Optimal Allocation of Programmatic Video Advertising Inventory

Avnish Brar
San Francisco State University, San Francisco, CA, USA

Olga Pipko
San Francisco State University, San Francisco, CA, USA

Robert Saltzman*
San Francisco State University, San Francisco, CA, USA

This article examines how a campaign manager can best allocate the inventory of programmatic video advertising on the Internet. Publishers, such as cnn.com, have different inventory spots that they make available via different purchasing methods, such as Open Exchange, Private Marketplace and YouTube, which have different costs and performance metrics. The goal of our linear programming model is to identify the optimal mix of inventory that satisfies thresholds for several campaign key performance indicators. In particular, we find the inventory mix that minimizes the cost per completed view subject to constraints on overall viewability and completion rates. The optimal allocation based on actual 2016 campaign data could have reduced the cost per completed view by more than 7%, potentially reducing actual total campaign costs by more than $1.5 million. We also present sensitivity analyses that could be useful for managers trying to set the parameters of their campaigns.

* Corresponding Author. E-mail address: saltzman@sfsu.edu

I. INTRODUCTION

Programmatic advertising is the algorithmic purchase and sale of advertising space in real time. During this process, software is used to automate the buying, placement, and optimization of media inventory via a bidding system. Different publishers who integrate ads into their online content have different amounts of advertising space (advertising inventory) available to sell to advertisers. Advertisers, in turn, can buy video advertising inventory via several different purchasing methods, as summarized in Table 1 below. For video ads, an impression is counted each time a video ad is placed on a Web page. Later, we'll distinguish between an impression, a *measurable* impression (an impression that a third party ad trafficking service can verify was placed on a Web page), and a *viewable* impression (a measurable impression that ran for two consecutive seconds in view of the user, as opposed to a separate part of the page that the user has not scrolled down to). Other relevant advertising terms are defined in the Appendix. For additional discussion of key concepts and components of programmatic advertising, see Busch (2016).

TABLE 1. VIDEO ADVERTISING INVENTORY PURCHASING METHODS.

Inventory Type	Purchasing Method	Price and Volume of Impressions
Open Exchange	Anyone can participate	Public auction with variable price and volume
PMP-Automated Guaranteed	Between one buyer and one publisher	Fixed price and fixed volume
PMP-Private	Between select buyers and one specific publisher	Private auction with a floor price
PMP-Preferred	Between several buyers and one publisher	Fixed price and variable volume
YouTube TrueView	Auction within YouTube among multiple potential buyers	Public auction within YouTube with variable price and volume

The different inventory purchasing methods vary in terms of their price and performance. The least costly purchasing method is through the Open Exchange, which is a real-time public auction. Here, advertisers bid on advertising spots (that don't include YouTube inventory) without knowing exactly what they're bidding for. The auctions take less than a millisecond to complete, with the advertising spot going to the highest bidder. One problem for advertisers with this purchasing method is that, because they don't know precisely what ad spot they're buying, they don't want to spend a big portion of their digital advertising budget on this type of inventory. Publishers, on the other hand, don't know precisely which advertisers are buying their advertising space and don't want to give their front-page ad spots to just any advertiser.

The Private Marketplace (PMP) solves the problem of the unknown for advertisers. Within the PMP, there are three different purchasing methods: Automated Guaranteed, Private, and Preferred deals. Automated Guaranteed offers a one-on-one deal negotiation directly between a publisher and a buyer with both inventory and pricing being fixed or guaranteed. Private, or Private Auction, occurs when one specific publisher holds an ad exchange auction with a group of select buyers. A Preferred deal is when buyers get a "first look" at the inventory that a seller offers for a pre-defined deal. Buyers can get the inventory by bidding at or above the negotiated fixed price; if no buyer chooses to buy the inventory, it will then potentially be available on the Open Exchange.

YouTube inventory is purchased through a separate method, called YouTube TrueView, which involves a public auction within YouTube videos among multiple potential buyers. YouTube TrueView (one of several YouTube purchasing methods) is transacted on a Cost Per View (CPV) basis, in which the buyer is only charged if the video is viewed for 30 seconds, or the user interacts with it. Buyers set the highest CPV they're willing to pay, or the maximum bid for the ad spots. Since YouTube is purchased with a CPV model, while the other inventory types are purchased with a cost per thousand impressions (CPM) model, we used cost per completed view (CPCV) as the common currency, where CPCV is the total media cost divided by the number of completed views. YouTube is a major part of many consumers' lives, so it is essential to include at least some YouTube inventory in a media budget even though its CPCV is considerably higher than that of the other types of advertising inventory.

