The Optimal Mix of TV and Online Ads to Maximize Reach

Yuxue Jin, Georg M. Goerg, Nicolas Remy, Jim Koehler

Google Inc.

Last update: September 20, 2013

Abstract

Brand marketers often wonder how they should allocate budget between TV and online ads in order to maximize reach or maintain the same reach at a lower cost. We use probability models based on historical cross media panel data to suggest the optimal budget allocation between TV and online ads to maximize reach to the target demographics. We take a historical TV campaign and estimate the reach and GRPs of a hypothetical cross-media campaign if some budget was shifted from TV to online. The models are validated against simulations and historical cross-media campaigns. They are illustrated on one case study to show how an optimized cross-media campaign can obtain a higher reach at the same cost or maintain the same reach at a lower cost than the TV-only campaign.

1 Introduction

In the era of digital media, more brand marketers are including online video ads in their advertising campaigns. They often wonder whether online ads should be included in the campaign, and what should be the right mix between online ads and TV ads in order to reach more of their target audience. To answer the first question, Jin et al. [1] proposed a statistical method to compare the cost efficiency of online ads over TV ads in historical cross-media campaigns. To get the right mix between online and TV ads, in this paper we present probability models to find the optimal budget allocation between TV and online ads in two media planning scenarios. In the first scenario, the advertiser aims to maximize reach with the total campaign budget fixed; in the second scenario, the advertiser aims to minimize total cost while maintaining the same reach as the original TV campaign.

The models are based on users' multi-media usage and ad exposure collected from cross media panels calibrated to the population. We use probability models to estimate the reach of hypothetical online campaigns. Furthermore, we take a historical TV campaign and estimate the reach of a hypothetical cross-media campaign if some TV programs were removed from the TV plan and the corresponding budget was shifted to online ads. Based on this model, we can find the optimal mix of TV and online budget to maximize reach.

In Section 2, we describe the probability models which are validated against historical cross-media campaigns in Section 3. In Section 4, we illustrate the methodology on a case study and show that the optimized cross-media plan is more cost efficient than the TV-only plan in the two media-planning scenarios.

2 Methodology

In cross media panels, the TV and online usage of individuals in the households are measured and linked together on individual level by matching personal login IDs. For illustration purpose, online ads are restricted to YouTube Homepage Mastheads and Watchpage ads, though the methodology can be extended to include other online ads as well. On YouTube, the watchpage is the homepage of a video, where all the public information of the video resides. For details on advertising products on YouTube, please refer to the web page in [2]. Our analysis is based on ad impressions of historical TV campaigns, visits to YouTube homepage and all monetizable watchpage views. Ad impressions of historical YouTube campaigns are not needed, except to validate the YouTube models, since we deal with hypothetical YouTube campaigns.

To estimate the reach of a hypothetical YouTube campaign, we need two user-level models: a YouTube watchpage model and a YouTube homepage model. They estimate the probability distribution of the number of ad impressions a user would have on YouTube, conditioned on the user's daily visits to YouTube homepage and all monetizable watchpage views during the campaign period. A third user-level model is needed to estimate the reach of a hypothetical cross-media campaign if some budget is shifted from TV to YouTube, the TV model. The TV model estimates the probability distribution of ad impressions a user would have on TV in the what-if scenario of a budget cut from TV, conditioned on the impressions the user had in the original TV plan.

These three models are combined to estimate the total number of ad impressions on TV and online for each user, and the probability of being reached at least k times. Aggregating these individual probabilities of being reached across all users in the target audience would lead to an estimate of the $k+$ reach of the hypothetical cross-media campaign. To find the optimal mix in the first media planning scenario, we try different budget allocations between TV and online while holding the total budget fixed and choose the mix with the largest combined reach. In the second

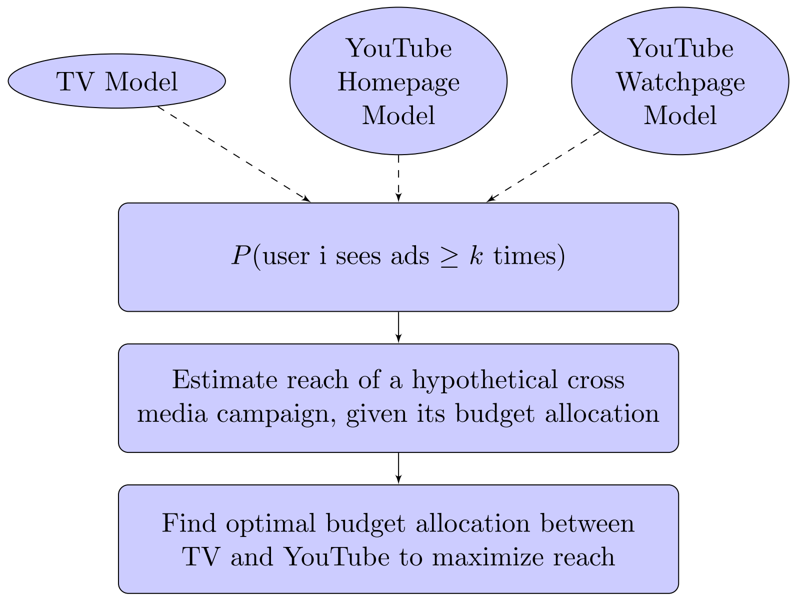

Figure 1: Hierarchy of probability models to evaluate and optimize media-mix strategies

scenario, we find the minimum total budget which can deliver the same reach as the original TV campaign, if the budget is allocated between TV and online optimally as in Scenario I. The hierarchy of the models are shown in Figure 1, and the details can be found in the following sections.

