What is the Cost of an Unseen Ad?

Steven Millman, comScore
ZhiWei Tan, comScore

Abstract

In order to determine the relationship between viewability and campaign lift, viewable and non-viewable ads - as defined by the Media Rating Council (MRC) - were measured independently using comScore's Smart Lift methodology. Four display campaigns were examined, among which there were nine metrics that displayed statistically significant lifts at the 90% confidence level. There were two main findings arising from this research. First, while it is hypothetically possible that the effect viewability could be non-linear for extremely large campaigns, in virtually all typical campaigns, the effect of viewability on lift will be linear in nature. Second, we found that non-viewable ads (those onscreen for less than a second and/or with less than 50% of the pixels on screen) were capable of contributing to the overall lift of a campaign. In all cases, viewable ads were substantially more effective than non-viewable ads, supporting both the definition of viewability (for display) and viewability measurement.

Introduction

Several years ago, an important step on the digital branding path was taken when three U.S. trade bodies representing the digital advertising ecosystem (IAB for publishers, ANA for advertisers, and 4As for agencies) announced an initiative named "*Making Measurement Make Sense*" (3MS). The focus of 3MS was to establish a new definition for digital ad impressions in which ad would only be counted only if it was "in-view" to the consumer. There are several reasons why an online ad might not be in-view; the most common being when an ad is loaded to a page by the ad server, but the viewer doesn't scroll down far enough to see the ad before navigating away. Another cause is invalid traffic (IVT) that fraudulently captures ad impressions through the use of bots or other techniques. The 3MS objective was to eliminate this low quality ad inventory and make digital directly comparable to TV, where "opportunity for the consumer to view" is an accepted tenet of brand advertising. Today, the MRC accredited definition for viewability requires at least 50% of pixels in view for at least one second for display advertising and at least 50% of pixels in view for at least two seconds for video ads. This definition is hoped to hold great promise for the future of the internet advertising industry, as it is based on simple, quantitative rules related to time onscreen and percent of pixels in view.

comScore has measured the viewability of billions of digital display ads from thousands of campaigns around the world. The results were startling to the industry; on average, only about 44% of ads were actually in view to the consumer. Obviously, this is an issue, as ads that are not in view have less chance of affecting consumer behavior. The viewability rate was higher on premium publisher sites (53%) than on the ad networks and exchanges (31%), with a substantial part of the difference being traceable to higher IVT on the exchanges.

There's substantial upside to be realized if viewability can be improved. comScore research has shown that increasing in-view rates can generate increases in ad impact. For example, Kellogg's realized a 75% increase in sales lift by increasing its viewability rates by 40%. Using tools and reporting systems such as validated Campaign Essentials (vCE) from comScore, viewability, exposure frequency, and the accuracy of demographic targeting can be monitored in real-time. Once campaign delivery measurements are delivered and shared with individual publishers, ad money can be shifted to those publishers and placements that will deliver the media plan as intended. The benefits of optimizing in this manner are substantial. Kellogg's has reported ROI increases by factors of five to six times from their investment in digital advertising since they and their agencies began optimizing digital campaigns in-flight.

As one can imagine, the use of viewability metrics has generated some advertiser demand for guaranteed digital audiences – such as those they are used to receiving in TV buys. While the jury is still out as to how publishers will ultimately respond, it's certainly clear that viewability is an issue whose time has come and one that promises success in attracting more brand advertising to the Internet.

IVT + Viewability

While IVT and viewability measurements are the result of two different methodologies focused on these two distinct challenges, they are not mutually exclusive and frequently collide. The most common problem arises when comparing single-point solutions where one accounts for both IVT and viewability, and the other just looks at viewability alone. Consider the example below of two viewability measurement providers, one of which applies IVT filters (Measurement Provider A) and one of which does not (Measurement Provider B).

	(A)	(B)	(C)	(D)	(E)	(F)
	Served impressions	Excluded as IVT/Fraud	Impressions post-IVT filtration	Disposition: Viewable	Disposition: Not Viewable	Reported Viewability
Measurement Provider A	13,000,000	4,000,000	9,000,000	7,200,000	1,800,000	55%
Measurement Provider B	13,000,000	0	13,000,000	10,400,000	2,600,000	80%

Measurement Provider A has identified 4 million impressions as fraudulent or another form of IVT; by definition these are treated as not viewable. Each classifies 80 percent of the "post-filtration" impressions as viewable. Since Measurement Provider A excluded 4 million impressions as IVT, their reported campaign viewability is 55 percent, or 7.2 million divided into the total impression count of 13 million. Measurement Provider B, filtering no traffic as IVT, reports a viewability of 80 percent, or 10.4 million divided into 13 million.

