

Beyond the Pair: Media Archetypes and Complex Channel Synergies in Advertising

J. Jason Bell

Saïd Business School
University of Oxford
jason.bell@sbs.ox.ac.uk

Felipe Thomaz

Saïd Business School
University of Oxford
felipe.thomaz@sbs.ox.ac.uk

Andrew T. Stephen

Saïd Business School
University of Oxford
andrew.stephen@sbs.ox.ac.uk

Beyond the Pair: Media Archetypes and Complex Channel Synergies in Advertising

Abstract

Prior research on advertising media mixes has mostly focused on single channels (e.g., television), pairwise cross-elasticities, or budget optimization within single campaigns. This is detached from practice where (i) marketers decide between an increasingly large number of media channels, (ii) media plans involve complex combinations of channels, and (iii) marketers manage complementarities among many channels. In this work we use Latent Class analysis to uncover tendencies in media allocations. Latent classes account for non-random selection of channels into campaigns, capture pairwise and higher-order interactions between channels, and allow for meaningful interpretation. We describe the most common media channel archetypes and estimate their relationship to the effectiveness of a set advertising campaigns on common brand-related performance metrics. We use a dataset of 1,083 advertising campaigns from around the world run between 2008 and 2019. We find no single media mix that consistently correlates with high performance across all metrics, but clear high-performing patterns emerge for specific metrics. We find that traditional channels (TV, outdoor) often appear together with digital channels (Facebook, YouTube) in high-performing campaigns. Additionally, current marketing practices appear suboptimal, with simple strategies predicted to improve lifts by 50% or more.

Keywords:

advertising, media planning, media complementarity, budget allocation, leading indicators

INTRODUCTION

As global advertising spend exceeded \$1T ([Statista 2024](#)) in 2023, marketing academics and practitioners need to understand effective media budget allocation approaches. Media selection is also important for recipients of media budgets in the vast and complex market for advertising. For instance Google reported \$224.47B in revenue from advertising sales in 2022 ([Inc. 2022](#)), while the entire global print publisher ad revenue was \$37.3B, just behind Amazon’s advertising sales of \$37.7B for that period ([Tobitt 2023](#)). Despite an often discussed shift to digital channels, traditional advertising products are hardly abandoned, with \$46.9B flowing to outdoor ([MarketWatch 2023](#)) and \$37.2B to radio advertising globally in 2022 ([IMARC 2023](#)).

Unsurprisingly, a \$1T market attracts entrants, competition, and innovation. As a result, marketers are met with a proliferation of media channels (e.g. the rise of retailer owned ad channels like sponsored listings on Amazon or Walmart), as well as an increase in channel-specific options to consider when developing media mixes. Marketers buying ads and media planners are faced with a complex resource allocation problem that is seemingly intractable with a vast number of possible media combinations from an expanding menu. The task is made more difficult by the fact that media planners cannot immediately evaluate the results of their choices, as “feedback” in the form of market results comes with both statistical noise and often a long delay. In addition, media planners must often decide on budget allocations in advance, and adapt to changing market conditions later. Practically, marketers must simplify media allocation. This simplification comes in the form of using heuristics, repeating previous patterns, imitating competitors, following industry trends, or outsourcing to the large support industry that offers help with media planning (e.g. marketing mix modeling) or validation (e.g. attribution).

Simplification introduces errors that, from the perspective of a manager, are hard to detect due to limited data relative to the high dimensionality of the problem. Unlike adjust-

ments to individual ads, A/B testing budget allocations on entire campaigns carries a high ratio of cost to information yield. It is worth recognizing that large firms are not in the habit of running experiments with multi-million campaign investments against their major brands, and tend to avoid risky experimentation. Furthermore, information about channel-specific return on ad spend usually comes from channel owners, and is often analyzed in isolation from concurrent channel spend decisions, which reduces its value for questions of channel interaction. In some ways, the overwhelming uncertainty provides “cover” for managers to justify approaches that are favorable to them for reasons other than performance. If two strategies plausibly produce similar results, managers are likely to prefer the one that is simpler to execute and explain. For example, buying ads based primarily on reach is one strategy which may appeal to managers facing very high uncertainty, particularly as it aligns with industry lay beliefs on reach-sufficiency (all results flow from achieving high reach) and the ease of measuring reach. Regardless of the criteria used, managers are likely to view the media allocation task holistically and intuitively. This contrasts with the notion of managers methodically simulating all possible channel combinations for every campaign and goal. The combinatorial nature of the problem makes such exhaustive testing time and cost prohibitive.

From the perspective of a researcher, differing budget allocation approaches taken by managers creates variation across firms and industries in media mix tendencies, which could potentially be exploited to uncover patterns in channel allocations and outcomes. We take advantage of the variation in campaign construction (channel usage and spend shares) using a unique dataset that includes 1,083 large multi-channel campaigns that ran between 2008 to 2019. These campaigns were undertaken by 557 global brands, and covered 23 major industries in 51 countries. The dataset is broad enough that it could be used, in principle, to detect patterns in campaign construction, and to estimate the relationship between different patterns and advertising outcomes. However, the data alone are insufficient because patterns do not make themselves readily available in the raw data. To see this, consider two campaigns with similar spend from similar brands. The first allocates 60% to TV and 40% to Youtube,

the second allocates 60% to TV, 30% to Youtube, and 10% to cinema ads. These campaigns are very similar to one another, but not exactly the same. To succeed in extracting patterns from the data, we need a model that can represent these two campaigns as “close” but not identical. We should avoid any model that would treat the two as totally distinct entities. At the same time, we should also avoid a model which reduces them to simply the sum of their parts, because that would fail to capture complementarities between the channels.

Our model relies on the concept of archetypes. An archetype is like a “Platonic ideal” of a campaign style. For example, a manager might imagine a campaign using digital channels with algorithmic targeting (e.g., Youtube). However, in practice, the manager might deviate slightly from the envisioned approach and include some spend on radio (neither digital nor algorithmically targeted). We could say the campaign turned out to be close to the “algorithmic digital” archetype, though not a perfect realization of it. From this perspective, archetypes are landmarks on a map of media allocation styles, and individual campaigns are characterized by their coordinates on the map, and thus their relative distances to different archetypes. Campaigns can be concisely described by projecting them into archetype space, and “close” campaigns can be treated as similar but not identical. Past research on media complementarity has been limited to main effects of channels plus pairwise interactions. This ignores interactions between more than two channels, and lacks parsimony. The lack of parsimony is a serious concern when there are more than three or four channels. For example, with eleven channels, there are 2047 possible channel combinations. Representing each of these combinations as an additional term in a regression is hopelessly infeasible.

To apply this approach we must be able to estimate archetypes from data, which requires that patterns in media allocation actually exist, and that we have a feasible estimation procedure for recovering them. Whether the patterns exist is an empirical question which we address in this work, using Bayesian latent class analysis as an estimation procedure. We combine the archetypes from the latent class analysis with a model predicting brand lifts to understand the relationship between archetypes and outcomes. The breadth of our

dataset is critical for estimating archetypes with high external validity. To make archetypes interpretable, we build an empirical Archetype Map, similar to the conceptual one described above, and display it in two dimensions. A given campaign on this map can be thought of as a mixture of the different more “pure” archetypes, and the empirical archetypes are given as landmarks.

We make three contributions. First, our dataset and archetype model allow us to come closer than prior research to addressing the critical practical question of how to best combine media channels for high-performance against different advertising goals. By going beyond simple pairwise combinations of media channels we are able to see how more complex combinations of channels have tended to perform, which is new to the marketing literature.

Second, our empirical study is, to the best of our knowledge, the largest academic study of advertising performance and media mixes. Our unique dataset includes 1,083 large multi-channel campaigns that took place during the period 2008 to 2019. These campaigns were undertaken by 557 global brands, covered 23 major industries, and ran in 51 countries. Finally our dataset allows us to consider several brand-related outcomes that are of strategic interest to marketers (e.g., brand awareness, associations, purchase intent). While the existing literature has been able to examine the implications of a few media channels, we focus on many simultaneous media channels.

Third, we produce several descriptive findings that strongly suggest the existence of media allocation tendencies among managers. We find that reach, which is relatively easy to understand and measure, is optimized to a greater degree than lift (i.e., improvement on brand metrics). We find seven media allocation archetypes that each relate differently to brand outcomes, implying that managers cannot simply apply a single archetype without respect to campaign goals. The archetypes have different risk profiles, with some archetypes displaying strong patterns of high average lifts but with high variation, while others associated with lower average lifts but much better worst-case outcomes. We also find that budget concentration (how unevenly managers allocate) relates favorably to lifts in Aided Awareness

but unfavorably to Association lifts. All of our substantive results are new to the literature.