2.1 Probability Models

2.1.1 TV

In the situation of a budget cut from the historical TV plan, we do not know which programs and associated spots would be removed. Instead of doing sophisticated TV planning for reduced TV budgets, we assume all the programs in the original TV plan are equally likely to be removed, an assumption that we will relax later on.

Suppose the portion of budget shifted from TV to online is s, $0 \leq s \leq 1$, and there are L programs in the original TV plan. The ith user in the panel had c_{ij} ad impressions from the jth program. After the budget shift, we use a Bernoulli random variable Y_j to denote whether the jth program remains in the plan. Let $f_{i,T}$ denote the number of ad impressions the ith user would have on TV after the budget cut, it can be represented as

$$f_{i,T} = \sum_{j=1}^{L} c_{ij} I_{(Y_j=1)}, \tag{1}$$

where $Y_j = 1$ means the jth program remains in the TV plan.

Let b_j denote the cost of the spots associated with the jth program, to have a budget shift of s, $Y_j, j = 1, \ldots, L$, should satisfy

$$\sum_{j=1}^{L} b_j I_{(Y_j=0)} = s \sum_{j=1}^{L} b_j. \tag{2}$$

Since all programs are equally likely to be removed from the original TV plan, Y_j's are symmetric and have identical distributions. Taking expectation on the left-hand side of (2), we get

$$P(Y_j = 1) = 1 - s, \qquad j = 1, \ldots, L. \tag{3}$$

As L is usually very large (larger than 100), the correlation imposed by the constraint (2) is very weak. We find that assuming independence of Y_j's is a good approximation and simplifies the calculation of the distribution of (1) into convolving L Bernoulli variables, each with a success probability of $1 - s$.

To get the expected number of ad impressions of the ith user on TV after the cut, we take expectations of (1),

$$\mathbb{E}(f_{i,T}) = \mathbb{E}\left(\sum_{j=1}^{L} c_{ij} I_{(Y_j=1)}\right) = \sum_{j=1}^{L} c_{ij} P(Y_j = 1) = (1-s) \sum_{j=1}^{L} c_{ij}. \tag{4}$$

In other words, on average, the number of ad impressions on TV for a user would be reduced by the same portion as the budget cut.

If the advertiser has a preference as to which programs or TV networks should be cut first, we can allow different probabilities for different programs to be removed. The distribution of (1) can be calculated similarly, except allowing Y_j's to have different success probabilities. Assuming all programs are equally likely to be removed would likely result in a sub-optimal TV plan in the case of a budget cut from TV. If inefficient programs are removed first or a TV plan better than randomly cutting is used, the optimal cross-media plan would deliver higher combined reach and higher reach uplift than using a randomly-cut TV plan.

2.1.2 YouTube Watchpages

We assume our hypothetical YouTube watchpage campaign is bought through the "reservation" process. With the budget shifted from TV to YouTube watchpages, at a given CPM (Cost per Mille[1]), the advertiser can buy M_0 impressions, out of the total available inventory M during the campaign period.

[1] the cost of one thousand impressions

Based on historical panel data, in the campaign period, the ith user had m_i monetizable watchpage views that satisfy the specific requirement of the advertiser on content or verticals. In other words, if the hypothetical campaign were indeed run, it would be possible to serve the campaign ad on any of the m_i views.

The mechanism by which the YouTube server decides which campaign to serve on a qualified watchpage view is complicated. We approximate it by assuming the probability that a qualified watchpage view serves the ad of the campaign in interest is proportional to the size of the campaign, and further assume that ad serving of different views of the same user is independent. Let $f_{i,W}$ denote the number of watchpage ad impressions the ith user would have in the hypothetical campaign, it follows the binomial distribution with m_i independent trials and success probability of M_0/M. Hence the expected number of watchpage ad impressions of the ith user is $\mathbb{E}(f_{i,W}) = m_i M_0/M$.

To avoid showing too many impressions to the same user, the advertiser can impose a frequency cap on the campaign. With a frequency cap of f_c per campaign, the server will not show the ads to the same cookie more than f_c times. To model such campaigns, we assume a user has only one cookie and modify the binomial model above into a truncated binomial model. The success probability is chosen so that the total expected number of impressions delivered in the population according to the truncated model is equal to the impressions bought, which is equivalent to solving the following equation in terms of p.

$$\sum_i w_i \left(\sum_{j=1}^{f_c-1} jP(f_{i,W} = j; p) + f_c P(f_{i,W} \geq f_c; p) \right) = M_0, \tag{5}$$

where the sum is taken over all individuals in the panel, w_i is the demographic weight of the ith user, and $f_{i,W}$ is as described above and follows the binomial distribution with success probability p.