Such discrepancies among measurement providers can be disconcerting for clients and contribute to a perceived disparity – and lack of confidence – in the leading viewability technologies in the market today. The reason for this disparity is not, nor should it be, a mystery – it is being driven by the increasing incidence of IVT and varying abilities between measurement providers to filter it appropriately. comScore's viewability metric treats IVT as non-viewable, but since the research described in this essay involves survey respondents, IVT is not generally an issue for the research design.

So with this in mind, how can one determine exactly how much ROI is lost or gained when viewability rates go down or up, and what does that relationship look like? Is the effect linear or curved? It is these question that this research attempts to answer.

Viewable v. Non-Viewable Ads

To examine the nature of the relationship between viewability and lift metrics, we took advantage of comScore's Smart Lift™ methodology, currently under patent consideration, and used in our Brand Survey Lift™ product. Rather than looking for advertising lifts in response to survey questions using test and control methodology and segmentation, Smart Lift is a non-linear, regression-based approach that looks at each impression as an independent event with individual characteristics. These characteristics include, but are not limited to, creative, publisher, placement, strategy, and, of course, viewability. Since there is a distribution of advertising exposures across the survey sample, we can construct a curve that defines the declining marginal utility of the ad campaign. Unique to the market and similar in design to widely-accepted Market Mix Models, Smart Lift models a baseline response, which provides the estimate for the non-exposed, eliminating the need for an actual control group. This allows us to accurately estimate what the expected lift of a campaign would be for any average frequency of exposure.

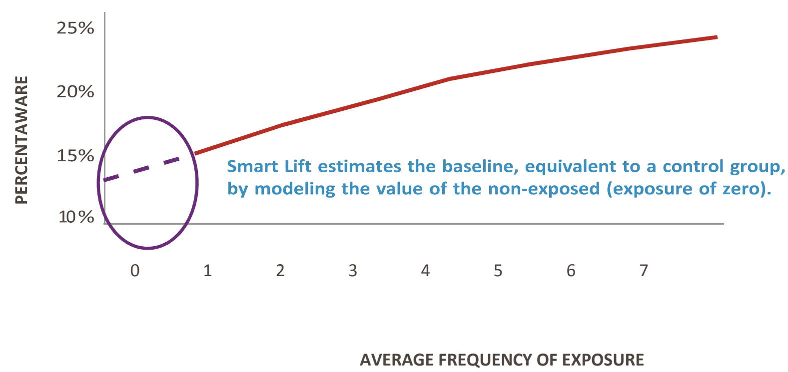

Because the analysis looks at each viewed ad or impression as an independent event, we can see separate curves emerge in attribution that control for cross-exposure between segment members. In the simple case that follows, you can see the ad effectiveness curves for two publishers, along with their lifts and baselines based on the average frequency of exposure for each publisher.