Different inventory types have different performance and costs associated with them. To run a successful video ad campaign, managers strive to utilize their entire campaign budget in the most effective way. Most video campaigns have goals, or key performance indicators (KPIs), with minimum thresholds. When campaign managers set up a campaign, they want to use their entire media budget, while achieving the campaign's minimum KPI thresholds. In doing so, a campaign manager may try to maximize the number of impressions, thus increasing the advertisement's reach, while meeting campaign KPIs. However, achieving more than the threshold KPI is usually undesirable since clients will often then increase the minimum threshold in future campaigns.

Video campaigns typically have two primary KPIs: viewability rate and completion rate, where the former is the percentage of ads actually seen by a user, and the latter the percentage of ads that are run to completion, regardless of whether a user might have seen them. Viewability rates are measured by a tracking code that is often not available on mobile advertisements. Completion rates are easier to track across different devices. Our objective is to minimize the CPCV. For example, if a $100 advertising campaign placed 1,000 impressions (video ads), and 800 of these ran to completion, then the CPCV would be $100/800 = $0.125. Thus, the more ads run to completion for a given total cost, the lower the CPCV. For a specific video campaign, managers use a blend of different inventory types to achieve their goals. To create this mix, they typically use the same inventory blend as in prior campaigns, or guess to the best of their knowledge. The problem with this approach is that the resulting inventory mix is usually not as efficient as it could be.

Our model allows campaign managers to input their campaign KPI thresholds to generate an optimal inventory mix that minimizes the overall CPCV and includes ad space purchased through Open Exchange, PMP, and YouTube. For example, a campaign might allocate 50% of its budget to Open Exchange, 30% to PMP, and 20% to YouTube, where the PMP category is actually composed of three different subtypes of PMP deals. As described in Section 3, our model incorporates actual data from past campaigns regarding media cost, the number of impressions, the number of completed views, and so on. Before formally presenting the model, we review some of the related literature.

II. RELATED LITERATURE

Much of the early published research focused on trying to understand the effect of advertising on consumers and determine how to allocate a budget to various advertising products and media. Vidale and Wolfe (1957) present a model of how sales respond to advertising based on data from carefully designed and controlled advertising experiments. Observing that the impact of advertising lingers but diminishes over time, they develop a three-parameter nonlinear model that incorporates a sales decay constant, saturation level, and response constant. Once these parameters have been estimated for each product, an advertising budget can be optimally allocated among products. Thomas (1971) develops a multi-period profit maximizing linear programming (LP) model that uses advertising decisions to alter demand, thereby allowing production to be smoothed. He assumes that demand is a piecewise linear (concave) function of advertising, approximating the continuous response curve with saturation of Vidale and Wolfe (1957). Summarizing 26 papers published between 1937 and 1979, Zoltners and Sinha (1980) present an integer LP framework for optimally allocating sales resources (*e.g.,* sales budget, representatives, calls) among sales entities (*e.g.,* sales districts, accounts, products) in the presence of a wide variety of sales response functions (*e.g.,* linear, power, modified exponential, logit, concave, S-shaped). While not based on advertising decisions, the models presented in this article and some of those it references could likely be adapted to allocating an advertising budget among products and markets.

The Internet's arrival in the early 1990s presented marketers with a new advertising environment. Previously, it had been difficult to precisely measure the impact of ads on consumer behavior, so advertising response functions typically could only be estimated. By contrast, the Internet permits advertisers to directly measure how consumers react to ads using metrics such as the number of impressions and click-through rates (CTR, the ratio of clicks to total impressions). Currently, the two most common types of online advertising are display ads and search-based ads (Stone 2015), where display ads include banner ads, video ads and all ads other than text-based ads. Banner ads, the online analog of images found in traditional newspapers, are inserted into Web pages in the hope that consumers will click through to the advertiser's Web site. Search-based ads, which began to appear in the late 1990s, are generated in response to a consumer's query on a search engine site such as Google. Based on the phrases or *keywords* entered, search engines yield not only *organic* results corresponding to the popularity or relevance of Web sites for the consumer's query, but also *sponsored* results paid for by advertiser campaigns.

A variety of published papers focus on banner ads and how they ought to be presented to consumers or priced by publishers, *e.g.,* Langheinrich et al. (1999) propose a method to capture a minimal amount of data from users (*i.e.,* recent keywords, current Web page) reflecting their short-term interests, and then select a suitable advertisement from a pool of available banner ads; the ad selection process involves solving an LP that maximizes the expected CTR subject to satisfying the required display rate of each ad. Similarly, Chickering and Heckerman (2003) report on finding a delivery schedule (*i.e.,* ads shown in various segments of a Web site like msn.com in a certain time period) using an LP that maximizes the expected overall click-through probability of the Web site while satisfying both impression quotas for each ad and supply limits for each segment. Chatterjee et al. (2003) analyze clickstream data to determine the degree to which several variables influence whether or not a consumer will click on banner ads. Mangani (2004) examines the relative impact on publisher revenue from selling advertising space using CPV and cost-per-click pricing models.