Some advertisers only want to show their ads to people in their target demographics. To minimize waste and spill-over of impressions to people out of the target demographics, the advertiser can use the option of advanced demographic targeting in the watchpage campaign. In this case, we would only count qualified monetizable views associated with cookies that have YouTube declared demographics within the advertiser's target. If a user never logs in or logs in with declared demographics out of the advertiser's target, he would never see the ad. The total inventory M will also change, and is typically much smaller than the inventory without advanced demographic targeting.

2.1.3 YouTube Homepage

Suppose the campaign ran for D days, and the advertiser bought the YouTube masthead on d days in the hypothetical cross-media campaign. We do not specify which d days are bought, and treat each day as equally likely to be bought. Based on panel data, the ith user had v_{ij} visits to the YouTube homepage on the jth day, $j = 1, \ldots, D$. The number of homepage ad impressions the ith user would have if d days were bought is

$$f_{i,H} = \sum_{j=1}^{D} v_{ij} I_{(Y_j=1)}, \tag{6}$$

where $I_{(Y_j=1)}$ denotes that the jth day is bought and satisfies $\sum_{j=1}^{D} I_{(Y_j=1)} = d$. Since the probability of each day to be bought is the same, the Y_j's are symmetric and

$$P(Y_j = 1) = d/D, \quad j = 1, \ldots, D. \tag{7}$$

If the campaign duration is long, i.e., D is large, we can use similar argument as in Section 2.1.1, treat Y_j's as independent and calculate the distribution of (6) by convolving D independent Bernoulli variables. However, for a short campaign, we cannot ignore the correlation among Y_j's. Hence, we use simulations to approximate the distribution of $f_{i,H}$. We randomly draw d out of D days, and for each draw count the number of homepage visits the ith user had on the chosen days. This is one realization of $f_{i,H}$; repeating the random draws for a large number would generate the distribution of $f_{i,H}$.

To calculate the expected number of homepage ad impressions of the ith user, we use a very simple formula without doing the simulations.

$$\mathbb{E}(f_{i,H}) = \mathbb{E}\left(\sum_{j=1}^{D} v_{ij} I_{(Y_j=1)}\right) = \sum_{j=1}^{D} v_{ij} P(Y_j = 1) = \frac{d}{D} \sum_{j=1}^{D} v_{ij}. \tag{8}$$

2.2 Combining the Probability Models

The above-mentioned probability models are combined to estimate the total number of ad impressions per user on TV and online in a hypothetical cross-media campaign. Let $f_i = f_{i,T} + f_{i,W} + f_{i,H}$ denote the number of ad impressions for the ith user. The three integer-valued random variables, $f_{i,T}$, $f_{i,W}$ and $f_{i,H}$, are independent given the user's media consumption on TV and YouTube. To see why, we briefly revisit their definitions in Sections 2.1.1 - 2.1.3. $f_{i,T}$ is a random variable conditioned on the user's exposure to the original TV campaign, and the randomness is from the execution of budget cut from TV, i.e., which programs to remove. $f_{i,W}$ is conditioned on the user's monetizable watchpage views during the campaign period, and the randomness is caused by the ad serving mechanism of YouTube. Similarly, $f_{i,H}$ is conditioned on the user's visits to YouTube homepage, and the randomness is caused by choosing the d days to buy the homepage. Therefore, we calculate the probability distribution of f_i by convolving the three independent distributions in Sections 2.1.1 - 2.1.3.

To calculate the k+ reach of the cross-media campaign R_k, we need $P(f_i \geq k)$ for all users in the target demographics in the panel. For $k = 1$, we can explicitly express it as

$$P(f_i \geq 1) = 1 - P(f_i = 0) = 1 - P(f_{i,T} = 0)P(f_{i,W} = 0)P(f_{i,H} = 0). \qquad (9)$$

For $k = 3$ or 5, the expansion of $P(f_i \geq k)$ into $f_{i,T}, f_{i,W}$ and $f_{i,H}$ is complicated, and we would rather use a fast Fourier transform than hard code the formulas.

Once we have calculated $P(f_i \geq k)$ for the N users in the target demographics, the k+ reach of the campaign in the target demographics is estimated by

$$\mathbb{E}(R_k) = \mathbb{E}\left(\frac{1}{W}\sum_{i=1}^{N} w_i I_{(f_i \geq k)}\right) = \frac{1}{W}\sum_{i=1}^{N} w_i P(f_i \geq k), \qquad (10)$$

where w_i is the demographic weight of the ith user, and $W = \sum_{i=1}^{N} w_i$ is the projected population size of the target demographics. The variance of the k+ reach is

$$\mathbb{V}(R_k) = \mathbb{V}\Big(\frac{1}{W}\sum_{i=1}^{N} w_i I_{(f_i \geq k)}\Big) = \frac{1}{W^2}\sum_{i=1}^{N} w_i^2 P(f_i \geq k)(1 - P(f_i \geq k)). \qquad (11)$$

Here we ignore the sampling variance and only focus on the variance of reach in the sampled users in the panel caused by shifting budget from TV to online. In other words, if there is no shift from TV to online, the reach would always be equal to that of the historical TV campaign measured in the panel and the variance would be zero.