Effect of Impression Frequency on Likelihood of Top-2 Box Response by Publisher

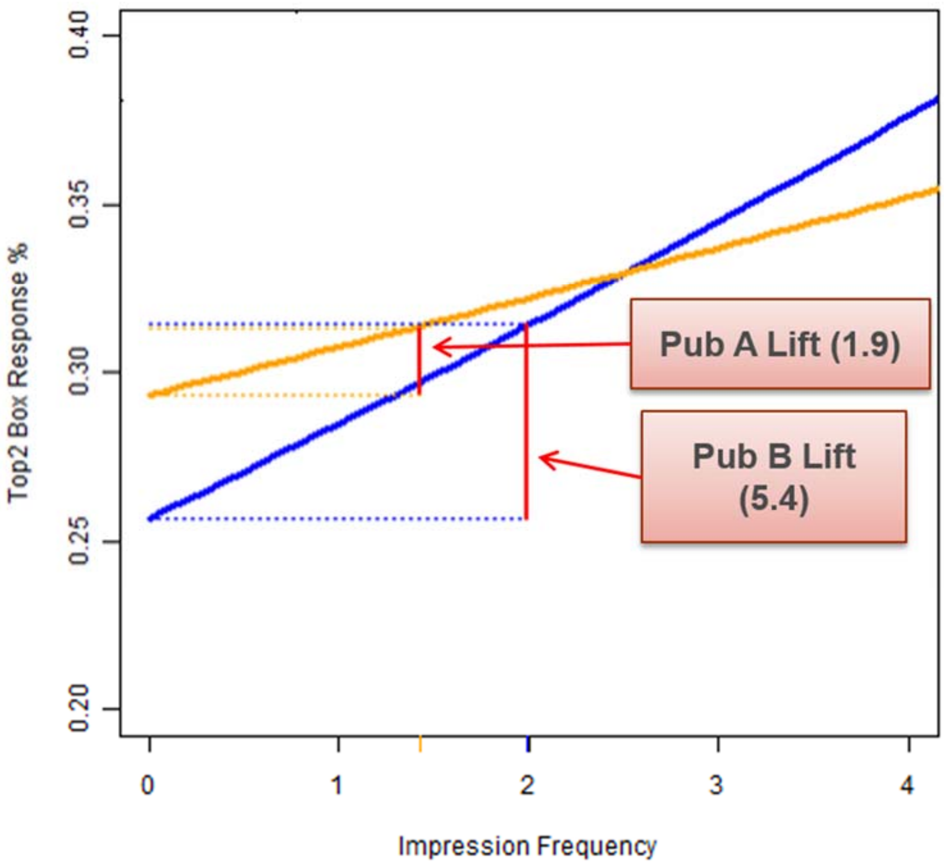

In order to measure the effectiveness of a campaign based on changes in viewability, we treated viewable and non-viewable ad exposures in much the same manner as we would if we were running an attribution analysis with two members of the attributive set such as the two publisher example above. The analysis was run this way on four campaigns across nine branding metrics in which a statistically significant lift at the 90% confidence level was observed. With these data, we examined what happened as the ratio of ineffective ads to effective ads was allowed to vary. Consider the following example:

Lift by Exposure for Viewable and Non-Viewable Ads
(High Average Frequency)

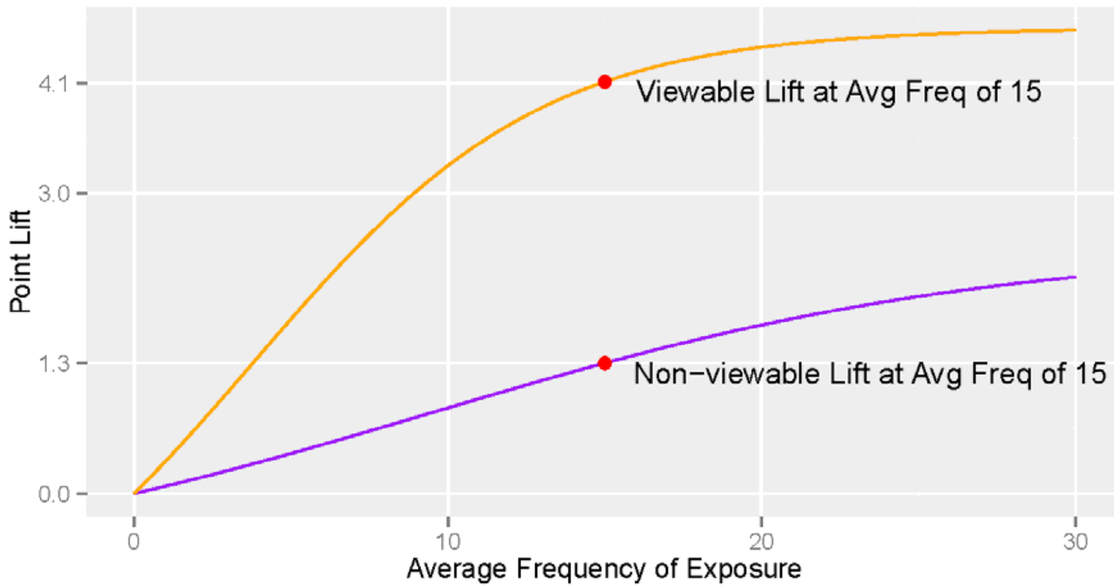

In this figure, the separate effectiveness curves for viewable and non-viewable ads are shown side by side. In this notional case, there is a relatively high average frequency of exposure in which the typical person was exposed 15 times. The red dot on the curves indicates the expected lift when 100% of ads are seen versus unseen. Both show the expected declining marginal utility curve. The following figure shows what happens to the overall lift measurement for the campaign as you vary the percent of ad impressions that are viewable:

