(0.0004)	(0.0003)	(0.0004)	(0.0005)

Table 8: Effects of display advertising (DID Matching Estimator)

	Search Brand	Direct Visit	Organic Search	Search Non-Brand	Display Click	Non-Immediate Conversion
Treat	-0.0024	-0.0027	-0.0028	-0.0012	-0.0021	-0.0009
	(0.0043)	(0.0042)	(0.0041)	(0.0038)	(0.0030)	(0.0022)
Time	-0.0557***	-0.0521***	-0.0495***	-0.0428***	-0.0249***	0.0134***
	(0.0013)	(0.0013)	(0.0012)	(0.0011)	(0.0009)	(0.0007)
ATE	0.0073*	0.0109***	0.0063*	0.0039	0.0044*	0.0025*
	(0.0029)	(0.0028)	(0.0027)	(0.0025)	(0.0020)	(0.0012)
Constant	0.0557***	0.0523***	0.0496***	0.0425***	0.0251***	0.0243***
	(0.0017)	(0.0017)	(0.0016)	(0.0015)	(0.0012)	(0.0009)
Individual-level Heterogeneity	✓	✓	✓	✓	✓	✓

Ad-level Heterogeneity	✓	✓	✓	✓	✓	✓
Centered R2	0.0483	0.0421	0.0430	0.0392	0.0194	0.0143
F-statistic	1209.87	1045.78	1071.12	971.66	471.06	345.36
N. of observ.	1232922	1232922	1232922	1232922	1232922	1232922

Table 9: IV (2SLS) Fixed Effect estimation with precipitation level, temperature and relative humidity as instruments

	Search Brand	Direct Visit	Organic Search	Search Non-Brand	Display Click	Non-Immediate Conversion
Time	-0.0551*** (0.0002)	-0.0547*** (0.0002)	-0.0501*** (0.0002)	-0.0418*** (0.0002)	-0.0246*** (0.0001)	0.0131*** (0.0001)
Treated	0.0020*** (0.0002)	0.0018*** (0.0002)	0.0025*** (0.0002)	0.0023*** (0.0002)	0.0006*** (0.0002)	-0.0020*** (0.0001)
ATE Branding Pre-Roll Video	0.0273 (0.1281)	0.0273 (0.1287)	1.0241*** (0.1228)	0.0199 (0.1113)	0.0124 (0.0907)	-0.0050 (0.0649)
ATE Retargeting	-0.0183*** (0.0003)	-0.0128*** (0.0003)	-0.0177*** (0.0003)	-0.0170*** (0.0003)	-0.0067*** (0.0002)	0.0081*** (0.0002)
ATE Affiliate	0.0038*** (0.0006)	0.0070*** (0.0006)	0.0005 (0.0006)	-0.0004 (0.0005)	0.0004 (0.0004)	0.0040*** (0.0003)
ATE Seeding CPM	0.0437*** (0.0004)	0.0431*** (0.0004)	0.0389*** (0.0004)	0.0320*** (0.0003)	0.0191*** (0.0003)	-0.0065*** (0.0002)
ATE Targeted CPA	0.0414*** (0.0010)	0.0407*** (0.0010)	0.0380*** (0.0009)	0.0303*** (0.0008)	0.0184*** (0.0007)	-0.0061*** (0.0005)
Constant	0.0507*** (0.0005)	0.0514*** (0.0005)	0.0464*** (0.0005)	0.0385*** (0.0004)	0.0235*** (0.0003)	0.0252*** (0.0002)
Individual-level Heterogeneity	✓	✓	✓	✓	✓	✓
Ad-level Heterogeneity	✓	✓	✓	✓	✓	✓
R-squared	0.0561	0.0513	0.0516	0.0457	0.0223	0.0161
F-statistic	3373.6421	3070.2128	3089.0339	2716.9083	1295.0329	929.5277
N. of observ.	3000000	3000000	3000000	3000000	3000000	3000000

Table 10: ATE of Display Advertising Exposure by Type of Targeting (Exposure in mins)