The past decade has seen a shift toward research on search-based advertising, with a large number of papers focusing on how to optimally select keywords to include in an ad campaign, and the related decision of how much to bid on each keyword in auctions of paid search results. Rusmeivichientong and Williamson (2006) present an algorithm that takes as given a fixed budget, a large set of potential keywords to bid on, and initially unknown click-through rates associated with the ad for each keyword, and proceeds to adaptively select a subset of keywords that maximizes total expected profit. Ozluk and

Cholette (2007) formulate the advertiser's problem of deciding how to allocate a fixed daily budget among multiple keywords using an optimization model that maximizes revenue. Keyword bid prices determine placement of the corresponding ad, which ultimately determines the number of clicks and revenues. They show that the greater the value a particular keyword generates for the advertiser, the higher the advertiser should bid on that keyword relative to other keywords. Cholette et al. (2012) extend this model to incorporate stochastic ad positions and costs while maximizing expected profits; they illustrate just how limiting it can be to enforce a probabilistic budget constraint. Selcuk and Ozluk (2013) explore two variants of the deterministic keyword bidding problem which (1) minimize expected total costs (based on a cost-per-click scheme) while satisfying a desired level of exposure that is measured by the average CTR weighted by the numbers of keyword impressions, and (2) minimize expected total costs (based on a CPM scheme) while satisfying a desired level of exposures that is measured by the total number of impressions weighted by the relative ad positions of all keywords.

In recent years, video advertising has become steadily more prominent. Stone (2015) reports that, while video advertising constituted less than 5% of all Internet ad revenue in 2014, it is projected to be the fastest growing sub-segment of the global Internet advertising market from 2014-2019, reaching an 8% share by 2019, with spending jumping from $6.3 billion in 2014 to $15.4 billion in 2019. As this paper is concerned with the strategic decisions of how much of their budgets campaign managers should allocate to various sources of programmatic video advertising inventory, the paper's results should be increasingly relevant for campaign managers in the near future. We have not seen any literature addressing this issue. Our use of linear programming extends a long tradition of attempting to optimally allocate a budget among competing advertising products, but we believe this is its first application to the allocation of programmatic video advertising inventory, which has only come into existence in the past few years.

III. DATA FOR THE MODEL

Demand side platforms (DSP) allow buyers of digital advertising inventory to manage multiple ad exchange and data exchange accounts through one interface (Interactive Advertising Bureau 2014). One strength of this project is that all of its data come from a DSP through which video ads were actually purchased in 2016. By running inventory reports, we were able to collect a year's worth of real campaign data. The campaigns are broken out by inventory source type: PMP (Automated Guaranteed, Private, and Preferred), Open Exchange, and YouTube. In particular, one campaign data file contained 22,519 records of PMP and Open Exchange video ads, a small portion of which is shown in Figure 1 below. A second file consisted of 9,458 records of YouTube video ads. Each record represents a programmatic video ad placed by an advertiser and contains more than 20 fields, such as the advertiser's name (hidden in Figure 1), device type, number of impressions, measurable impressions, and viewable impressions, media cost, and number of completed views.

We created pivot tables in Excel to find the total cost and performance across all ads of each inventory source type. The resulting data, shown in Table 2 below, serve as input to the model. The report can be updated regularly by inputting more current campaign data, allowing campaign managers to easily revise this table and then generate a new inventory mix based on the most recent information.

	A	E	F	G	I	K	L	M	R	S	T	U
1	Month	Inventory Source Type	Max Video Duration (seconds)	Device Type	Impressions	Measurable Impressions	Viewable Impressions	Clicks	Media Cost	Click Rate (CTR)	Completed Views (Video)	Completion Rate (Video)
2	2015/11	Public Excha	15	Tablet	1864	1	1	12	25.74	0.64%	1378	73.93%
3	2015/11	Public Excha	15	Tablet	92132	34	24	30	370.73	0.03%	64742	70.27%
4	2015/12	Public Excha	30	Smart Phone	288	0	0	0	1.62	0.00%	198	68.75%
5	2016/01	Private Exch	15	Desktop	18736	18541	16685	41	326.23	0.22%	17827	95.15%
6	2016/01	Public Excha	15	Tablet	5269	0	0	55	15.91	1.04%	4915	93.28%
7	2016/02	Private Exch	30	Smart Phone	57	4	4	0	0.27	0.00%	49	85.97%
8	2016/02	Public Excha	20	Smart Phone	12	2	0	0	0.09	0.00%	8	66.67%
9	2016/03	Private Exch	20	Tablet	606	600	43	0	7.86	0.00%	415	68.48%
10	2016/03	Public Excha	20	Smart Phone	3	0	0	0	0.04	0.00%	1	33.33%
11	2016/03	Public Excha	15	Desktop	704956	689218	124347	233	7674.12	0.03%	542133	76.90%

FIGURE 1. A SMALL PORTION OF ONE OF THE CAMPAIGN DATA FILES.

