Lift Measurement

A Data-Driven Approach with VideoAmp

Table of Contents

Introduction

"Half the money I spend on advertising is wasted; the trouble is I don't know which half."

– John Wanamaker

Advertisers spend billions of dollars each year to reach consumers. In 2021, advertisers in the US spent an estimated total of $278 billion, a 23% increase from the year prior.[1] And yet, the industry as a whole is still searching to uncover an accurate means of answering the fundamental questions that keep advertisers up at night such as, "Did my advertising drive consumers to purchase our product" or "Did we advertise in the right places to have an impact on conversions?" As we start to layer on the complexities of defining and measuring lift across channels, it becomes increasingly difficult to answer these questions. In order to empower advertisers to make smarter data-driven decisions, we must begin by unraveling the data to uncover how an ad impacts a key business metric.

The purpose of ads is to influence individuals to perform a specific action, whether purchasing a product, signing up for a service, tuning into a program, downloading a mobile app or visiting a website. This is known as a **conversion**.

By disentangling the effect of an ad from an individual's inherent predisposition to perform an action, we uncover what is known as **lift**. When measured accurately, lift provides the most meaningful metric by which to optimize an ad campaign.

In the most basic of terms, lift refers to the number of individuals who converted due to their exposure to an ad campaign. Or, how many more individuals converted after seeing an ad who otherwise would not have. This may seem like a fairly basic question, but in actuality, it is extremely complex and requires sophisticated tools to answer with accuracy.

It is important to consider the inherent challenges of measuring lift across channels. The gold standard approach to measuring lift within digital environments is utilizing experiments, or randomly assigning individuals to a test or control group before the campaign runs. While this approach is effective for digital campaign measurement, implementing a similar approach for linear TV is simply not feasible since it is primarily purchased as a one-to-many medium, meaning that ads are not targeted to individual households. Addressable TV is the exception to that rule; however, in 2021, addressable TV campaigns only represented a small portion of total TV advertising spend at 4.3%.[2] Therefore, it is critical to have an approach to measuring lift that works for the vast majority of linear TV spend and adapts to the nuances of how linear TV, streaming and digital campaigns are executed.

In this whitepaper, we will demonstrate the effectiveness of VideoAmp's data-driven approach to measuring lift across linear TV, streaming and digital channels.

1 Source: eMarketer, "US Total Media Ad Spending and Growth, by Format/Media, H1 2021 & 2020-2022". 27 September 2021. LINK
2 Source: eMarketer, "Linear addressable TV ad spending will grow 33.1% this year". 11 May 2021. LINK

Measuring Lift Across Channels

A common industry-wide attempt at determining the effectiveness of cross-channel advertising is known as **naive lift**, or comparing the conversion rate of an exposed group to an unexposed group of people. The difference in these two conversion rates, multiplied by the number of people exposed, equals the naive lift of the campaign.

It is important to consider that all individuals are, in fact, people who possess unique behavioral and demographic attributes. There is no guarantee that two individuals would behave the same way after seeing an ad. We see this issue occur often with linear TV advertising. Naturally, those who are most likely to be exposed to linear ads are those who watch the most TV. Illustrated in Figure 1.1, VideoAmp conducted a study to compare the TV viewership habits for an exposed and unexposed group over a two month period. As shown below, the exposed group contains a larger portion of individuals who are heavy TV watchers. The median household's linear TV consumption within the exposed group was 46 hours per week, compared to only 13 hours per week for the unexposed group.

If watching TV correlates to the predisposition to convert, and high volume TV watchers are more likely to be in the exposed group, the naive method will overestimate the lift. For example, if the ads are for an upcoming TV show premiere, it's conceivable that TV watchers are more likely to tune in. This leads to those TV watchers seeing the ads at a higher rate. A naive comparison between exposed and unexposed groups in this example would not account for the exposed group containing more likely TV watchers.

Human behaviors are diverse and ever-evolving, and consumers are not aware of which ad exposure, or combination of ad exposures, resulted in them converting. Therefore, it is important to have an approach to lift that takes into account each person's unique attributes that may impact their likelihood of converting, while being channel-agnostic. In order to accurately measure lift across linear TV and digital, VideoAmp has developed a new approach that puts these channels on the same playing field.