LITERATURE REVIEW AND CONCEPTUALIZATION

Deighton (1996, p.243) ponders “Why would one not integrate a communications campaign? Integration is surely the natural condition, the simple and obvious way of doing things. Most marketing problems are solved with a combination of tools, even if the blending is no more complex than synchronizing advertising with trade promotion schedules. Unintegrated communication campaigns are aberrations.” This perspective highlights that the combination of marketing tools is a longstanding part of marketing thinking, captured in both the concept of the marketing mix (Culliton et al. 1948; Borden 1964) as well as the Marketing Concept itself (Houston 1986). Therefore, it is unsurprising that the combinatorial aspect of integration is connected to critical facets of marketing, such as budget allocation (Venkatesan and Kumar 2004; Fischer, Albers, Wagner, and Frie 2011), marketing mix and media planning (Wedel and Kannan 2016), attribution (Berman 2018), as well as advertising effectiveness (Danaher, Danaher, Smith, and Loaiza-Maya 2020) and the financial impacts of marketing (McAlister et al. 2016). Since marketing in practice has combinatorial or “mixing” principles at its core, it is also unsurprising that much of the literature on media effectiveness and related decisions has considered how different combinations of activities work. The IMC literature has focused extensively on this; see reviews by Naik and Peters (2009), and Batra and Keller (2016). These works lay out a framework of components for successful marketing communication integration: (i) consistency, (ii) cross-effects, and (iii) complementarity.

Interestingly, the concept of consistency of communication across media channels is deemed a “solved problem” on the basis of numerous studies supporting its advantage (Duncan and Moriarty 1998; Batra and Keller 2016), or placing it as a tactical part of integration (Naik and Peters 2009), thus requiring no further investigation. However, these perspectives also mark a significant distinction between the three components of IMC. While the complementarity and cross-effects of channels relate to their inclusion in a media plan, the

consistency of communications relates to content development and management that occurs alongside media channel selection and planning. Therefore, it is in the interplay of complementarity and cross-channel effects, or what [Naik and Peters \(2009\)](#) term strategic synergy, where combinatorial aspects of media planning are most relevant.

The questions in this type of research are relatively straightforward, such as seeking to understand how different media channels operate together in order to see if and when they “play nicely together” (i.e., are synergistic). However, conflating cross-effects and complementarity of media channels, or collapsing both into a single construct (synergy) is not always ideal. Indeed, we argue that in doing so prior research on media selection and integration has been overly concerned with simple media channel combinations (e.g., pairwise combinations that can be captured with cross-effects) at the expense of considering more realistic mixtures of media channels that go beyond what can be empirically specified with pairwise cross-effects in advertising response models.

The importance of this distinction comes from [Batra and Keller’s \(2016\)](#) observation that, while the literature has provided a significant amount of empirical evidence on the existence and importance of media interactions and cross-effects, “it has not yet clearly explicated the relative strengths and weaknesses of different media in influencing different communication outcomes” (p.123). The reason for this gap might be more apparent from the body of literature on media complementarity and cross-effects, a selection of which is shown in [Table ??](#) (in [Web Appendix F](#)). It highlights a strong dedication to explaining advertising impact on sales ([Gopalakrishna and Chatterjee 1992](#); [Naik and Raman 2003](#); [Danaher and Dagger 2013](#); [Srinivasan, Rutz, and Pauwels 2016](#); [Kolsarici and Vakratsas 2018](#)), or analogous metrics like conversion and recruitment ([Gatignon and Hanssens 1987](#)). However, not all media campaigns are designed for activation or sales. Brand-related outcomes are often important to practitioners. Yet the number of studies on brand outcomes is much smaller than those on sales outcomes ([Edell and Keller 1989](#); [Chang and Thorson 2004](#); [Dijkstra, Buijtelts, and Van Raaij 2005](#); [Fischer 2019](#)).

This focus reduces the question of media selection and combinatorial channel effects down to a single advertising outcome. This concentration on one outcome creates a potentially false equivalence between channels, as they are all measured as better or worse delivery mechanisms for one metric; effectively creating a literature-level environment where it is difficult to identify complementarity (hence the reasonable simplification to synergy). If the construct requires “difference” and “differentially useful” as [Batra and Keller \(2016\)](#) suggested, then the second component is only observable when conditions vary (outcomes).

Much of this equivalence is observed in studies concerned with within media plan budget allocation and re-allocation toward an optimal position ([Fischer 2019](#)), where the trade-off towards one specific goal is justified. However, it translates into a concerning industry practice where media planning and purchase is similarly simplified to the bases of reach and frequency (GRPs), rather than a combination of GRPs and the functionality or purpose of the channel. Practitioner guidance encourages this behavior by focusing on campaign reach measures ([Binet and Field 2013](#)). In a way, it is similar to treating an ad exposure as equal across channels, or perfectly fungible, such that an ad exposure on Facebook is identical to an ad view on television or in a newspaper. This assumption goes against the notion of differentiation of media channel exposures ([Batra and Keller 2016](#)), as well as findings showing differential consumer behaviors and the resulting importance in respecting resulting channel differences ([Schweidel and Moe 2014](#)).

Further, prior literature has tended to leverage fairly limited datasets in terms of scale and scope. The majority of prior studies rely on case studies, focusing on a single brand and single campaign. Notable exceptions to this are [Joo et al. \(2014\)](#) and work by [Kolsarici and Vakratsas \(2018\)](#). Similarly, studies have strongly focused on simple dyadic relationships between channels, such as the interaction between sales force and advertising ([Gatignon and Hanssens 1987](#)), television and radio ([Edell and Keller 1989](#)), television and print ([Naik and Raman 2003](#)), television and search ([Joo et al. 2014](#)), or television and social ([Fossen and Schweidel 2019](#)). Even those fewer studies that consider more than two channels simulta-

neously are restricted to the pairwise interactions (Danaher and Dagger 2013; Srinivasan, Rutz, and Pauwels 2016; Kolarici and Vakratsas 2018; Fischer 2019) or have aggregated data on multiple channels of the same category (e.g., traditional media and social media) to allow for pairwise cross-effects estimations (Stephen and Galak 2012). In fact, Danaher and Dagger (2013, p.530) specifically describe the inability to demonstrate synergies among the channels in their study, as the pairwise interactions among all ten channels became an econometric problem. And in that vein, Batra and Keller’s (2016) observed lack in guidance on the appropriate selection and complementarity in media channels for a diverse set of outcomes is not necessarily due to a lack of interest, or motivation; but rather from a simple lack of opportunity to observe a sufficient amount of brands, campaigns, and outcomes in order to make meaningful statements.

DATA

Our dataset was provided by Kantar. All campaigns in our dataset were generated by Kantar’s Insights Division in their CrossMedia Research practice, which is dedicated to global market research on media campaign performance, advertising effectiveness, and brand impact. The dataset includes measures captured in brand lift and advertising effectiveness studies for 1,083 large multimedia advertising campaigns. The earliest campaign in the dataset ran in 2008, and the most recent was in 2019. A total of 557 global brands are represented in these campaigns, with 23 different industries covered (e.g. automotive, financial services, retail). Campaigns come from 51 different countries, capturing all major markets and regions (e.g. North America, Latin America, Europe, Asia Pacific). A noteworthy feature of this dataset is that all campaigns are significant in their budgets and size, meaning that our sample is focused on large, professional advertising campaigns. For instance, the average media-only spend (i.e. ignoring production, management, and ancillary costs) of the U.S. campaigns in our dataset is USD 12 Million.¹ Kantar, on behalf of their global

¹The dataset does not have comparable figures across countries.

brand clients, measured media budget allocation for each campaign across eleven types of advertising channels: (i) Cinema, (ii) Facebook, (iii) Magazines, (iv) Newspapers, (v) Online Display (banners), (vi) Online Video, (vii) Outdoor, (viii) Point of Sale, (ix) Radio, (x) Television, and (xi) YouTube.

Table 1: Advertising Campaign Dataset Overview

Total number of campaigns	1,083	
Mean Number of campaigns per year	92.1	(60.8)
Average Spend (U.S. campaigns only):	\$12M	
Total number of brands:	557	
Mean Lift (difference between percentages)		
Motivation	2%	(2%)
Association	3%	(3%)
Unaided Awareness	2%	(3%)
Aided Awareness	1%	(4%)
Number of Channels Used: Distribution		
1 channel	6%	
2 channels	13%	
3 channels	26%	
4 channels	25%	
5 channels	19%	
6+ channels	12%	
Maximum Number of Channels Used	9	(of 11)

Notes: Standard deviations in parentheses.

Each observation in the dataset is an advertising campaign and an associated Kantar study. As the studies were conducted in real-time against live campaigns across a varied number of media channels, true random-assigned control was impossible. The approach Kantar used is standard industry practice for multi-channel campaigns across digital and non-digital channels. In this situation true exposure (or non-exposure) to a given campaign cannot be fully controlled in channels such as television and outdoor like it can be in digital channels such as YouTube and Facebook. In each campaign Kantar controlled for background variation by surveying non-exposed people both before the campaign begins and throughout the campaign. Non-exposure during the time of a campaign is estimated by Kantar using




Figure 1: Prevalence of Unique Channel Combinations

media consumption self-reports from survey respondents. This is a typical “opportunity to see” approach in the industry. For example, if a respondent reports that she does not use Facebook, she is eligible to be in the control group for a Facebook advertising campaign. Kantar also directly asked respondents about seeing each campaign. Kantar additionally accounted for the “predisposition” of a respondent towards each brand with information about exposure to previous campaigns, use or awareness of the brand by family members, brand equity in the survey region, and similar measures to determine a lift attributable to the specific campaign alone.