TABLE 2. MODEL INPUTS FOR EACH OF THE FIVE INVENTORY SOURCE TYPES.

Model Input	**1. PMP-Private**	**2. PMP-Preferred**	**3. PMP-Guaranteed**	**4. Open Exchange**	**5. YouTube**
Number of Impressions	330,265,904	50,818,812	121,102,354	1,470,549,867	319,030,362
Total Cost ($)	5,369,378	780,376	2,230,053	13,050,003	4,408,550
Viewable Impressions	146,418,879	25,566,898	95,290,156	368,415,767	262,940,677
Measureable Impressions	253,098,686	37,067,395	101,780,563	935,424,395	279,490,281
Completed Views	247,230,657	40,793,109	103,027,630	1,054,532,155	77,212,329

IV. THE PROGRAMMATIC VIDEO ADVERTISING INVENTORY MODEL

We developed an LP model to find the best allocation of programmatic video advertising inventory. The algebraic formulation of the model will be presented first, followed by its implementation in Microsoft Excel. In this and subsequent sections, cells ranges refer to the Excel model shown in Figure 2 below.

Indices:

i = Advertising inventory source type (i = 1, 2, …, 5),
where 1 = PMP-Private;
2 = PMP-Preferred;
3 = PMP-Guaranteed;
4 = Open Exchange;
and 5 = YouTube.

Decision Variables (cells B21:F21):

$\boldsymbol{X}_i$ = Percentage of overall mix allocated to inventory type i (i = 1, 2, …, 5)

Input Data (cells B5:F10 and D25:D27):

NI_i = Number of impressions of inventory type i (i = 1, 2, …, 5)
TC_i = Total cost of inventory type i (i = 1, 2, …, 5)
VI_i = Viewable impressions of inventory type i (i = 1, 2, …, 5)
MI_i = Measurable impressions of inventory type i (i = 1, 2, …, 5)
CV_i = Completed views of inventory type i (i = 1, 2, …, 5)
MVR = Minimum viewability rate of the overall mix
MCR = Minimum completion rate of the overall mix
MYT = Minimum YouTube inventory Allocation

Calculated Quantities (cells B14:F17):

$CPM_i = (TC_i/NI_i)*1000$ = Cost per 1000 impressions of inventory type i (i = 1, 2, …, 5)
$VR_i = VI_i/MI_i$ = Viewability rate of inventory type i (i = 1, 2, …, 5)
$CR_i = CV_i/NI_i$ = Completion rate of inventory type i (i = 1, 2, …, 5)
$CPCV_i = TC_i/CV_i$ = Cost per completed view of inventory type i (i = 1, 2, …, 5)

Programmatic Video Advertising Inventory LP Model

Minimize $$\sum_{i=1}^{5} CPCV_i \boldsymbol{X}_i \quad (1)$$

Subject to:

$$\sum_{i=1}^{5} \boldsymbol{X}_i = 1 \quad (2)$$

$$\sum_{i=1}^{5} VR_i \boldsymbol{X}_i \geq MVR \quad (3)$$

$$\sum_{i=1}^{5} CR_i \boldsymbol{X}_i \geq MCR \quad (4)$$

$$\boldsymbol{X}_5 \geq MYT \quad (5)$$

$$0 \leq \boldsymbol{X}_i \leq 1, \text{ for } i = 1, 2, \ldots, 5 \quad (6)$$

The objective function (1) minimizes the total cost per completed view. Constraint (2) requires all of the inventory mix variables to add to 100%. Constraints (3) and (4) force the overall viewability and completion rates to meet or exceed their respective threshold values; constraint (5) likewise forces the allocation of YouTube inventory to meet or exceed its target value. We note that while YouTube requires advertisers to specify a daily budget that's greater than \$0, and also possibly enter a maximum bid (such as \$0.10 CPV), it does not require a minimum number of impressions. If an advertiser's bid is too low to win any impressions on YouTube, then it will get 0 impressions. There are minimum spends required for Automated Guaranteed and the other PMP inventory sources, but since most of the negotiations are done at the advertiser or agency level, an individual campaign can still have 0 impressions or spend $0 on a given inventory source. The decision variables are continuous and clearly must take on values between 0 and 100%, as given in (6). The LP was implemented in Excel (Figure 2); with just five variables and four constraints, it is optimally solved quickly using Excel's well-known Solver add-in (Fylstra et al., 1998).