Furthermore, we estimate the total number of GRPs delivered to the target demographics by

$$\frac{1}{W}\sum_{i=1}^{N} w_i \mathbb{E}(f_i) \times 100 = \frac{100}{W}\sum_{i=1}^{N} w_i \left(\mathbb{E}(f_{i,T}) + \mathbb{E}(f_{i,W}) + \mathbb{E}(f_{i,H})\right), \qquad (12)$$

where $\sum_{i=1}^{N} w_i \mathbb{E}(f_i)$ is the estimated number of impressions delivered to the target audience.

Besides estimating reach and GRPs in the entire target demographics, we can break it into subgroups and estimate reach and GRPs in each subgroup. For example, a user can be assigned to a light, medium, or heavy TV viewing group according to his average TV viewing time per day. Reach and GRPs in a subgroup are estimated similarly using (10) and (12), except only summing over users in the subgroup. An example can be found in Section 4.

2.3 Optimization

Given a budget reallocation between TV and YouTube, we can estimate the reach and GRPs of the hypothetical cross-media campaign. In Scenario I where the total campaign budget is held constant, the media mix is determined by two parameters, the TV budget share and the number of YouTube mastheads bought. The number of watchpage impressions bought can be inferred from these two parameters plus the total budget, YouTube watchpage CPM and the cost of one masthead.

Optimization is conducted over the two parameters to find the mix that gives the largest combined reach. Possible candidates of the optimization algorithm are brute force search on a pre-specified grid or the Newton-Raphson method. The advertiser decides the target demographics of the campaign, whether certain programs or channels should be cut from TV first, and whether the watchpage campaign should use advanced demographic targeting or not. These additional features in designing the cross-media campaign are treated as input into the optimization algorithm.

In Scenario II where the advertiser wants to maintain the reach while cutting costs, we do the optimization in Scenario I for each reduced campaign budget and estimate the reach of the optimized cross-media plan. Then we search for the budget whose optimal reach is equal to the reach of the original TV plan. Details are explained in Section 4.

3 Model Validation

To validate the TV model, we simulate reduced TV plans by randomly removing TV programs from the original plan. For each simulated plan, we calculate its cost, reach and GRPs based on panel data. Then we compute the average reach and GRPs of plans with the same or similar cost. This is compared with the estimated reach of the TV model after a budget cut, using (10) and (12) with f_i replaced by $f_{i,T}$. Two reach curves are generated, one based on simulations, the other based on the TV model.

Figure 2 shows the two reach curves as well as the cumulative reach curve of the original TV campaign. The left figures are the 1-plus, 3-plus, and 5-plus reach curves of a CPG campaign in Germany in Q3-Q4 2012 targeting females 30-59; the right figures are those of a campaign of a large retailer in the US in Q4 2012 targeting adults 20-59. The German campaign is based on GfK's cross media panel [3], with the TV media plan from Ebiquity [4]; the US campaign is based on Nielsen's Cross Media Panel [5], with the TV plan from Nielsen's Monitor Plus. The reach curve based on simulations is almost the same as that based on the TV model.

The YouTube watchpage model is validated on historical cross-media campaigns. Given the number of impressions bought and other features of the campaign, we

Table 1: Reach (in %) of historical cross-media campaigns. Numbers in the parentheses are half width of the 95% confidence intervals.

Campaign / Media	TV	Watchpage		Combined	
	Observed	Observed	Estimate	Observed	Estimate
A web browser	71.33	13.69	14.38 (0.47)	74.91	75.07 (0.22)
CPG	65.07	2.23	2.98 (0.90)	65.53	65.85 (0.43)
CPG	66.48	3.21	4.21 (0.78)	67.19	67.72 (0.39)
Telecommunications	48.46	0.72	0.94 (0.41)	48.80	48.93 (0.29)
Movie Release	61.63	0.49	0.65 (0.35)	61.77	61.86 (0.20)

estimate its reach and the combined reach with TV using (10) and (12) with f_i replaced by $f_{i,W}$ and $f_{i,W} + f_{i,T}$. The numbers are compared with the panel-measured numbers based on watchpage impressions identified through tagging the online ads. Table 1 shows the results of five historical cross-media campaigns in Germany. The estimated reach matches the panel-measured reach reasonably well. For three campaigns, the difference is within the confidence interval. For the other two campaigns, the YouTube watchpage campaign was targeted at some demographic groups, while the model assumes there is no demographic targeting. This probably causes the model to overestimate watchpage reach.

For the YouTube homepage model, we do not do any validation since it simulates the scenario of randomly picking d out of D days to buy the homepage masthead.

4 Case Study

We illustrate the method on a typical TV campaign running in the United States from 2012-10-15 to 2012-12-31 (77 days).