Point Lift by Viewability at an Average Frequency of Exposure of 15

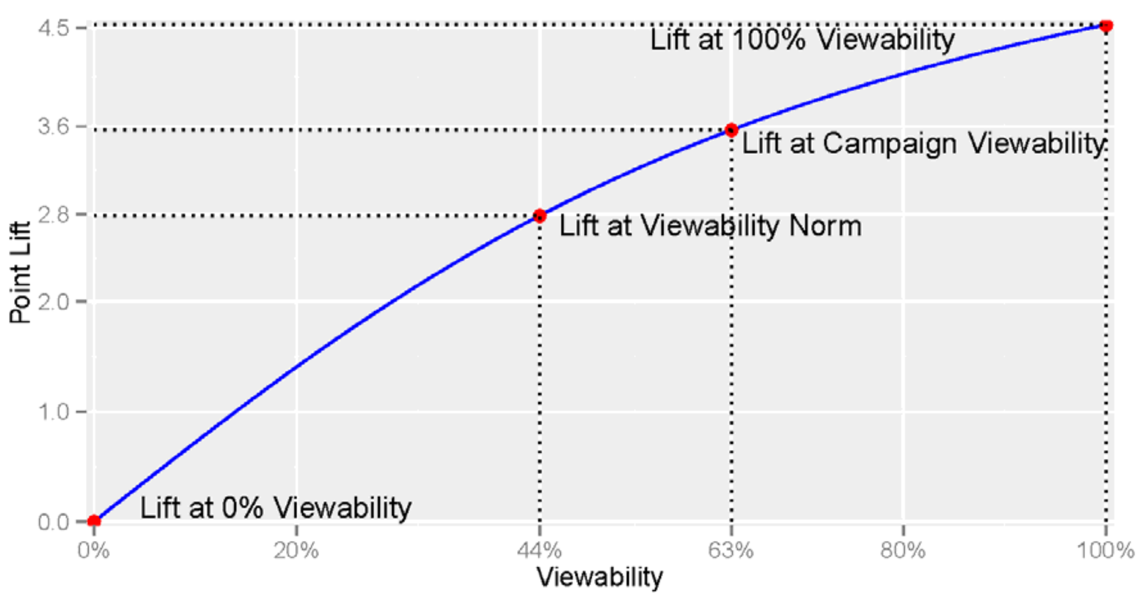

In this campaign, 63% of the ads served were viewable, a substantial improvement over the industry benchmark of 44%. As delivered, this campaign had a point lift of 3.6, but had the campaign performed only as well as the average campaign in terms of viewability (44%) they would have seen a diminished lift as the result of a smaller number of viewable ads, down to a 2.8 point lift. Additionally, if the difference between the effectiveness of viewable and non-viewable lift curves had been greater, this slope of the line would have also been steeper and the curve more pronounced. In this example, we're seeing a 29% increase in lift from a 43% increase in viewability. Due to the non-linear relationship expressed by the curve above, at lower viewability values the ROI improvement is more rapid than at higher values. This exposure rate of 15 impressions per exposed individual, however, does not describe a typical campaign. Typical campaigns generally have average frequencies of exposure that vary between two and seven per exposed individual. At exposure levels in this range, there is virtually no degradation in the marginal effectiveness of digital advertising, especially with display ads. As a consequence, the resulting viewable and non-viewable effectiveness curves look essentially linear in the range of observed behavior.

Lift by Exposure for Viewable and Non-Viewable Ads
(Low Average Frequency)