	A	B	C	D	E	F
1	**Optimal Video Advertising Inventory Mix**					Avnish Brar, Olga Pipko
2						
3	**Inputs**			Inventory Source Type		
4		PMP-Private	PMP-Preferred	PMP-Guaranteed	Open Exchange	YouTube
5	Number of Impressions	330,265,904	50,818,812	121,102,354	1,470,549,867	319,030,362
6	Total Cost	5,369,378	780,376	2,230,053	13,050,003	4,408,550
7	Viewable Impressions	146,418,879	25,566,898	95,290,156	368,415,767	262,940,677
8	Measureable Impressions	253,098,686	37,067,395	101,780,563	935,424,395	279,490,281
9	Completed Views	247,230,657	40,793,109	103,027,630	1,054,532,155	77,212,329
10	Media Budget	$200,000				
11						
12	**Calculated Quantities**			Inventory Source Type		
13		PMP-Private	PMP-Preferred	PMP-Guaranteed	Open Exchange	YouTube
14	CPM	$16.26	$15.36	$18.41	$8.87	$13.82
15	Viewability Rate	58%	69%	94%	39%	94%
16	Completion Rate	75%	80%	85%	72%	24%
17	CPCV	$0.0217	$0.0191	$0.0216	$0.0124	$0.0571
18						
19	**Decision Variables**			Inventory Source Type		
20		PMP-Private	PMP-Preferred	PMP-Guaranteed	Open Exchange	YouTube
21	Inventory Mix	**0.0%**	**0.0%**	**41.3%**	**43.7%**	**15.0%**
22						
23	**Constraints**	LHS		RHS		
24	Inventory Mix	100.0%	=	100.0%		
25	Viewability Rate	70.0%	≥	70.0%		
26	Completion Rate	70.1%	≥	70.0%		
27	YouTube Threshold	15.0%	≥	15.0%		
28						
29	**Objective to Minimize**					
30	Overall CPCV	**0.0229**				
31						
32	**Other Outputs**					
33	Total Impressions	14,751,507				
34	Overall CPM	$13.56				
35		PMP	Open Exchange	YouTube		
36	Budget Allocation	$82,639	$87,361	$30,000		

FIGURE 2. THE PROGRAMMATIC VIDEO ADVERTISING LP MODEL IN EXCEL.

In the Excel model, the decision variables are contained in cells B21:F21, while the left and right hand sides of constraints (2-5) are located in cells B24:B27 and D24:D27, respectively. The objective function is calculated in cell B30. Other outputs are also generated by the model to assist the campaign manager, *e.g.,* the overall CPM is computed in cell B34 as a weighted average of the individual inventory source type CPMs (B14:F14) and the inventory mix. The total number of impressions (B33) can then be computed as 1000 times the ratio of the media budget (B10) to the overall CPM. Finally, applying the optimal allocation to the overall media budget gives the amount of budget allocated to each major inventory type (B36:D36).

V. RESULTS

For the base case, we set both the minimum viewability and completion rate thresholds to 70%, which would be considered a high-performing campaign; the minimum YouTube requirement was set to 15%, since YouTube is a necessary inventory source. For the initial budget, we used a typical campaign budget of $200,000 (cell B10). Under these conditions, solving the LP yields a minimum overall CPCV of $0.0229, achieved with 41.3% of the budget allocated to PMP-Guaranteed, 43.7% to Open Exchange, and 15.0% to YouTube. In the optimal solution, all of the PMP budget is allocated to PMP-Guaranteed. This makes logical sense because, based on the calculated quantities in B14:F17, PMP-Guaranteed inventory has significantly higher average performance (viewability and completion rates) than both PMP-Preferred and PMP-Private, while its CPCV is about the same as PMP-Private and just slightly higher than PMP-Preferred.

In 2016, the company's average viewability and completion rates for video campaigns actually run through this particular DSP were 47.3% and 73.3%, respectively; the overall CPCV without allocating anything to YouTube was $0.0148. When we run the model with identical viewability and completion rate thresholds of 47.3% and 73.3%, respectively, and YouTube's threshold set to 0%, we find an optimal inventory mix that achieves the same performance as in 2016, but with a CPCV of $0.0137, a 7.43% reduction over the cost actually incurred in 2016. While this reduction may not seem like a big difference, when applied to the total media spend in 2016 of $21.4 million, it means that the optimal inventory mix could have saved the company more than $1.5 million in total media spend. Consequently, we believe the LP model is capable of providing significantly better solutions than those found by manually allocating budget to the different inventory source types.