Our dataset illustrates the complexity of the media allocation task. There are a total of 332 unique combinations of channels used in the dataset. The most common six combinations together only account for around 16% of all campaigns. The complexity of the media allocation landscape is such that simply enumerating the most common combinations is not a feasible way to discover generalizable patterns in the data. Figure 1 plots each unique channel combination in the data, ordered by frequency of usage, with the y-axis showing the percentage of all campaigns each accounts for.

In terms of campaign performance metrics, all of these campaigns had brand-related objectives and therefore the studies that Kantar ran for their clients were all “brand lift”

studies. Four different brand-related outcomes were measured for each campaign: (i) Unaided Awareness, (ii) Aided Awareness, (iii) Association, and (iv) Motivation. The dataset does not include sales performance data, since the measurement work Kantar did on behalf of their clients focused solely on brand outcomes. Of course, it is entirely possible that advertisers had a desire to impact sales as well as brand outcomes. Sales data, however, is not part of this proprietary Kantar dataset and therefore is not available. We have four dependent variables, which are all lifts (differences between proportions). Each question used to measure the dependent variables is coded so that a value of 1 indicates a desirable response for the brand. Lift is then a proportion, and lies between 0 and 1. A lift of 1, the maximum, means that every respondent in the exposed group answered affirmatively and all control respondents answered negatively. Table 1 summarizes key features of the dataset. Figure 2 shows histograms of our four outcomes. Note that each is greater than or equal to zero, with the exception of one observation, which we drop. The histograms make clear two modeling challenges. First, there is a large mass of observations at zero, for three of the four outcomes. We must account for this zero-inflation. Second, the data distributions have relatively long right tails.

Figure 3 shows that after a mild decrease in the number of channels per campaign from 2008 to 2013, the number of channels began increasing to reach a high in 2019. As more new media channels and formats appear, complexity is likely to continue increasing.

Finally, we provide some model-free evidence that brands have idiosyncratic tendencies in media selection, which suggests we should be able to identify archetypes in the dataset. For eleven brands most commonly appearing in our dataset, we plot the average share allocated to each channel, as well as the usage probability of each channel, in Figures 4 and 5. It is evident that whereas brands exhibit some idiosyncrasies in share of spend across channels, there is much more variation when looking at campaigns in terms of *usage* (a binary indicator of zero vs. non-zero spend). This supports our choice of usage as the input into the archetype model, as discussed below.

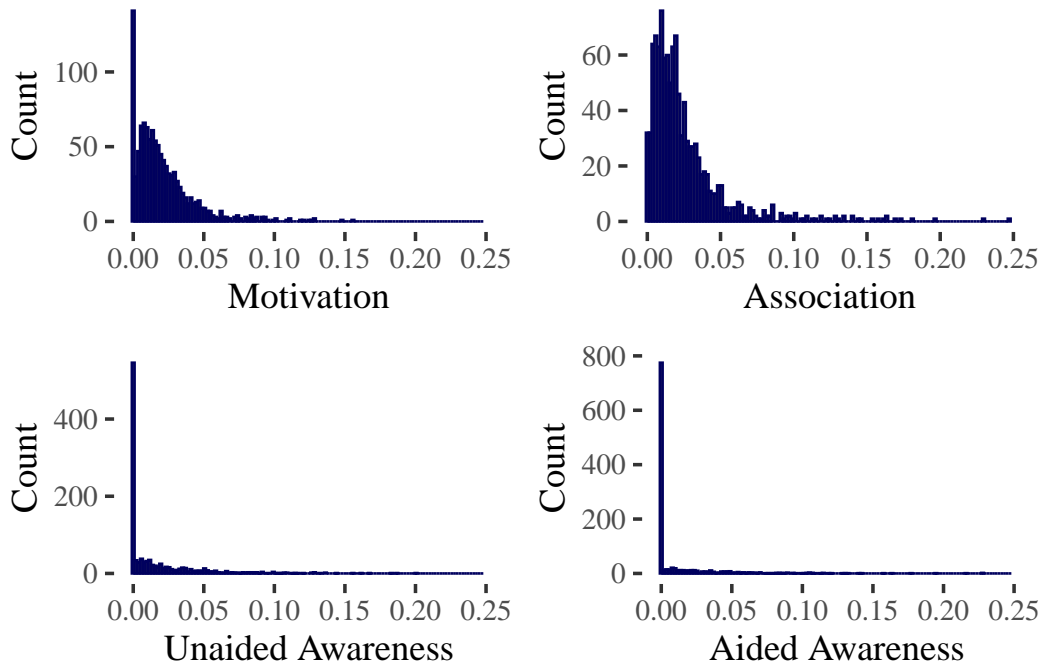


Figure 2: Outcome Distributions are Positive and Right Skewed




Figure 3: Channel Consolidation then Expansion

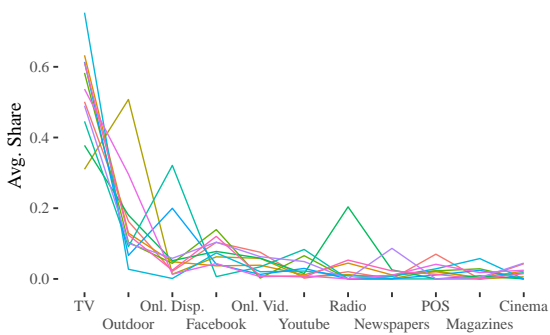


Figure 4: Patterns in Share

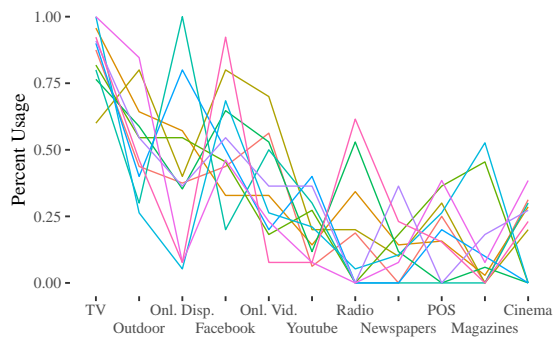


Figure 5: Patterns in Usage

Measures

Keller (1993) suggests that marketers should track changes to different dimensions of brand knowledge over time, and importantly, how these changes might be related to different marketing mix choices. The dependent variables in our study follow this pattern, and in a structure consistent with Keller’s brand equity framework. Each of the following four outcome measures is described in terms of its campaign lift, or the (Kantar-estimated) campaign-only attributable change in the metric as measured by the differences between the response of pre-campaign and during-campaign consumers in exposed versus non-exposed matched control groups. These campaign lifts can theoretically take values between -1 and 1, but in practice are almost exclusively non-negative. Lastly, as all brand outcomes are measured for all campaigns, regardless of goal, and then passed through a significance filter, the resulting data include a number of exact zeros (i.e. non-significant quasi-experimental lifts). We must appropriately handle these exact zeros in the modeling, and we do so with a hurdle model. The four outcome measures we use are defined below.

Unaided awareness measures knowledge and memory of the brand without direct cues or prompts. While we are not permitted to list the exact questions used on the Kantar survey, we can say that the question used by Kantar is very similar to many used across the industry. The question usually takes the form “which brands in [category] are you familiar with?”

Aided awareness measures recognition of the brand with the aid of a prompt or a direct cue. Like the question used to measure unaided awareness, Kantar follows typical industry practices. Most questions in the industry are of the form “which of the following brands in [category] are you familiar with?”, followed by a list of brands to select from, where respondents are allowed to indicate multiple brands.

Association is related to the second dimension of the Keller (1993) brand knowledge framework, the brand image. Association captures the presence and strength of associative connections to the brand, regardless of favorability. Explicitly, it considers whether respondents can connect a particular message, positioning statement, slogan or other stimuli used

in advertisements to the focal brand. The questions used to measure this are often dependent on the category or even the specific brand. Some very general associations such as “high quality” or “trustworthiness” can be measured for a wide variety of brands and categories. Two examples of this type of question have the form “which of these brands is the most [association]” or “which of these brands has the slogan: [slogan]”

Motivation is an index of brand metrics meant to proxy for near-term sales. The construction of the exact metric is proprietary, but it is highly correlated with Purchase Intent (which is one of its constituent metrics), so we interpret findings as if Motivation is essentially Purchase Intent. Of the four metrics, Motivation is the most like a sales measure, thus providing the closest (or at times identical) measure to the existing body of work on cross-media effects described in Table 1. We then make use of both the subset of the eleven channels selected (binary indicators of positive spend) as well as the share of campaign budget spent on each channel as the focal point of our analysis, and in the empirical study of Archetypes (or clusters) that we describe next. We also control for time (year) and spend concentration (defined below), and we use the brand and category (industry) as control variables in a random intercepts specification.