Figure 3 shows that the campaign achieves a maximum 3+ reach of 63.72% to the target audience (males 18-49) after 990 GRPs and a total budget of $40 million ($40.4 thousand / GRP).[2] The red curve shows our model estimate for evaluating TV reach of campaign with lower budget [see 1, for details]. Note that as the estimated red curve lies above the data points, our models are based on a conservative view on the effects of shifting money from TV to online advertising.

4.1 Scenario I: Maximize Reach at Fixed Budget

In this scenario we hold budget constant and maximize TV & YouTube combined reach by shifting budget from TV to YouTube. Since the number of days to buy

[2] To estimate the campaign budget we use Nielsen's rate-card prices with details on costs of individual TV spots [5].

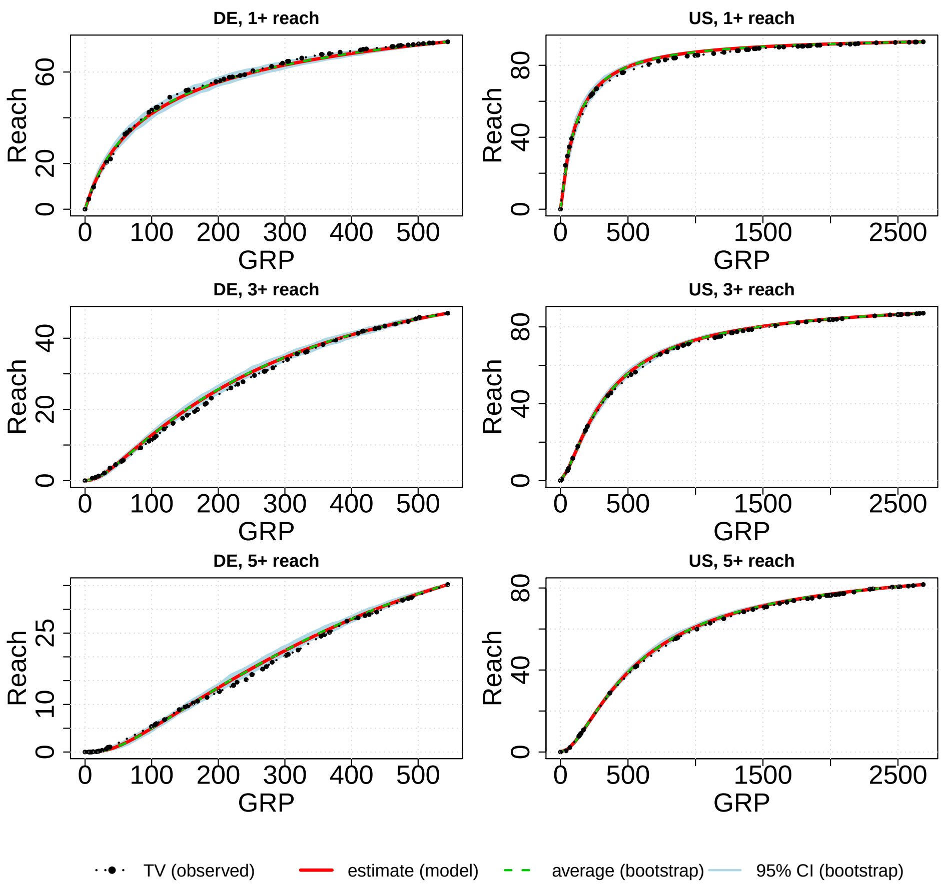

Figure 2: GRP / Reach curves for DE and US for different k+ reach

homepage masthead are discrete ($k = 0, 1, \ldots, k_{\max}$), we fix the number of days, and find the optimal share of the remaining budget (total budget − k · daily masthead cost) to buy YouTube watchpage impressions instead of TV. We then compare the different optima over all $k = 0, 1, \ldots, k_{\max}$, and choose the one with the largest combined reach.[3]

Figure 4 shows combined 3+ reach as a function of TV budget share (each curve represents a different k). The largest combined reach can be achieved by using 17.25% of the total TV budget ($6.9 million) for online advertising instead: optimally, the advertiser should buy 2 homepage mastheads (at a cost of $500,000 per day) and use the remaining online budget to show 268 million YouTube watchpage impressions.

[3] We limit the maximum possible number of HP masthead buys to $k_{\max} = 3$ days (out of 77 days).