VI. SENSITIVITY ANALYSES

In initial experiments with the model, we noticed that the minimum YouTube allocation constraint was always binding on the optimal solution, so we first considered the model's sensitivity to the YouTube threshold (MYT). While YouTube is generally viewed as the most valuable video Internet platform, some businesses have recently started to pull their advertising off YouTube, not so much because of its high cost, but because its automated system sometimes places ads for their brands next to offensive material such as hate speech (Wakabayashi and Maheshwari, 2017). Nonetheless, YouTube remains an extremely popular site for a wide variety of web users, so it seems wise to allocate at least part of one's advertising budget to YouTube inventory. We therefore varied the minimum allocation for YouTube inventory from 1% to 24% (keeping the viewability and completion rate parameters at their base case values) using the SolverTable add-in for Excel (Winston and Albright 2016), with results shown in Figure 3 and Table 3 below.

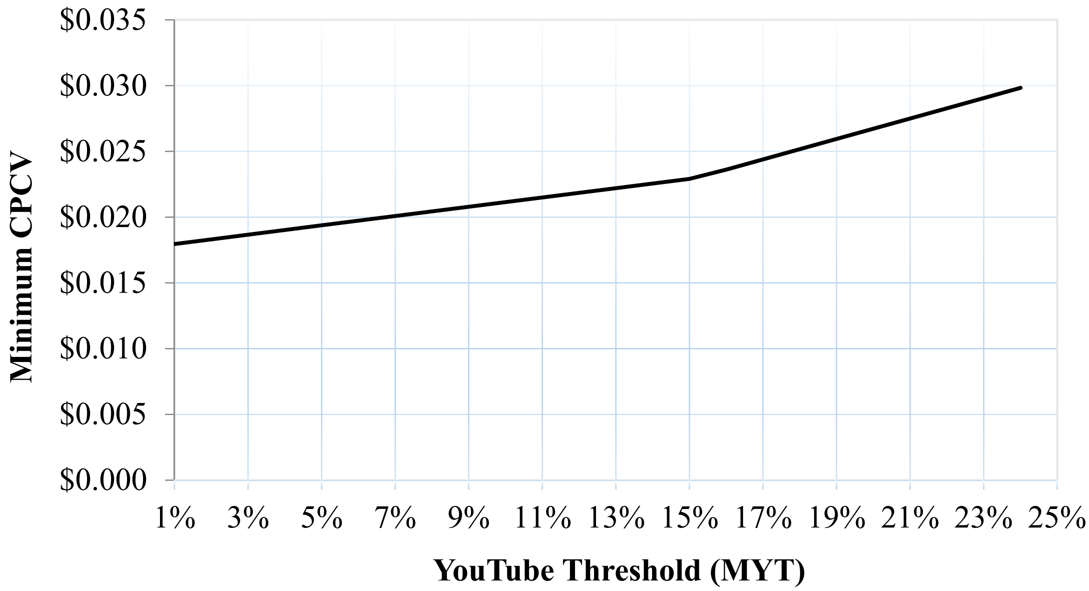

FIGURE 3. GRAPHICAL VIEW OF YOUTUBE THRESHOLD'S IMPACT.

The YouTube threshold clearly has a strong impact on the optimal inventory mix and CPCV. The company can achieve a CPCV as low as $0.0180 when YouTube constitutes just 1% of the mix. For each additional percentage point increase in MYT up to its base case value of 15%, the minimum CPCV rises about $0.00035. For each additional percentage point increase in MYT above 15%, the minimum CPCV increases by $0.00077, more than twice as fast as when MYT is below 15%. By the time MYT gets to 24%, the minimum CPCV reaches $0.0298, a 30% increase over its base case value of $0.0229. (For values of MYT above 24%, the LP is infeasible.) As for the optimal mix itself, PMP-Guaranteed's allocation drops by the amount allocated to YouTube, as it grows from 1% to 15%, while Open Exchange's allocation remains steady at just under 44%. However, once YouTube's allocation rises above 15%, Open Exchange's allocation drops quickly with a corresponding increase in PMP-Guaranteed's allocation.

Campaign managers can use the model to input their own campaign budgets and desired campaign viewability and completion rate thresholds, so we also examined how various combinations of these two parameters (MVR and MCR) would impact the minimum CPCV (leaving MYT at its base case value of 15%). Confirming expectations, Table 4 shows that higher viewability and completion rate thresholds generally lead to a higher minimum CPCV. The table quantifies for campaign managers the cost tradeoffs that would have to be made in order to achieve better campaign results. For example, to reach viewability and completion rates of 75% instead of 70%, the minimum CPCV would rise by 15%, from $0.0229 to $0.0263. On the other hand, if the campaign manager feels that viewability and completion rates of 65% are adequate, the minimum CPCV would fall by less than 4%, from $0.0229 to $0.0212.