FIGURE 1.1

TV Viewership Habits for an Exposed and Unexposed Group Over a Two Month Period

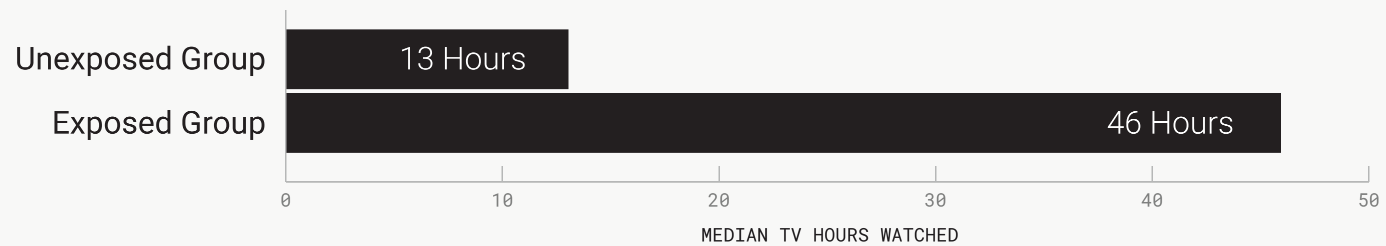

VideoAmp's Data-Driven Lift Approach:

Lifting Measurement Standards to New Heights

VideoAmp's data-driven approach to measuring lift across channels has two key components: probability modeling and assigning lift credit to each touchpoint.

Probability Modeling

VideoAmp first identifies all households that were exposed to an ad within VideoAmp's TV Viewership footprint and determines their probability of converting based on a combination of household and touchpoint attributes.

- **Household Attributes:** Age, gender, income, presence of children, historical TV viewing and past purchase behavior.
- **Touchpoint Attributes**: Frequency, recency, website and network.

VideoAmp then estimates the household's probability of converting when they are unexposed to an ad based on the same attributes. The difference in conversion probabilities represents each household's lift for the campaign; the sum of each household's lift represents the **total lift** of the campaign. This approach removes potential confounding variables that would impact a household's likelihood of converting in order to determine whether ad exposure resulted in a conversion.

As an example, let's consider an ad campaign for an upcoming episode of a TV show. It is important to account for past viewing behaviors when modeling the conversion probability since individuals who watch similar shows will have a higher predisposition to tune in. If we assume this person was not exposed to the ad campaign, we can use the model we created to infer how likely they are to convert in the absence of ads. The difference between the probabilities of them tuning into the episode when they are exposed to the ad, versus when they are not exposed to the ad, is the lift assigned to that individual. The sum of each individual's lift is the total lift of the campaign.

Assigning Lift Credit to Each Touchpoint

In order to assign lift credit to each touchpoint across the campaign, we leverage **Shapley Values**, a solution concept in cooperative game theory that finds each player's marginal contribution averaged over every possible sequence in which the players could have been added to the group. When applied to advertising, this model considers all possible conversion paths to determine each touchpoint's contribution to the total lift.

"Every interaction creates a data point, and every data point tells a piece of the customer's story."

– Paul Roetzer

Putting Our Data to the Test

Since there is no truth set to compare lift methodologies to, we can never be truly certain of what a person would do if they had not been exposed to an ad. However, we can use simulated data to investigate the effectiveness and accuracy of different lift methodologies.

To compare the naive and data-driven lift methodologies, VideoAmp simulated a tune-in campaign for an upcoming episode of a TV show where ads for the episode could run on three linear TV networks. This was a blinded study, so we will refer to each network as Network A, Network B and Network C, with the new episode set to premiere on Network A. To keep it simple, the only differentiating attributes that the individuals within the simulation had were how much they watched each TV network. All other attributes for these individuals were identical.