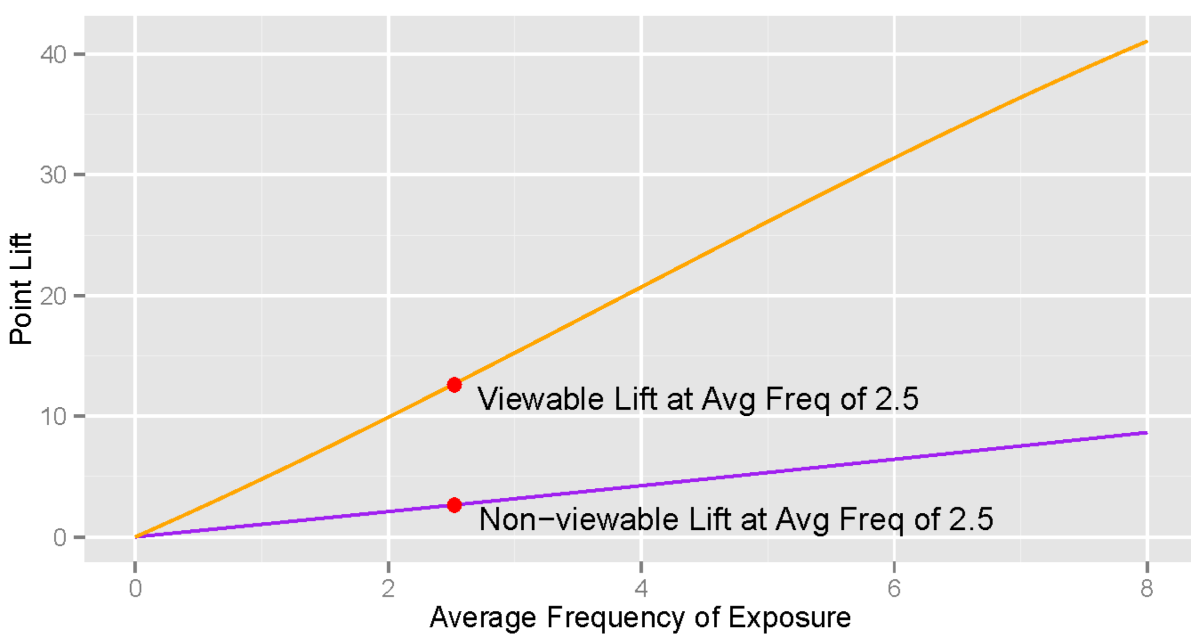

This is a real-world example of a highly effective digital campaign with a typical average frequency of about 2.5 impressions per exposed individual. As can be clearly seen in this exposure range, there is not yet any diminishing return and the effectiveness curves are essentially linear. The lift curve resulting from two linear contributors must also be linear, and therefore the viewability curve seen below is defined by a straight line.