METHODS

Descriptive Stochastic Frontier Analysis

In our initial discussion, we highlighted that practitioners are concerned with mindset metrics consistent with theory and empirics (Vakratsas and Ambler 1999; Srinivasan, Vanhuele, and Pauwels 2010). As such, it would be logical to assume that managers are therefore optimizing their media selection towards this declared goal. However, we also described a practitioner focus and guidance on “reach maximization and sufficiency,” which carries with it the implication that channel complementarities are inconsequential, and ad exposure are at least somewhat fungible across media channels in the construction of an effective media campaign seeking to achieve these declared intermediate brand goals (e.g. awareness,

associations, etc).

Conversely, we put forward that firms (and academics) are missing the complementarity aspect of planning and focusing on coarse audience maximization strategies, rather than goal and outcome-oriented approaches. So, in order to interrogate these assumptions, we will rely on a “model-light” description of our data via Stochastic Frontier Analysis (SFA). This approach is used to describe the efficiency of resource allocation for the creation of a given output. It estimates an efficient frontier, and ranks the inefficiency of observed resource allocation decisions according to their proximity to this estimated theoretical line (see [Krasnikov, Jayachandran, and Kumar 2009](#) for a detailed description of the approach). In our case, we consider channels as inputs and lifts in any of the above mentioned brand mindset metrics as the outcome. And if the assumption holds that managers’ strategic choices are correct and aligned with their stated and demonstrated intent, then campaigns should be efficient in delivering brand outcomes, with an expected large mass of campaigns relatively close to the most efficient observed campaigns. Columns two and three of [Table 2](#) shows the distribution of campaigns under the efficient frontier by distance, along with their associated mean observed lifts (outcome).

Table 2: SFA Mindset Metric Lift

Decile to Efficient Frontier	Lift SFA		Reach SFA	
	% of Campaigns	Mean Lift	% of Campaigns	Mean Reach
0.0	3.72%	0.13%	–	–
0.1	8.81%	0.42%	–	–
0.2	11.89%	0.74%	0.62%	24.64%
0.3	13.89%	1.18%	0.62%	33.07%
0.4	17.69%	1.75%	2.07%	45.14%
0.5	18.23%	2.61%	1.65%	55.50%
0.6	17.08%	4.16%	5.99%	65.65%
0.7	7.60%	7.97%	8.88%	74.77%
0.8	1.08%	18.46%	20.66%	85.42%
0.9	–	–	59.50%	95.23%
@ Frontier	–	–	–	–

Notes: Reach is stated as % of available audience exposed to the ad.

First, we must recognize that there is no *ex ante* expectation that campaigns would reach the efficient frontier, but the best performers do highlight what is practically achievable. However, it is rather surprising that so few campaigns exist at the top of the effectiveness ladder, and that the disparity of average outcomes is so large. Furthermore, efficiency/in-efficiency with respect to brand metrics appears to be approximately normally distributed around the mid-efficiency point. This result provides some evidence against the assumption that managers and media planners are already making optimal choices, despite the practical guidance that 60% of their advertising should move these metrics.²

Next, we consider a different construction of the stochastic frontier where the intended outcome is campaign reach. As noted previously, practitioners are also given the guidance to maximize reach, which carries the risk of downgrading or entirely dismissing complementarities, as planners can opt for channels with larger audiences, and continuously pay more to reach increasingly greater proportions of them, with no consideration for any mechanical or functional differences between channels and their consequences for the campaign goal. Columns four and five of Table 2 show these results.

This second frontier shows a markedly different pattern. With respect to reach, 80% of campaigns fall near the two highest reach-efficiency scores. Compared to 8% of campaigns on the two highest efficiency bands for brand metrics (again, metrics that are important enough for brands to pay an agency to measure their shift as a result of multi-million dollar campaigns). It is clear that optimization is happening against reach, but not campaign goal. Managers are indeed following industry practice and guidance by organizing campaigns around achieving highest reach. Yet, the disconnect between both frontiers suggests that managers' belief of reach sufficiency is incorrect, otherwise brand metrics would similarly improve with reach. Lastly, as the reach sufficiency (and channel/impression fungibility) argument stood in direct opposition to the concept of complementarities, we take this as

²At the request of an anonymous reviewer, we reconsider the SFA with the 10 largest campaign lifts removed to examine the influence of extreme outlier performances. The resulting SFA from this restricted set maintains the same distribution of results, again peaking at the 80% proximity to the efficient frontier, but offering an average 15.23% lift in outcomes at this level.

some initial evidence supporting our argument that complementarity is a crucial component of planning media mixes. We next we turn to the main analysis of complementarity effects.

Main Analysis

Our modeling approach must account for several important features of the dataset. These features are: (i) complex interactions between channels through campaign types, (ii) the skewed and bounded shape of the outcome distributions, and (iii) the presence of a group of exact zeros on each of the outcome variables. We handle the first using a Bayesian Latent Class approach, and the second two with a Gamma hurdle regression model.

Campaign Types and Channel Interactions

Including higher-order interactions comes with a combinatorial challenge. With eleven media channels in our data, specifying every possible pair of channels, triple of channels, and so on, up to all eleven, would require $2^{11} - 1 = 2047$ interaction terms. In the cross-effects research tradition, a typical solution to this combinatorial explosion is to limit the analysis to pairwise interactions between media channels. Limiting interactions to pairs would reduce the number of combinations to 55. However, even this number of combinations is intractable for modeling and interpretation, and would result in the same issues previously documented by [Danaher and Dagger \(2013\)](#). For the sake of illustration, we report our attempt at this approach using a regression approach similar to what is found in much of the media cross-effects literature in Web Appendix D.

Instead of blindly or exhaustively including interaction terms, which is doomed to fail, we took a practical approach. We are interested in how different campaign types compare across the outcome measures. This means that we should only be concerned with the combinations of channels that actually appear in data (i.e., a vast number of the theoretically possible combinations in fact are never observed). Further, we are interested in comparing the effectiveness among “typical” media plans. This type of comparison is distinct from try-

ing to statistically estimate interactions between individual channels to ascertain the extent to which they operate in synergy.

Taking this logic into account, our modeling approach empirically finds common media plan “archetypes” across all of the campaigns in our data. These archetypes embody different media channel combinations. We used Bayesian Latent Class analysis to find the archetypes, which allows us to determine the most typical and therefore practically relevant combinations. Latent Class analysis finds campaign archetypes, and each individual campaign is assumed to be some mixture of the archetypes. The degree of each archetype is given by membership probabilities which are estimated for each campaign. Latent Classes can then be used for clustering by assigning each campaign to the archetype for which it has the highest membership probability. Using Latent Classes, we limit the number of combinations to consider, and the archetypes are interpretable and relate to theory. Finally, by estimating the archetypes from actual campaign spends, we ensure that the combinations we study are substantially present in the data. At the same time Latent Classes exhibit channel combinations of more than two, so that we are not confined to dyadic relationships.

One alternative to Latent Class analysis is a neural network. This would allow us to handle a high degree of complexity and determine combinations empirically, but neural networks are notoriously difficult to interpret. Latent classes or archetypes, on the other hand, are easily interpretable and not only do they aid in the main task of understanding effectiveness, the patterns of channel usage they uncover are findings in their own right.

Bayesian Latent Class Model

We estimate the latent class model twice, once as an exploratory exercise to determine the best number of latent classes, and a second time in a *joint* estimation with the regression. The outcomes or *dependent variables* in the latent class model are indicators of which channels were used on a given campaign, which will later become *independent variables* in the Gamma hurdle model. In the exploratory latent class model — estimated by itself — mem-

bership probabilities and class parameters are estimated. In the second latent class model — integrated with the regression — the membership probabilities and class parameters are estimated, just as before. However, in the integrated version, the membership probabilities for each campaign are included as covariates in the regression. These parameters will reflect the differential performance associated with each archetype.