First, various inputs were placed into the simulation, including the following:

- **Ad Impact:** The strength of an ad's influence when seen on each network. A network with a high ad impact indicates that the ads seen on that network increase the viewer's likelihood of tuning in more than a network with a low ad impact.
- **Watch Impact:** The correlation between the amount and type of content the viewer has watched on a network and their likelihood of tuning into the advertised episode. A network with a high watch impact indicates that people who watch that network have a higher propensity to tune into the advertised episode, while a network with a low watch impact indicates that watching that network typically does not correlate to a higher propensity to tune in. For example, if the advertised episode is on a network that only airs sporting events, other networks that air sporting events will have a high watch impact since the content of the two networks are similar. However, a network that only airs children's programming might have a lower watch impact since that is a very different type of content.
- **Number of Ads**: The number of ads each network runs.

The values of these inputs for each of the three networks were placed into the simulation to calculate the expected lift for each network, as shown in Figure 2.1.

FIGURE 2.1
VideoAmp Simulated Tune-In Campaign Inputs

	Number of Ads	Ad Impact	Watch Impact
Network A	20	0.1	0.3
Network B	100	0.005	0.1
Network C	40	0.4	0.001

Figure 2.2 illustrates the results from 100 simulations. This plot contains the expected lift calculated using the simulation inputs, along with the measured lift for each touchpoint using VideoAmp's data-driven lift method. For comparison, the plot also contains the measured lift using two different naive methodologies:

- **Full Naive Method**: Requires the unexposed group to have not seen an ad across any touchpoint.
- **Siloed Naive Method:** Requires the unexposed group to have not seen an ad across a single touchpoint at a time.

The placement of the points in relation to the dashed line represents alignment between measured and expected lift. The closer the points are to the dashed line, the better the performance.

The results of the simulation showcase the accuracy of VideoAmp's data-driven lift method across all touchpoints. VideoAmp's data-driven lift method performed better than both naive methods across Networks A and B. For Network C, VideoAmp's performance is on par with the siloed naive method, but far exceeds the full naive method. While results for Network A show VideoAmp's measured lift is 15% higher than the expected lift, the difference is much more significant for the siloed naive method, which is 177% higher, and the full naive method, which is 578% higher than the expected lift.

In summary, touchpoint attributes, such as network watch time, which correlate with both the ad exposure and the conversion event, are not handled well by any naive method of lift measurement. However, VideoAmp's data-driven lift method handles this scenario exceptionally well, which is important as these attributes are seen very often in the real world, especially when measuring tune-in.

FIGURE 2.2

VideoAmp Simulated Tune-In Campaign Results

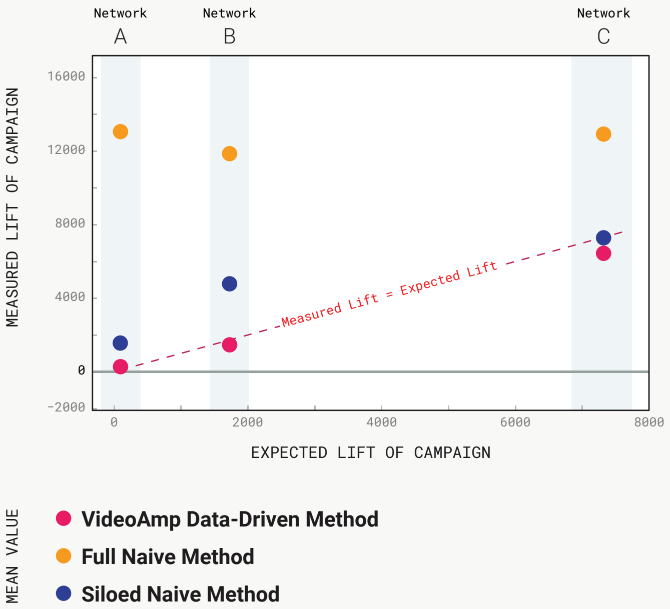

Data-Driven Lift: For Online & Offline Conversions

While the previous section centered around validating VideoAmp's data-driven lift methodology using a tune-in campaign simulation, this section applies VideoAmp's methodology to online and offline conversion campaigns. There are subtle, albeit critical, differences between tune-in campaigns and online or offline conversion campaigns which require an additional step in our lift methodology: conversion path treatment.