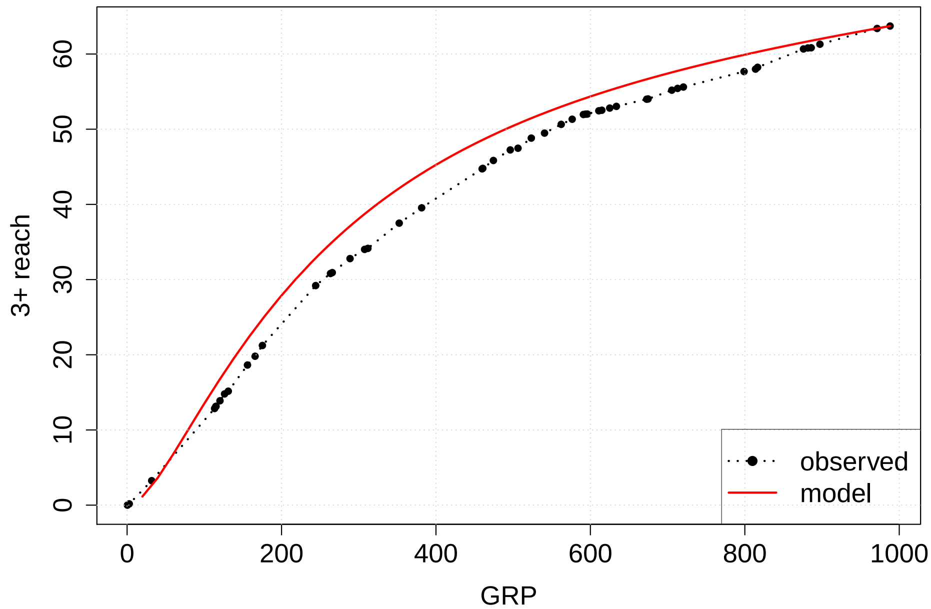

Figure 3: TV GRPs vs 3+ Reach based on Nielsen data.

This mix yields a 67.51% combined reach, compared to 63.72% TV-only reach. The incremental reach of 3.79% is a combination of two opposing effects:

-3.41% TV reach ($63.72\% \rightarrow 60.31\%$) due to budget cuts, but

$+7.2\%$ gain in YT-only reach.

To understand this trade-off it is worthwhile to re-examine Figure 3: since TV reach curve shows diminishing returns (slope is decreasing as GRPs increase), cutting budget does not reduce TV reach too much, but YouTube can cost-effectively reach new (light TV) viewers. Recall again that this increased reach can be obtained at *no* additional cost.

We also note that we use demographic information only to estimate typical viewing behavior of the target demo. We do not use it for demographic targeting of YouTube campaigns. If an advertiser uses demo-targeted ads online, the optimal shift and the attained extra reach will be, in general, higher.

4.1.1 Results by TV Viewing and Age Groups

By definition, TV-only campaigns have a hard time reaching light TV viewers, while spending a large amount of money on serving ads to already reached heavy

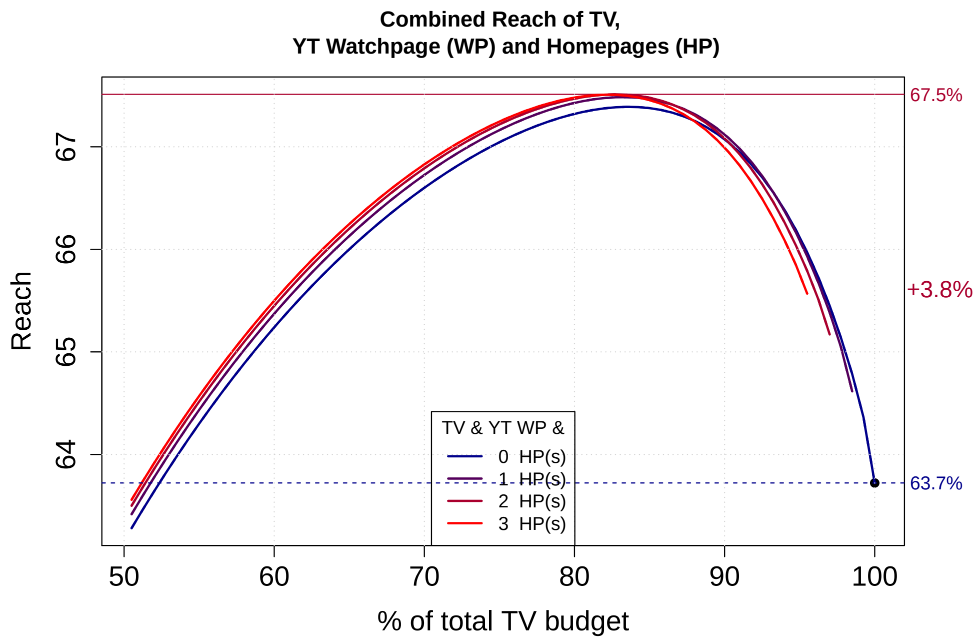

Figure 4: Combined 3+ reach of TV, YouTube WP & HP as a function of TV share (100 % means TV-only). Moving to the left puts more ads to YouTube and increases combined reach.

TV viewers (i.e., adding more frequency). YouTube, on the other hand, can efficiently reaching those light TV viewers, while reducing average frequency of heavy TV-viewers.

To understand why adding YouTube can increase the combined reach of a campaign, it is useful to study effects of the media mix across different TV viewing time groups. Here we divide TV viewing time in 3 equally sized groups (3 quantiles). The panel-estimates among males 18-49, for the group cutoffs are 1.47 and 3.49 hours/day, with group averages of 0.6, 2.4 and 6.1 hours/day.