To explain the Latent Class model, we start with the likelihood contribution for a single observation. A single observation is a vector of channel usage indicators. Let $n \in \{1, \dots, N\}$ be the campaign index, and $c \in \{1, \dots, C\}$ be the channel index, and s_{nc} is the share of spend allocated to channel c in campaign n . Then the vector of channel usage indicators, $u_n = (u_{n1}, \dots, u_{nC})$ is defined as:

$$u_{nc} = \begin{cases} 1 & \text{if } s_{nc} > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Index classes with $d \in \{1, \dots, D\}$, and define z_n as the class label for observation (campaign) n . The likelihood contribution for observation n is:

$$\Pr(u_n) = \sum_d^D \nu_d \prod_c^C \Pr(u_{nc} | z_n = d) \quad (2)$$

Since u_{nc} is binary, we use a Bernoulli distribution to model the usage probability, with parameters $w_{cd} = \Pr(u_{nc} = 1 | z_n = d)$. Collect u_{nc} into a matrix \mathbf{U} . If $\boldsymbol{\nu} = (\nu_1, \dots, \nu_D)$ and \mathbf{W} is a $K \times D$ matrix containing the set of parameters $\{w_{cd}\}$ with rows corresponding to channels and columns to classes, then the Latent Class likelihood for the full dataset is:

$$\ell_{lc}(\mathbf{U} | \boldsymbol{\nu}, \mathbf{W}) = \prod_n^N \left[\sum_d^D \nu_d \prod_c^C w_{cd} \right] \quad (3)$$

EXPLORATORY LATENT CLASS ANALYSIS

To determine the number of classes, we estimated an exploratory Latent Class model 7 different times, with the number of classes varying from 2 to 8. Each time we computed the PSIS-LOO CV, the approximate leave-one-out cross-validation for Bayesian models using Pareto smoothed importance sampling (Vehtari, Gelman, and Gabry 2017, Vehtari et al. 2022). This allows for model comparisons using the pointwise expected log predictive density (ELPD). The best ELPD is reached at 7 classes. The ELPD values are in Web Appendix C.

The likelihood function in Equation (3) is invariant to permutations of the class labels. When running separate chains, convergence for class-specific parameters can fail because of this, since separate chains often explore relative to different orderings, which has the effect of reordering the columns of \mathbf{W} and the elements of $\boldsymbol{\nu}$. This problem is known as label switching. Label switching does not affect model outputs for which the discrete class probabilities have been marginalized out (such as the PSIS-LOO CV values used to determine the number of classes). However, to estimate class-specific parameters, label switching must be prevented. In practice, different approaches are used to solve this problem, but one of them is to use a “non-exchangeable prior”, meaning a prior which differs across classes, thus keeping each chain on the same relative ordering (Betancourt 2017) by penalizing the sampler if it moves toward different orderings. We use a non-exchangeable prior by first getting point estimates via the EM algorithm, a standard approach for Latent Class. We use the R package `poLCA` (Linzer and Lewis 2011) to get these parameter point estimates. Then we put priors on \mathbf{W} and $\boldsymbol{\nu}$ centered around the EM point estimates. This solves the label switching problem for our dataset. We provide details of the prior specifications and code in Web Appendices A and B.

Outcome Model: Bayesian Gamma Hurdle Regression

We have four dependent variables, which are all lifts expressed as decimals (differences between proportions). The outcome (lift) distributions are positively skewed. In principle a lift could be negative, but we observe this only once in the entire dataset.³ Thus, in practical terms, lifts are bounded at 0. The lifts are also bounded above by 1, as noted above. There are some exact zeros, however, which represent campaigns that had no detectable lift for a given outcome. To appropriately model each of these aspects, we use a hurdle model with a Gamma density for the response distribution.

We selected a Gamma regression by visual posterior predictive checks using only shares as variables. Candidate distributions were: Normal, Log-Normal, Exponential, Exponentially Modified Gaussian, Laplace, and Gamma. The Gamma best fit the lift distributions above zero. We then add the hurdle component to account for the mass at zero.⁴

Gamma regression is a generalized linear model (GLM) that assumes a Gamma distribution for the conditional distribution of the response. GLMs use a link function to transform the conditional mean of the response, which in our case can only take non-negative values, to the same space as the linear predictor $(x_n\beta)$. We pair this with a hurdle component parameterized by δ_{0n} . If y_n is an observation of an outcome, and x_n contains covariates, then the generative model for y_n is

$$y_n = \begin{cases} 0 & \text{with probability } \text{logistic}(\delta_{0n}) \\ \sim \Gamma(\exp(\beta_{0n} + \beta_1 x_n)\lambda, \lambda) & \text{otherwise.} \end{cases} \quad (4)$$

The variable λ is the shape parameter of the Gamma distribution. The expectation of the Gamma distribution in Equation (4) is $[\exp(\beta_{0n} + \beta_1 x_n)\lambda] / \lambda = \exp(\beta_{0n} + \beta_1 x_n)$. The link function $g(\cdot)$ for a GLM satisfies $E[y_n | x_n \beta_n] = \mu = g^{-1}(x_n \beta_n)$, our link function is $\log(\cdot)$.

³We drop this observation.

⁴The Gamma distribution can in principle produce values above 1, which is impossible for a lift. In practice this happens extremely rarely, perhaps 1 time per 1,000 draws. In all simulations from the model, we drop any draws above 1.

Intuitively, the outcome model has two components. The first component accounts for whether a lift value is exactly zero or positive. The second component models the magnitude of the outcome conditional on the outcome being positive. The first component, the hurdle, has a simple specification with random intercepts that vary according to category. Letting a index categories and letting a_n stand for the category of campaign n :

$$\begin{aligned}\delta_{0n} &= \delta_0 + \delta_{a_n} \\ \delta_{a_n} &\stackrel{iid}{\sim} N(0, \sigma_{\delta,a})\end{aligned}\tag{5}$$

In words, each category has a constant effect that shifts the probability of exact zeroes. The category constants share a common distribution.

The second component, the Gamma GLM, models the values of the outcomes conditional on those values being positive. The intercept β_{0n} has a random intercepts specification. Recalling that a indexes categories and letting j index brands:

$$\begin{aligned}\beta_{0n} &= \beta_0 + \beta_{a_n} + \beta_{j_n} \\ \beta_{a_n} &\stackrel{iid}{\sim} N(0, \sigma_{\beta,a}), \quad \beta_{j_n} \stackrel{iid}{\sim} N(0, \sigma_{\beta,j})\end{aligned}\tag{6}$$

Define $\eta_{0n} = \text{logistic}(\delta_{0n})$ and $\boldsymbol{\beta}_n = (\beta_{0n}, \beta_1)$. Then conditional on the observation-specific parameters, the likelihood contribution for observation n is

$$\Pr(y_n | \eta_{0n}, \boldsymbol{\beta}_n, \lambda) = \eta_{0n} \mathbb{I}\{y_n = 0\} + (1 - \eta_{0n}) \mathbb{I}\{y_n > 0\} \Gamma(y_n | \exp(\beta_{0n} + \beta_1 x_n) \lambda, \lambda)\tag{7}$$

Collect y_n into a vector \mathbf{y} , and similarly for x_n into a matrix \mathbf{X} , and for $\boldsymbol{\eta}$ and $\boldsymbol{\beta}$. Then the Gamma hurdle likelihood is:

$$\ell_{gh}(\mathbf{y}, \mathbf{X} | \boldsymbol{\eta}, \boldsymbol{\beta}, \lambda) = \prod_n^N \Pr(y_n | \eta_{0n}, \boldsymbol{\beta}_n, \lambda).\tag{8}$$

The full likelihood for the joint model is the product of the Latent Class likelihood and

the Gamma hurdle likelihood:

$$\ell(\mathbf{y}, \mathbf{X}, \mathbf{U} | \boldsymbol{\eta}, \boldsymbol{\beta}, \lambda, \boldsymbol{\nu}, \mathbf{W}) = \ell_{lc}(\mathbf{U} | \boldsymbol{\nu}, \mathbf{W}) \ell_{gh}(\mathbf{y}, \mathbf{X} | \boldsymbol{\eta}, \boldsymbol{\beta}, \lambda). \quad (9)$$

We estimate the standalone Gamma hurdle regression using the NUTS algorithm (Hoffman and Gelman 2014) from the `Stan` probabilistic programming language (Carpenter et al. 2017). We also estimate the joint Gamma hurdle and Latent Class model with NUTS.

Variables

The Gamma hurdle regression includes a random intercept by category for the hurdle as shown in Equation (5), and a random intercept for the outcome regression by category and brand as shown in Equation (6). In addition, the Gamma regression includes Archetype membership probabilities for each campaign, the concentration of spend for the campaign, a time trend, and the spend shares for each channel.

The membership probabilities link the Latent Class model to the Gamma hurdle model. At each iteration of the sampler, membership probabilities for each campaign are calculated conditional on the current Latent Class parameter draws and included in the Gamma hurdle as covariates. We leave Archetype 2 out as the reference class.⁵

We define

$$\text{concentration}_{nc} = \sum_c s_{nc}^2.$$

This is a version of the Herfindahl-Hirschman Index used as a metric for market concentration in economics (Tuckman and Chang 1991). Here we use the shares as decimals rather than whole numbers to keep consistent scales across variables. A concentration value of 1 means all budget is allocated to a single channel. For 11 channels, the minimum value of concentration is achieved when a campaign spends equally on all 11 channels: $(1/11)^2 + \dots + (1/11)^2 = 11 \times (1/11)^2 = 1/11 \approx 0.09$.

⁵Since the membership probabilities sum to one, the design matrix is not full rank if all are included.

We include a time trend, ranging from 2008 to 2019. Finally we include the shares $\{s_{nc}\}$ for each campaign and channel. We leave Online Video out as the reference channel.⁶ Table 3 summarizes the variables used and which models they enter.