Basic Path Treatment

The industry standard for conversion path treatment is known as the **basic path treatment**, illustrated in Figure 3.1. If an individual converts, their path will contain all impressions delivered to them prior to the conversion. If the individual converts multiple times, they will have a separate path for each conversion. If an individual does not convert, their path will simply be all impressions delivered to them. Modeling conversions using a dataset constructed from a basic path treatment is not able to accurately determine conversion probabilities, and can therefore lead to less accurate lift measurement. VideoAmp has developed an advanced handling of the conversion path to help mitigate this.

Advanced Path Treatment

VideoAmp's approach to conversion path treatment is referred to as the **advanced path treatment**, and estimates the probability that the individual converts after each impression. VideoAmp's simulations have demonstrated that this treatment largely increases the accuracy of lift measurement compared to the basic path treatment.

FIGURE 3.1
Basic Path Treatment

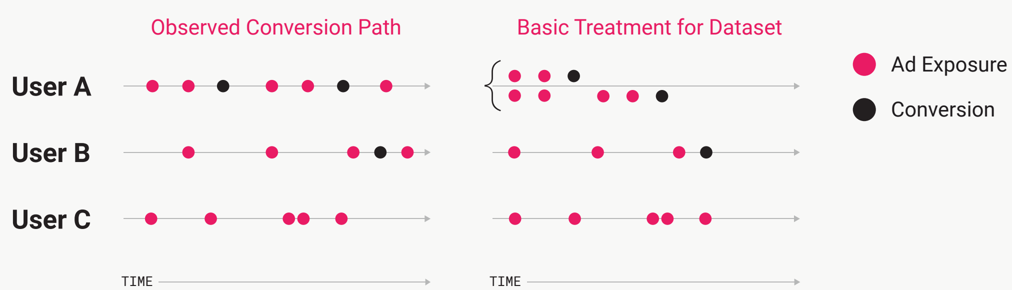

Path Treatment Comparision

Once again, VideoAmp performed a simulation, this time emulating a campaign associated with online conversions. While the inputs remain unchanged from the previous simulation, there was a notable difference in values, and multiple conversions were allowed.

Lift was measured using VideoAmp's data-driven method, with and without the advanced path treatment. As explained previously, the expected lift is calculated using the inputs of the simulation. In Figure 3.2, the placement of the points in relation to the dashed line represents alignment between measured and expected lift. The closer the points are to the dashed line, the better the performance.

Simply changing the conversion path treatment can produce vastly different results. VideoAmp's data-driven lift method, when combined with the basic path treatment, underestimates the lift to such a degree that it becomes negative. This is seen frequently with the basic path treatment. However, VideoAmp's data-driven lift method combined with the advanced path treatment performs very well, with strong alignment between the expected lift and measured lift of the campaign.

Now, let's consider competing methods for measuring the lift of an online conversion campaign. Figure 3.3 incorporates the siloed naive and full naive methods using basic path treatments, two competing methods with the basic path treatment, as well as VideoAmp's implementation of the logistic regression method with the advanced path treatment.

The basic path treatment underestimates lift on all accounts. The two competing methods using the basic path treatment fail in the same way VideoAmp's method with the basic path treatment does. The basic path treatment underestimates the lift to such a degree that it becomes negative. The naive methods either overestimate or underestimate the lift, while VideoAmp's method is the most accurate across all touchpoints.