Figure 5a shows how frequency, GRPs, and reach change across the 3 buckets for 3 different scenarios (differently colored bars). The first bar corresponds to the TV-only plan; the second to the optimal plan (use 17.25% of TV budget to buy 2 days of homepage mastheads and rest for YouTube watchpage impressions), and the third bar shows the optimal plan plus 10% additional budget shift. Comparing reach and frequency for each bar confirms that the optimal media-mix plan increases reach for light TV viewers and reduces frequency for heavy TV viewers. For example, the optimal media-mix plan efficiently reaches light TV viewers (TV-only 22.84%; with YouTube 33.08%), while it can reduce average frequency of heavy TV-viewers (from

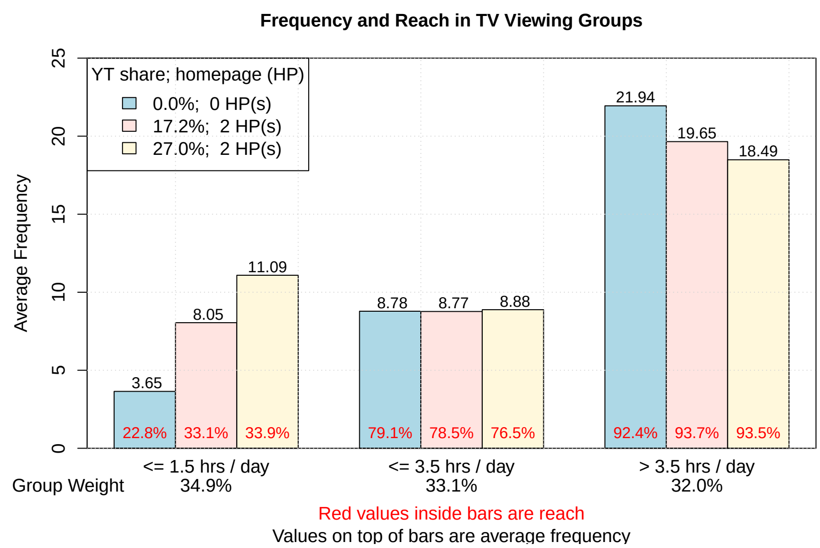

(a) TV viewing times

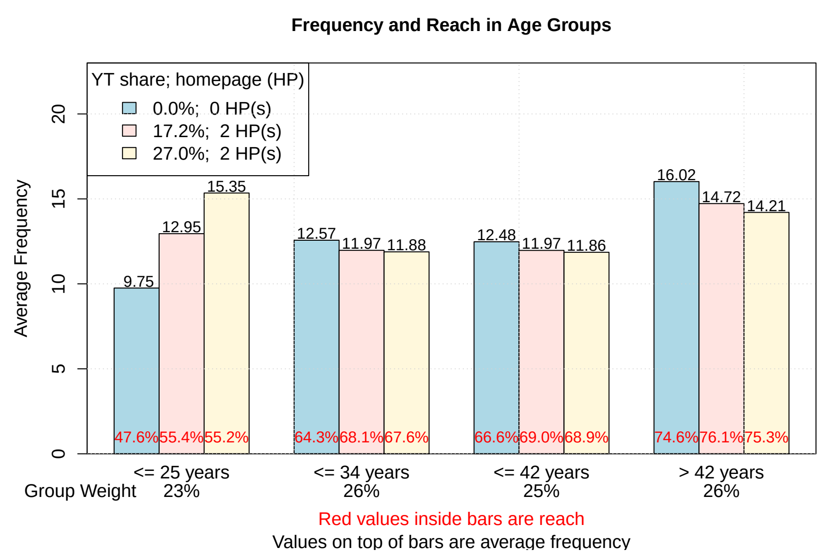

(b) Age groups

Figure 5: Combined reach and frequency across a) TV viewing groups and b) age groups for three budget shift scenarios.

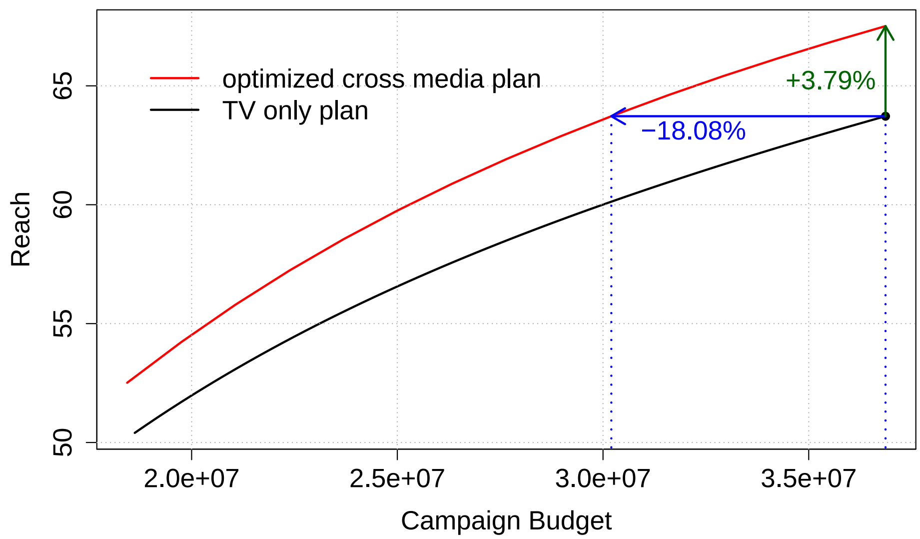

Figure 6: Shareshift optimization results for both scenarios. Scenario I (green arrow): obtain 3.79% incremental reach at constant budget (black dot); scenario II (blue arrow): save 18.08% of the total budget while maintaining the TV-only reach (black dot).