Table 3: Variable Definitions

Variable	Definition	Enters which model
Channel indicators	Binary values: 1 means positive spend	Latent Class
Membership Probabilities	Campaign membership probability for each Archetype	Gamma hurdle
Concentration	Sum of squared spend shares	Gamma hurdle
Time trend	The campaign year (centered)	Gamma hurdle
Channel shares	Share of spend on each channel	Gamma hurdle

RESULTS

There are two sets of results to explore. The first are descriptive and detail different campaign archetypes. These results show which groups of media channels are commonly used together, and typical spend concentrations. These descriptive results are shown next.

The second set of results show the estimates of the regression coefficients relating archetypes, campaign concentrations, and allocation shares to the four outcomes. These results show which clusters and individual channels are associated with better advertising performance on each of the outcomes. We call this the outcome model and the results are shown below.

Descriptive Results

Table 4 gives the archetype profiles numerically, and the same results are plotted in Figure 6. The archetype prevalence is the parameter estimate of (ν_c) , which empirically is almost the same as the percent of campaigns assigned to each archetype.⁷ For each archetype, we

⁶Since the shares sum to one, the design matrix is not full rank if all shares are included.

⁷The difference exists because assignment to classes is stochastic.

also include in Table 4 the probability of usage for each channel, as well as the average concentration of spend among campaigns assigned to each archetype.

Table 4: Archetype Profiles

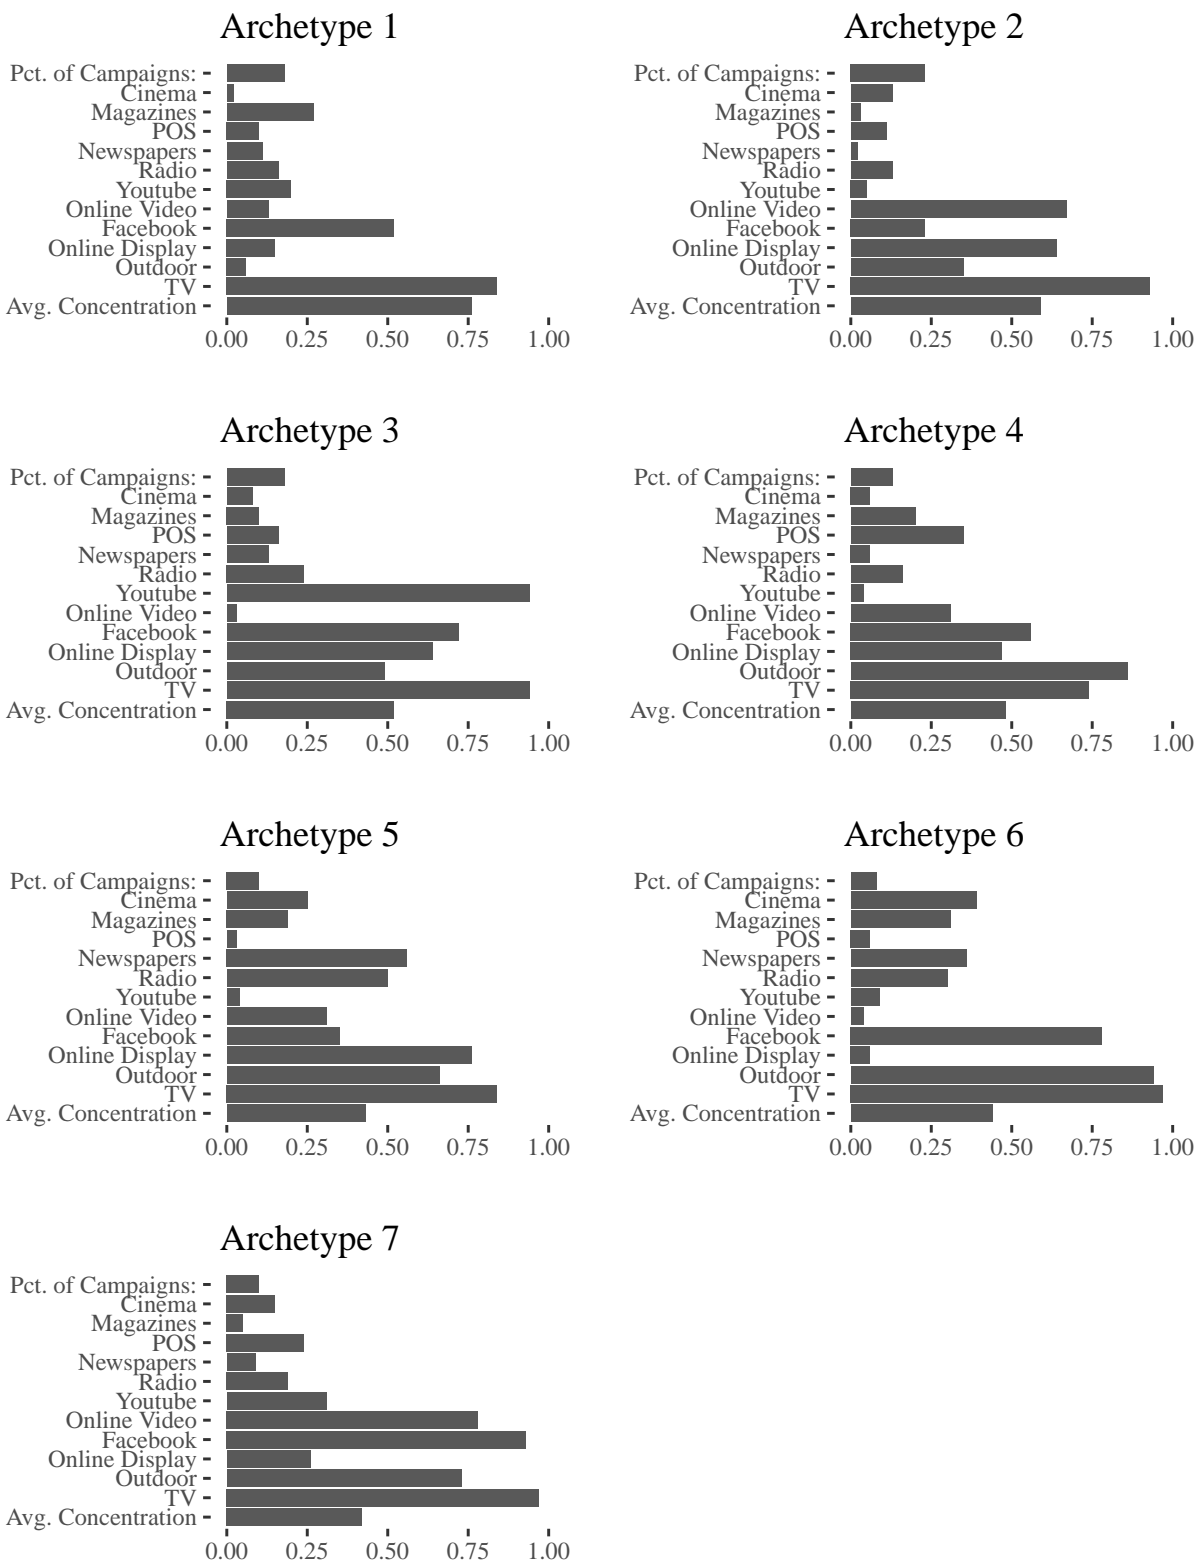
Archetype →	1	2	3	4	5	6	7
Prevalence:	.18	.23	.18	.13	.10	.08	.10
Usage Prob: Cinema	.02	.13	.08	.06	.25	.39	.15
Usage Prob: Magazines	.27	.03	.10	.20	.19	.31	.05
Usage Prob: POS	.10	.11	.16	.35	.03	.06	.24
Usage Prob: Newspapers	.11	.02	.13	.06	.56	.36	.09
Usage Prob: Radio	.16	.13	.24	.16	.50	.30	.19
Usage Prob: Youtube	.20	.05	.94	.04	.04	.09	.31
Usage Prob: Online Video	.13	.67	.03	.31	.31	.04	.78
Usage Prob: Facebook	.52	.23	.72	.56	.35	.78	.93
Usage Prob: Online Display	.15	.64	.64	.47	.76	.06	.26
Usage Prob: Outdoor	.06	.35	.49	.86	.66	.94	.73
Usage Prob: TV	.84	.93	.94	.74	.84	.97	.97
Avg. Concentration	.76	.59	.52	.48	.43	.44	.42

TV is the most commonly used channel, and cinema the least. Three channels see usage less than half the time for any archetype: cinema, magazines, and point of sale. Archetype 2 is the most prevalent (23%). To ease interpretation, we also provide Figure 7.

The map was created by multidimensional scaling with the L1 norm for measuring distance (also known as city block or Manhattan distance). The campaigns are plotted by weighting the Archetype coordinates by the campaign membership probabilities. An important takeaway from the archetype map is that campaigns are not easy to organize in terms of a few dimensions such as legacy vs. digital or broad vs. narrow. Archetype 5 is mostly legacy channels, but with a high tendency to use online display. Archetype 6 is similarly reliant on legacy channels but with a heavy use of Facebook.