FIGURE 3.2

VideoAmp Simulated Online Campaign Results Across Path Treatments

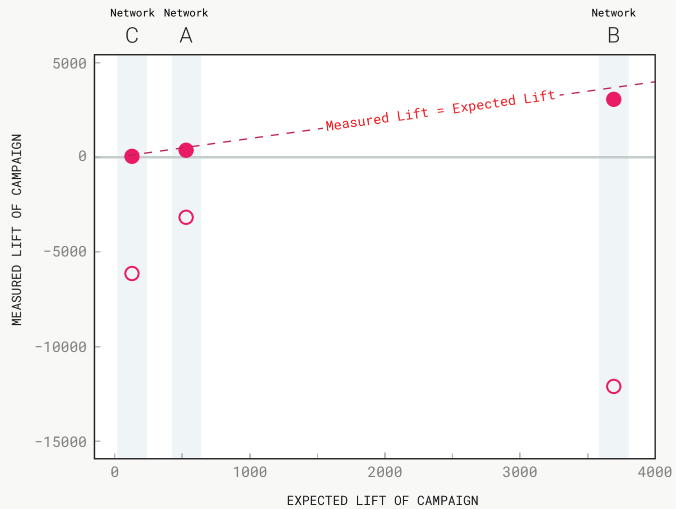

MEAN VALUE

- ○ **VideoAmp Data-Driven Method** with **Basic** Path Treatment
- ● **VideoAmp Data-Driven Method** with **Advanced** Path Treatment

FIGURE 3.3

VideoAmp Simulated Online Campaign Results Across Competing Methods

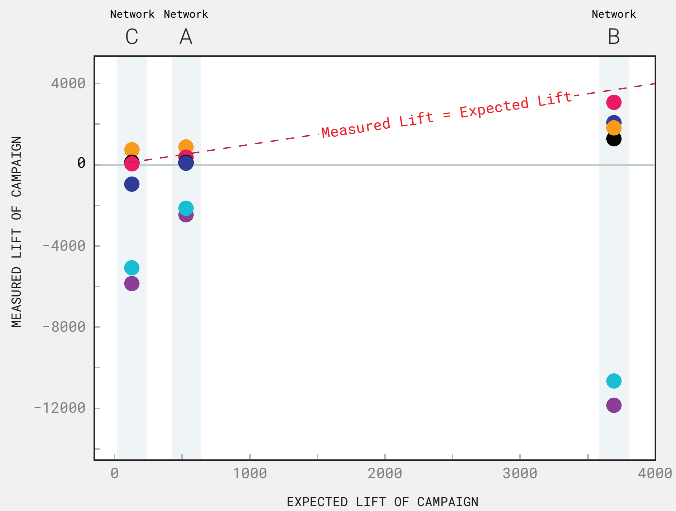

MEAN VALUE

- **VideoAmp Data-Driven Method** with **Advanced** Path Treatment
- **Competing Method 1** with **Basic** Path Treatment
- **Competing Method 2** with **Basic** Path Treatment
- **Logistic Regression** with **Advanced** Path Treatment
- **Siloed Naive Method** with **Basic** Path Treatment
- **Full Naive Method** with **Basic** Path Treatment

Figure 3.4 illustrates the results for another simulation that is conducted, which uses the advanced path treatment for all methodologies, omitting the naive methods. Even when we apply our advanced path treatment to the competing methods, VideoAmp's data-driven lift method continues to outperform. The most difficult touchpoint to measure is Network B, where even VideoAmp's measurement underestimates lift. However, VideoAmp's method still delivers a measured lift that is much closer to the expected lift for that touchpoint.

FIGURE 3.4

VideoAmp Simulated Online Campaign Results Across Competing Methods with Advanced Path Treatment