Point Lift by Viewability at an Average Frequency of Exposure of 2.5

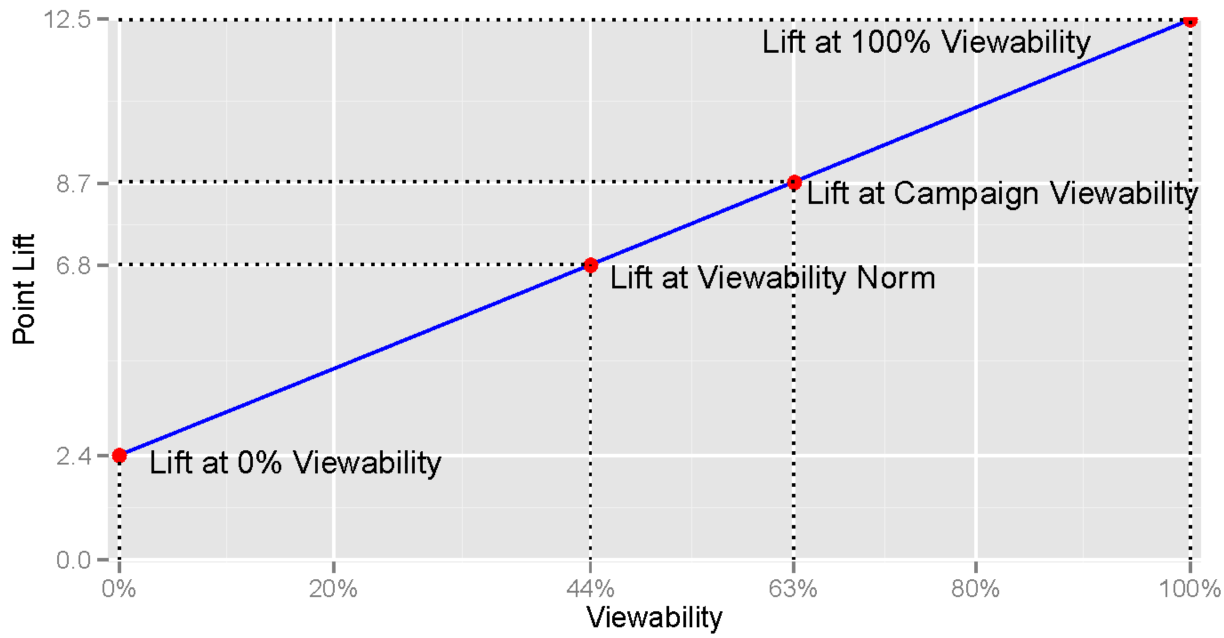

In every example studied for this research, the same result was observed. Because all of the campaigns had average frequencies of exposure in the expected range, no declining utility is observed. So while we have a rational expectation of a non-linear effect of viewability on lift based on the declining marginal nature of advertising impact, as a practical matter digital campaigns do not deliver at frequency rates high enough to reach the part of the distribution where the curvilinear form starts to take shape. The result is that the effect of viewability on lift appears to be nearly uniformly linear.

Are Non-Viewable Ads Worthless?

An additional finding in this research relates to the MRC definition of viewable impressions. The point of a viewability standard is to differentiate between ad exposures that do and do not have the opportunity to impact a consumer. In most cases this was true: in six of the nine cases studied where there was a statistically significant lift overall, non-viewable impressions had an effect that was statistically indistinguishable from zero. In the other three cases, there was a statistically significant impact from both the viewable and non-viewable ads, although the effect of the viewable ads was substantially higher. In the case above, had the campaign had a zero percent viewability, the lift observed would still have been 2.4 points. These results are strong evidence supporting both the MRC definition and vCE's measurement capability. But how are non-viewable ads effective? It is certainly possible that a simple visual ad could have a cognitive impact in less than a second, possibly even at the subconscious level, which is the basis for subliminal messaging. In a 2012 study, researchers Brooks, Savov, Allzén, Benedict, Fredriksson, and Schiöth found that subliminal messages can elicit reactions in the brain that are measureable on magnetic resonance imagers (MRI) even though the subjects were not consciously aware of their exposure to them (*Neuroimage*, Volume 59, Issue 3, 1 February 2012, Pages 2962–2973). With respect to video ads, the two second definition has an interesting consequence. Because video ads appear as a static image prior to the video rolling, there is essentially a display ad on the screen which could provide value to the campaign, and would not be measured as viewable if it failed to run.

While these may serve to strengthen the hypothesis that very short exposures could have some effect, a more likely explanation is that simple graphic-heavy ads do not require 50% of pixels on page in order to have an impact. Neither Mars nor Snickers campaigns were used in this analysis, but the Snickers internet ad below serves as an excellent hypothetical example of how an ad impression could be effective without meeting the MRC definition of in-view.

This Snickers brand is among the more universally aware in the market, and so it requires little cognitive effort to recognize it’s logo and trademark. When you cut the ad off one side, someone familiar with the brand would have no difficulty in recognizing the campaign.

In this version with less than 50% of pixels, the message that if you are hungry you might like a Snickers is still pretty clear. Even more effective would be cutting the ad in half lengthwise as though the ad had scrolled up just past the edge of the screen.

Taken together, it may be a reasonable expectation to assume that simple, graphic-heavy ads with simple messaging will likely provide at least some lift among brand metrics. Future research will seek to investigate these effects further.

Appendix – the Model

Let $Y_i = 1$ if person i has a positive response to a specific survey question and 0 otherwise. We assume Y_i is a random variable that comes from a Bernoulli distribution with probability p_i of occurring. This probability can be modelled according to the following,

$$p_i = \varphi(\alpha + \beta_D \vec{D}_i + \beta_V V_i + \beta_{NV} NV_i)$$

where α, β_D, β_V,and β_{NV} are all parameters to be estimated. $\vec{D}_i$ is a vector of demographic variables (age, gender, income) for person i; V_i (or NV_i) is the number of times person i has been exposed to an viewable (or non-viewable) ad prior to taking the survey; and φ is the inverse-logit function used as the link function in logistic regression.

Note that the gross number of times person i has been exposed to an ad prior to taking a survey is also tracked. Gross exposures consist of viewable and non-viewable exposures. For person i, $X_i = V_i + NV_i$, where X_i denotes gross exposures.

Estimation of Parameters

We fit a binomial logistic regression for each Y_i using $\vec{D}_i$, V_i, and NV_i as predictors. Given Y_i is a Bernoulli random variable with probability p_i and that Y_i are independent across Y_i, we may write a likelihood function for the entire set of respondents. Parameters are then estimated using maximum likelihood estimators (MLE).