The map is drawn and labeled according to what best discriminates archetypes. This does not mean that archetypes toward the extremes are “pure-play” but rather that they are distinctive for using some channels more or less. For example, Archetypes 7, 3, and 1 make high use of channels with algorithmic targeting (e.g., YouTube), but vary in degree of legacy channel usage. Archetype 7 uses outdoor 73% of the time while Archetype 1 uses

Figure 6: Archetype Profiles






Figure 7: Archetype Map

Notes: Filled dots are individual campaigns. Open dots are archetypes and size indicates prevalence. “Legacy” refers to cinema, point-of-sale, outdoor, magazines, radio. “Generic digital” refers to non-YouTube online video and non-Facebook online display. “Algorithmic” refers to algorithmic targeting (e.g., Youtube and Facebook)

outdoor 6% of the time. Despite this, they are still close because of the high tendency to include algorithmic digital channels.

Integrated Model Results

In Table 5, we present the results of the integrated Latent Class and Bayesian regression. As a reminder, the link between the two models is that the archetype membership probabilities for each campaign are included as covariates in the regression. Also recall that the regression incorporates a hurdle for zero vs. non-zero response, and fits a Gamma density to the positive portion of the response distribution. The results in the table are conditional on clearing the hurdle. Each coefficient should be viewed as the estimated effect on a campaign outcome *given* that the campaign has any detectable impact at all. Since we use a Bayesian approach, the results take the form of probability distributions for each parameter. For the coefficient tables in Table 5, we report the posterior means and 95% credible intervals.

Interpreting the coefficients here is not as straightforward as in conventional linear regression. As explained above, in order to model an outcome with a Gamma distribution, a link function is necessary to bridge between the right-hand-side variables and the responses. This is because the right-hand-side variables can often take negative values, which are not in the support of the Gamma density. The link function we use is the natural logarithm. The inverse link — the exponential in our case — is applied to the linear predictor $(x_n\beta_n)$ to transform values into \mathbb{R}^+ . For interpretation, we pass the coefficients through the exponential function, which allows them to be interpreted as factors. Explicitly, a coefficient β_x for a right-hand side variable x is interpreted as the factor by which the response is multiplied given a 1 unit increase in x . In some cases, this will create an impression of unrealistically large effect sizes, but in interpretation it is important to note the intercept term, which always applies. As will be evident, in each model the intercept balances out large coefficients on the covariates. With a typical regression coefficient, where effects are additive, zero (the additive identity) indicates a null effect. Here, effects are multiplicative, so a null effect is

one (the multiplicative identity). In Table 5 coefficients have been exponentiated already, so each number can be interpreted directly as a scaling factor.

Archetype membership probability coefficients are relative to Archetype 2, the reference level (which therefore has no coefficient). Each coefficient shown is the factor of improvement over Archetype 2. Similarly, spend share coefficients are relative to Online Video.

Starting in block (a) of Table 5, we have the results for Motivation. Archetype 4 has the largest significant coefficient. However, the share of spend on cinema has a larger coefficient, and share spent on YouTube has a coefficient nearly as large, on average. Also note that the *lower bound* of the credible interval for Archetype 3 is higher than for 4. This suggests the possibility that different channel selections come with different risk profiles, which highlights a reality where the concept of “best strategy” is more nuanced, with the potential to at least consider the risk appetite of brands. In fact, scanning across the four dependent variables in Table 5, this relationship between Archetypes 3 and 4 holds in all but one case. Note that concentration is not significantly different from 1 (multiplicative identity) for Motivation.

Coefficients for Association are in block (b). We see that again Archetype 4 has the largest coefficient among the archetypes, but four main effects (spend shares) are at least as large: cinema, YouTube, newspapers, and point of sale (POS). Concentration has a significant coefficient below 1, meaning more concentrated spend relates to *reduced* lifts for Association. This may be due to the fact that higher concentrations tend to be common among campaigns which spend more on TV, outdoor, or other “legacy” media, whereas lower concentration reflects experimentation and adoption of new channels. Heavy use of TV and other legacy media may in turn constrain or dampen the ability for brands to create changes in association, while on the other hand, more diverse and newer channels may afford more variability in exposure, new means of communication, as well as new signaling mechanisms. All conditions and consequences that are consistent with Keller’s (1993) theory and practical implications, as well as some early empirical evidence on advertising variance (Haugtvedt et al. 1994).

See coefficients for Unaided Awareness in block (c). The largest archetype coefficient is for 4, but by a very small margin. If the reference class were changed from 2 to 5, it is likely that Archetypes 3, 4, 5, and 7 would not be significantly different from one another. For Unaided Awareness, the coefficients on Archetypes are generally larger than the coefficients on spend shares (with the exception of cinema). In fact, several of the share coefficients are below 1, with magazines having an especially low coefficient.

Block (d) shows coefficients for Aided Awareness. As shown in Table 1, Aided Awareness has a lower average lift and a higher standard deviation than the others. Aided Awareness also has more zeroes, consistent with a ceiling effect for large brands (Srinivasan, Vanhuele, and Pauwels 2010). The larger variation and reduced signal in the Aided Awareness lift data manifests in much larger credible intervals for the coefficients. However, relative effects still reveal interpretable results, such as the fact that Archetype 4 and share of spend on cinema have distinctly larger impacts on Aided Awareness than the other variables.

While the raw coefficients are very informative, model simulations can show the relationships between outcomes and more holistic changes to campaign structure, encompassing channel selection, spend shares, and concentration. We show these holistic relationships in Figures 8 and 9. The x-axis of each figure starts on the left with the posterior predictive median for a campaign of Archetype 2: membership probability of 1 for Archetype 2, and the average channel selection, spend, and concentration of Archetype 2 campaigns in the data. Then we move the membership probability, average spends, and average concentration values progressively toward different clusters, and plot the predicted outcome. In Figure 8 we plot the median value of the posterior predictive distribution, and for Figure 9 we plot the 2.75th percentile. The differences indicate that sometimes an approach optimized for a median result would not be the same as more risk averse ‘maximin’ approach. For example, Archetype 4 delivers by far the best median result on Aided Awareness, but by far the worst 2.75th percentile result — it is a high risk, high reward campaign structure for Aided Awareness.