TABLE 3. BASE MODEL'S SENSITIVITY TO YOUTUBE THRESHOLD.

MYT	PMP-Private	PMP-Preferred	PMP-Guaranteed	Open Exchange	YouTube	Min. CPCV
1%	0.0%	0.0%	55.4%	43.6%	1.0%	$0.0180
2%	0.0%	0.0%	54.4%	43.6%	2.0%	$0.0183
3%	0.0%	0.0%	53.4%	43.6%	3.0%	$0.0187
4%	0.0%	0.0%	52.4%	43.6%	4.0%	$0.0190
5%	0.0%	0.0%	51.4%	43.6%	5.0%	$0.0194
6%	0.0%	0.0%	50.4%	43.6%	6.0%	$0.0197
7%	0.0%	0.0%	49.4%	43.6%	7.0%	$0.0201
8%	0.0%	0.0%	48.4%	43.6%	8.0%	$0.0204
9%	0.0%	0.0%	47.4%	43.6%	9.0%	$0.0208
10%	0.0%	0.0%	46.4%	43.6%	10.0%	$0.0211
11%	0.0%	0.0%	45.4%	43.6%	11.0%	$0.0215
12%	0.0%	0.0%	44.3%	43.7%	12.0%	$0.0219
13%	0.0%	0.0%	43.3%	43.7%	13.0%	$0.0222
14%	0.0%	0.0%	42.3%	43.7%	14.0%	$0.0226
15%	**0.0%**	**0.0%**	**41.3%**	**43.7%**	**15.0%**	**$0.0229**
16%	0.0%	0.0%	44.1%	39.9%	16.0%	$0.0236
17%	0.0%	0.0%	47.6%	35.4%	17.0%	$0.0244
18%	0.0%	0.0%	51.2%	30.8%	18.0%	$0.0252
19%	0.0%	0.0%	54.7%	26.3%	19.0%	$0.0259
20%	0.0%	0.0%	58.3%	21.7%	20.0%	$0.0267
21%	0.0%	0.0%	61.9%	17.1%	21.0%	$0.0275
22%	0.0%	0.0%	65.4%	12.6%	22.0%	$0.0283
23%	0.0%	0.0%	69.0%	8.0%	23.0%	$0.0291
24%	0.0%	0.0%	72.5%	3.5%	24.0%	$0.0298

TABLE 4. SENSITIVITY OF MINIMUM CPCV TO MVR AND MCR THRESHOLDS.

	MCR					
MVR	50%	55%	60%	65%	**70%**	75%
50%	$0.0195	$0.0195	$0.0195	$0.0195	$0.0228	$0.0263
55%	$0.0203	$0.0203	$0.0203	$0.0203	$0.0228	$0.0263
60%	$0.0212	$0.0212	$0.0212	$0.0212	$0.0228	$0.0263
65%	$0.0221	$0.0221	$0.0221	$0.0221	$0.0228	$0.0263
70%	$0.0229	$0.0229	$0.0229	$0.0229	**$0.0229**	$0.0263
75%	$0.0238	$0.0238	$0.0238	$0.0238	$0.0238	$0.0263
80%	$0.0246	$0.0246	$0.0246	$0.0246	$0.0246	$0.0263
85%	$0.0255	$0.0255	$0.0255	$0.0255	$0.0255	$0.0263
90%	$0.0263	$0.0263	$0.0263	$0.0263	$0.0263	$0.0263

VII. CONCLUSION AND FUTURE WORK

The goal of any advertising campaign is to reach the largest audience possible in the most effective way. With many different advertising inventory types available, it can be difficult for campaign managers to manually decide how to allocate their advertising budget. We've attempted to solve this problem by developing a model that can help campaign managers choose the optimal inventory mix that minimizes the overall CPCV while meeting the campaign KPIs. This method has significant advantages over choosing the optimal inventory mix based on intuition because it takes into account live data instead of just guessing based on a previous campaign's success or failure. Overall, our model should prove to be a very useful tool that will allow campaign managers to develop advertising inventory mixes based on actual performance data.

An immediate extension to the current model could be to include different device types to allow the campaign manager to select an optimal inventory mix by device. By comparing inventory performance by device, we could create a model that breaks down the optimal inventory mix for desktop and mobile devices, although this may be risky if there are not enough data available per device type. Also, if one device performs better than the others, then the model will prefer to run as much of the inventory through that device as possible. In order to accommodate that, the model might need to include a minimum percentage mix per device, and also find the maximum number of impressions available for each device.

REFERENCES

Busch, O. (ed.), *Programmatic Advertising: The Successful Transformation to Automated, Data-Driven Marketing in Real-Time*, Springer International, Cham, Switzerland, 2016.