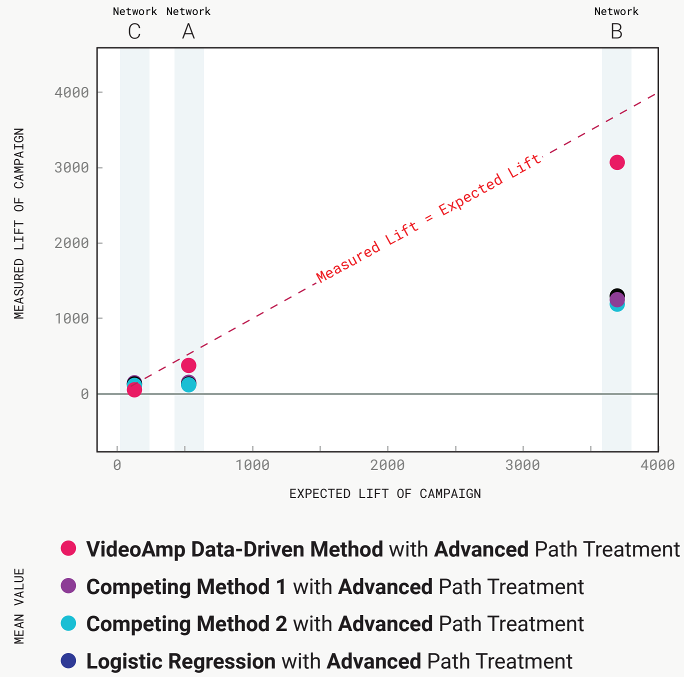

Closing the Gap: The Future is Data-Driven

In this whitepaper, we have introduced the meaning and value of lift, along with VideoAmp's data-driven lift measurement methodology. Using simulations, VideoAmp has been able to validate the accuracy of our data-driven lift method, which includes probability modeling, assignment of lift credit to each touchpoint, and conversion path treatment for online and offline conversion campaigns.

While it is clear that lift is a very difficult metric to measure, VideoAmp is continually focused on innovating, testing, and validating new methods to challenge the industry standard and bring the most accurate cross-channel solution to market. These advancements in lift measurement bring us all closer to understanding the effectiveness of advertising, and empower advertisers to make smarter, data-driven decisions.

Glossary

Ad Impact	The strength of an ad's influence when seen on each network. A network with a high ad impact means that ads seen on that network increase the viewer's likelihood of tuning in more than a network with a low ad impact.
Advanced Path Treatment	VideoAmp's approach to conversion path treatment that estimates the probability that the individual converts after each impression they were served.
Basic Path Treatment	An approach to conversion path treatment where an individual's path will contain all impressions delivered to them prior to the conversion. If the individual converts multiple times, they will have a separate path for each conversion. If an individual does not convert, their path will simply be all impressions delivered to them.
Conversion	The completion of an advertiser's desired action by an individual, such as purchasing a product online, visiting a website, submitting a form, tuning into a TV show, etc.
Full Naive Method	An approach to naive lift where the unexposed group of individuals are unexposed across all touchpoints.
Lift	A numerical calculation that represents an increase in a desired business outcome (i.e. sales), achieved by advertising. Measuring lift helps advertisers identify and understand the impact of ad exposure on conversions, and how to optimize their marketing efforts.
Logistic Regression	Statistical model used to predict the probability of a binary outcome (one where there are only two possible scenarios—either the event happens or it does not happen) based on a set of independent variables.
Naive Lift	Method that compares the conversion rate of an exposed group to an unexposed group of people. The difference in these two conversion rates, multiplied by the number of people exposed, equals the naive lift of the campaign.
Number of Ads	The number of ads each network runs.
Shapley Values	A solution concept in cooperative game theory that finds each player's marginal contribution, averaged over every possible sequence in which the players could have been added to the group. When applied to advertising, Shapley Values determines the marginal contribution of each ad exposure in contributing to conversions.
Siloed Naive Method	An approach to naive lift where the unexposed group of individuals are unexposed for a single touchpoint at a time.
Total Lift	The sum of each household's lift for the campaign.
Touchpoint	An individual network or platform by which a household was exposed to an ad.
Watch Impact	The correlation between the amount and type of content the viewer has watched on a network and their likelihood of tuning into the advertised episode. A network with a high watch impact indicates that people who watch that network have a higher propensity to tune into the advertised episode, while a network with a low watch impact indicates that watching that network typically does not correlate to a higher propensity to tune in.

About VideoAmp

VideoAmp is a software and data company creating a more sophisticated data driven advertising ecosystem that redefines how media is valued, bought and sold, creating a source of truth and standard for advertisers and media sellers.

The VideoAmp platform provides measurement and optimization tools that unify audiences across the disparate systems of traditional TV, streaming video and digital media. Unlocking new value for those currently operating within a siloed view of their audiences, VideoAmp creates efficiencies for the entire industry.

VideoAmp is transforming a 100-year old industry by powering a more effective three-way value exchange that results in advertisers increasing their return on investment, publishers increasing their revenues and improving the viewing experience for consumers.

VideoAmp is headquartered in Los Angeles with offices across the United States.

To learn more, visit videoamp.com.