21.94 to 19.65).

Splitting by age (3 equally sized buckets) reveals major differences between TV and YouTube (Figure 5b). The TV-only plan (blue bars) is heavily skewed towards older viewers – both in reach and frequency. An optimized media-mix plan on the other hand (red bars) can substantially increase reach and frequency of the young targets, while lowering them only by a bit for older viewers.

4.2 Scenario II: Minimize Cost While Maintaining 3+ Reach

Now we consider the scenario where an advertiser wants to maintain the reach while cutting costs.

Figure 6 shows that in this case the optimal strategy is to first cut the total disposable budget by 18.08%. If campaign would only run on TV, a lower budget would reduce the campaign's reach. However, using 17.38% of the remaining budget for the more effective online advertising (2 days of homepage mastheads plus rest on watchpage impressions), the advertiser can maintain the original TV-only reach of 63.72%. Note that this leads to cost savings of $7.23 million for the advertiser.

5 Conclusion and Extensions

The models presented here can be used to show how a past TV campaign could have been improved by shifting budget from TV to online ads. They are based on individual-level media consumption and ad exposure collected from cross media panels. For each user in the panel in the target audience, the models estimate the number of ad impressions on TV and online in the hypothetical cross media campaign after the budget shift. These individual probabilities of being reached at least k times are then aggregated to generate an estimate of the $k+$ reach of the cross media campaign. The models are validated with simulations and historical online campaigns. Furthermore, they are illustrated on a case study to show the optimized media mix can obtain a higher reach at the same cost or maintain the same reach at a lower cost than the original TV campaign. They can be used to provide valuable guidelines for future media planning.

A very common question in cross-media planning is whether ad exposures on different platforms or in different ad formats are equivalent in effectiveness. Some research has been done on this topic, such as [6] and [7]. In Section 2.2, we assume ads on TV are of the same effectiveness as those on YouTube watchpages and mastheads. If the advertiser has an opinion on the ad effectiveness of TV vs online, we can calculate the effective ad frequency by modifying f_i into

$$f_i = e_t f_{i,T} + e_w f_{i,W} + e_h f_{i,H}, \tag{13}$$

where e_t, e_w and e_h are the ad effectiveness multipliers of TV, YouTube watchpage and masthead. The rest of the models still apply.

The optimal cross-media plan is obtained by maximizing reach in the target demographics. Some advertisers may be concerned with the possibility of losing GRPs and thus the share-of-voice in the market when shifting budget to online. To avoid a deep cut on GRPs, we can constrain the loss of combined GRPs compared to the original TV plan to be within a certain threshold, such as 15%, in the optimization step. In this way the optimal cross-media would deliver a possibly higher reach and maintain a reasonable level of GRPs.

Acknowledgments

Special thanks go to Christoph Best, Simon Morris and Sheethal Shobowale for productionizing the models and running the analysis on thousands of historical campaigns. We would also like to thank Tony Fagan and Penny Chu for their encouragement and support, Raimundo Mirisola, Andras Orban, Jim Stewart, Sergio Sancho, Joris Merks, Vanessa Bohn, Elissa Lee, Daniel Meyer, Taylan Yildiz, Simon Row, Harry Case and Jim Dravillas for their insightful discussion and constructive feedback.

References

[1] Yuxue Jin, Sheethal Shobowale, Jim Koehler, and Harry Case. The incremental reach and cost efficiency of online video ads over tv ads. Technical report, Google Inc, 2012. URL http://research.google.com/pubs/pub40426.html.

[2] YouTube Advertise. YouTube Solutions for Advertisers, 2013. URL http://www.youtube.com/yt/advertise/our-solutions.html.

[3] GfK Consumer Panels. Media Efficiency Panel, 2013. URL http://www.gfk.com/solutions/consumer-panels/Pages/media-efficiency-panel.aspx.

[4] Ebiquity. Portfolio plus. URL http://us.ebiquity.com/portfolio-plus/.

[5] Nielsen Solutions. Nielsen Presents: Cross-Platform Home Panels and Extended Screen Ratings, 2013. URL http://nielsen.com/us/en/nielsen-solutions/nielsen-measurement/nielsen-tv-measurement.html.

[6] Google Think Insights. YouTube and TV are Better Together Infographic, 2012. URL http://www.google.com/think/infographics/youtube-tv-better-together-infographic.html.

[7] D. Varan, J. Murphy, C. Hofacker, J. Robinson, R. Potter, and S. Bellman. Cross-device synergy versus cross-media synergy, 2012. URL http://researchrepository.murdoch.edu.au/9153/.