Chatterjee, P., Hoffman, D. and Novak T., "Modeling the Clickstream: Implications for Web-Based Advertising Efforts," *Marketing Science*, 22(4), 2003, 520-541.

Chickering, D. and Heckerman, D., “Targeted Advertising on the Web with Inventory Management,” *Interfaces*, 33(5), 2003, 71-77.

Cholette, S., Ozluk, O., and Parlar, M., “Optimal Keyword Bids in Search-Based Advertising with Stochastic Ad Positions, *Journal of Optimization Theory and Applications*, 152(1), 2012, 225-244.

Fylstra, D., Lasdon, L., Watson, J. and Waren, A., “Design and Use of the Microsoft Excel Solver,” *Interfaces*, 28(5), 1998, 29-55.

Interactive Advertising Bureau, “How an Ad is Served with Real Time Bidding,” IAB Digital Simplified, 6/19/14**,** https://www.youtube.com/watch?v=-Glgi9RRuJs (accessed 6/23/17)

Langheinrich, M., Nakamura, A., Abe, N., Kamba, T., and Koseki, Y., “Unintrusive Customization Techniques for Web Advertising.” *Computer Networks* 31, 1999, 1259-1272.

Mangani, A., “Online Advertising: Pay-per-View Versus Pay-per-Click,” *Journal of Revenue and Pricing Management*, 2(4), 2004, 295-302.

Ozluk, O. and Cholette, S., “Allocating Expenditures Across Keywords in Search Advertising,” *Journal of Revenue and Pricing Management*, 6(4), 2007, 347-356.

Rusmeivichientong, P. and Williamson, D., “An Adaptive Algorithm for Selecting Profitable Keywords for Search-Based Advertising Services,” In *Proceedings of the 7th ACM Conference on Electronic Commerce* (EC 2006). ACM, New York, NY, 260-269.

Selcuk, B. and Ozluk, O., “Optimal Keyword Bidding in Search-Based Advertising with Target Exposure Levels,” *European Journal of Operational Research*, 226, 2013, 163-172.

Stone, M., “World Media Trends: Global Digital Revenue Trends 2015,” World Newsmedia Network, http://www.magazine.org.tw/ImagesUploaded/news/14404738842160.pdf (accessed June 22, 2017).

Thomas, J., “Linear Programming Models for Production-Advertising Decisions,” *Management Science*, 17(8), 1971, B474-B484.

Vidale, M. L. and Wolfe, H.B., “An Operations Research Study of Sales Response to Advertising,” *Operations Research*, 5(3), 1957, 370-381.

Wakabayashi, D. and Maheshwari, S., “Why Some Advertisers Flee YouTube,” *San Francisco Chronicle,* March 25, 2017, D2.

Winston, W. L. and Albright, S. C., *Practical Management Science* (5th ed.), Cengage Learning, Boston, MA, 2016.

Zoltners, A. and Sinha, P., “Integer Programming Models for Sales Resource Allocation,” *Management Science*, 26(3), 1980, 242-260.

APPENDIX:

DEFINITIONS OF KEY TERMS

Advertiser: An organization, entity, or individual that provides advertisements to be displayed on the publisher's content.

Advertising Inventory: The number of advertisements, or amount of ad space, a publisher has available to sell to an advertiser.

Completion Rate: The percentage of impressions that have completed in full. A completed ad runs its full length, whether 7, 15, or 30 seconds, before the user exits the web page or closes the ad window. DSPs can log this information on ads placed on different websites, and determine how much of each video ad ran (which they typically show as 25%, 50%, or 100% of each ad).

CPM: The cost per 1000 impressions, *i.e.,* 1000*Total Cost / Number of Impressions.

CPCV: The cost per completed view, *i.e.,* Total Cost / Number of Completed Views.

DSP: Demand Side Platform, or system that allows buyers of digital advertising inventory to manage multiple ad exchange and data exchange accounts through one interface.

Impression: When an ad is fetched from its source and is countable. Whether the ad is clicked is not taken into account. Impressions are the same for video and banner ads.

Measurable Impression: An impression that a third party ad trafficking service can verify was placed on a Web page. Third party tags track how long the video ad ran, and if it was 50% in view.

PMP: Private Marketplace, which includes three different purchasing methods for advertisers: Automated Guaranteed, Private, and Preferred.

Publisher: An organization, entity, or individual that integrates advertisements into its online content.

Viewable Impression: An impression where 50% of the video ad is seen by the user for at least two consecutive seconds. The definition may depend on the type of ad unit and reporting system.

Viewability Rate: The ratio of the number of viewable impressions to the number of measurable impressions.