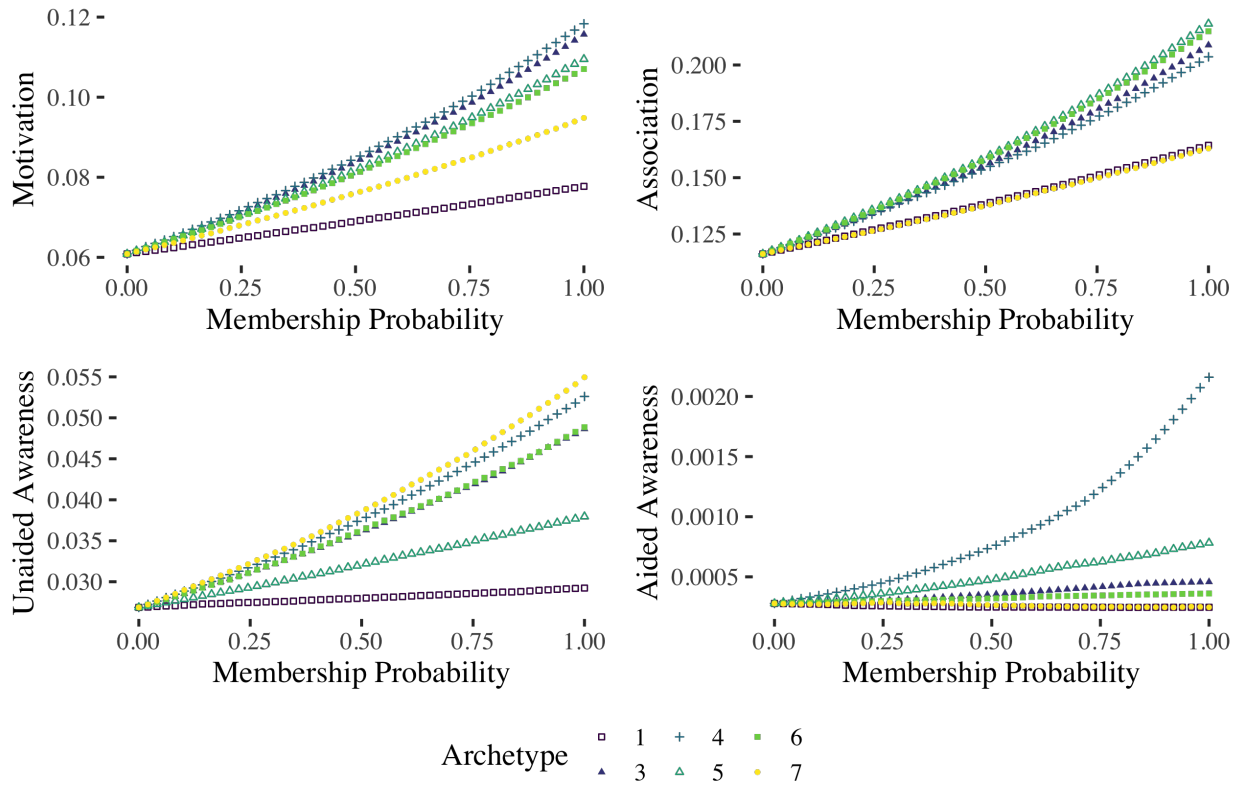


Figure 8: Archetype Outcome Profiles: Median

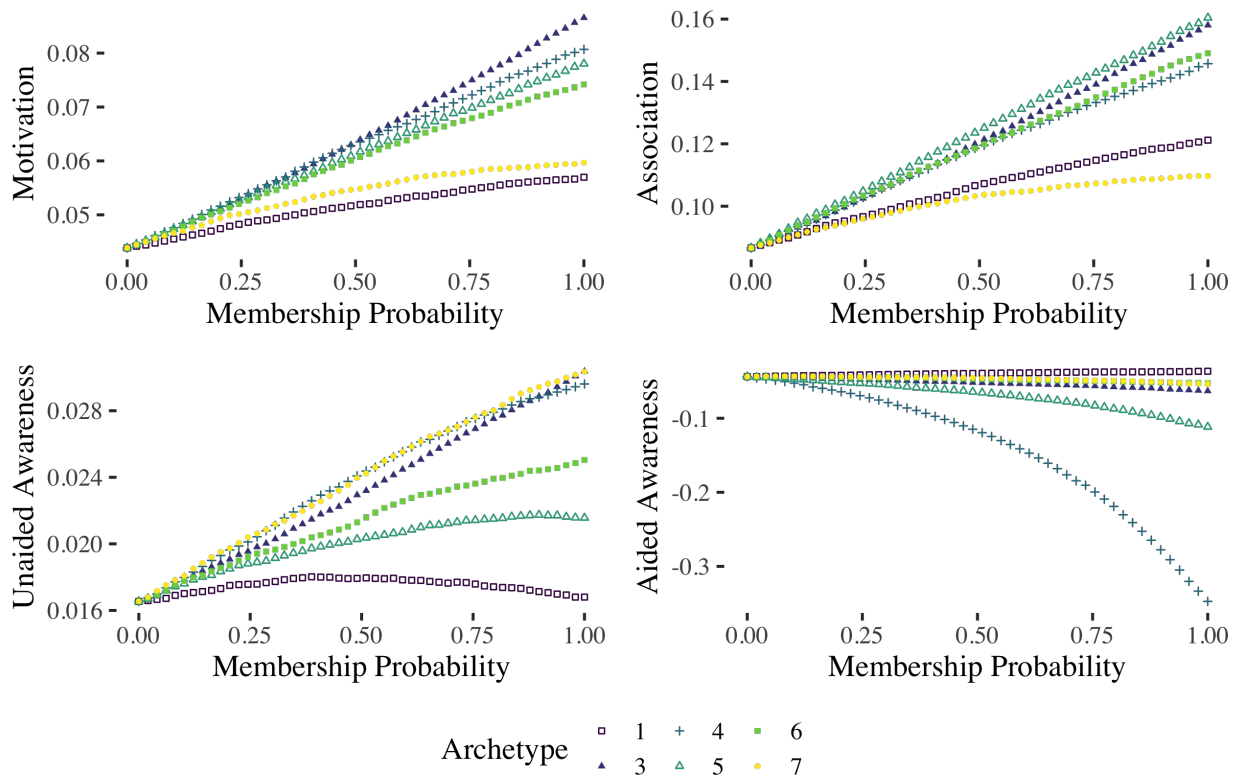


Figure 9: Archetype Outcome Profiles: 2.75th Percentiles

Table 5: Outcome Model Results

	(a) Motivation			(b) Association		
	Mean	CI:L	CI:U	Mean	CI:L	CI:U
Intercept	0.01	0.01	0.02	0.01	0.01	0.02
Archetype 1	1.22	0.91	1.65	1.38	1.04	1.85
Archetype 3	1.99	1.56	2.60	1.83	1.45	2.34
Archetype 4	2.25	1.51	3.51	2.02	1.41	2.94
Archetype 5	1.74	1.25	2.49	1.68	1.22	2.31
Archetype 6	1.66	1.18	2.45	1.59	1.12	2.32
Archetype 7	1.65	1.05	2.60	1.45	0.99	2.14
Concentrator	0.95	0.80	1.13	0.68	0.58	0.81
Time Trend	1.02	1.00	1.05	1.00	0.98	1.03
Share: Cinema	3.34	2.78	4.04	6.05	4.98	7.33
Share: Facebook	1.33	1.12	1.58	1.68	1.42	2.01
Share: Magazines	0.69	0.57	0.83	0.60	0.49	0.72
Share: Newspapers	1.98	1.66	2.37	3.34	2.77	4.03
Share: Online Display	1.09	0.91	1.32	1.77	1.50	2.10
Share: Outdoor	0.65	0.55	0.78	0.77	0.64	0.92
Share: POS	1.09	0.90	1.32	2.02	1.68	2.43
Share: Radio	0.89	0.74	1.07	1.54	1.28	1.86
Share: TV	1.17	1.01	1.36	1.73	1.49	2.00
Share: Youtube	2.08	1.74	2.48	3.55	2.97	4.27
	(c) Unaided Awareness			(d) Aided Awareness		
	Mean	CI:L	CI:U	Mean	CI:L	CI:U
Intercept	0.01	0.01	0.02	0.01	0.00	0.01
Archetype 1	1.12	0.67	1.82	0.68	0.39	1.20
Archetype 3	2.03	1.38	3.02	1.83	1.10	3.11
Archetype 4	2.10	1.19	3.77	10.30	4.51	27.22
Archetype 5	1.90	1.13	3.28	2.89	1.37	6.10
Archetype 6	1.76	0.97	3.33	1.23	0.65	2.38
Archetype 7	2.06	1.18	3.74	1.35	0.49	3.58
Concentrator	1.20	1.00	1.43	1.76	1.47	2.11
Time Trend	1.06	1.02	1.10	1.04	0.98	1.10
Share: Cinema	6.69	5.54	8.14	26.11	21.44	31.48
Share: Facebook	1.40	1.16	1.69	4.26	3.54	5.18
Share: Magazines	0.07	0.06	0.08	1.20	0.99	1.46
Share: Newspapers	1.74	1.44	2.12	2.59	2.14	3.14
Share: Online Display	0.79	0.66	0.96	0.88	0.73	1.06
Share: Outdoor	0.74	0.62	0.90	0.45	0.38	0.55
Share: POS	0.39	0.32	0.47	2.79	2.31	3.39
Share: Radio	0.81	0.66	0.98	1.81	1.50	2.20
Share: TV	1.07	0.90	1.26	1.74	1.47	2.07
Share: Youtube	1.59	1.31	1.93	6.98	5.78	8.41

Notes: CI:L is the 2.75th percentile and CI:U is the 97.5th percentile, which together form the 95% credible interval.

To give a sense for how the best archetype differs by industry or category, we show in Table 6 the best Archetype coefficient when the model is estimated on subsamples of the data. While some aggregation is necessary to maintain reasonably large sample sizes, we are able to provide results for three mega-category subsamples: consumer packages goods (CPG), services, as well as technology & durables. Notably, different categories have distinctly different profiles. Archetype 6 is best for all four outcomes in CPG, but for none of the outcomes in the other categories. Archetype 4, though best in the overall dataset, is only best in one case (Unaided Awareness in the Services category) for the subsamples. Archetype 6 is notable for having a very high probability of combining TV, outdoor, and Facebook. That this triple performs extremely well for CPG calls into question the common conceptual division between traditional vs. digital media. It appears traditional and digital media can be highly complementary.

Technology and Durables shows a clear division between awareness (“top of funnel”) and motivation and association (“lower funnel”). Archetype 7, notably high in Facebook and YouTube, unsurprisingly is associated with lower funnel metrics, while Archetype 5, which is distinctly lower in YouTube and Facebook, associates with upper funnel awareness metrics.

Table 6: Archetypes with Highest Coefficient by Industry

	CPG	Services	Technology & Durables
Motivation	6	5	7
Association	6	5	7
Unaided Awareness	6	4	5
Aided Awareness	6	7	5

Notes. CPG: personal care, food and drink, household.
 Services: financial services, airlines, utilities, telecoms, leisure.
 Technology & durables: automobiles, consumer electronics,
 apparel, search, social media, ecommerce.

PRESCRIPTIVE RESULTS

The model can be used to simulate the impact of possible changes to marketing media planning. For each of the four dependent variables (Motivation, Association, Unaided Awareness, and Aided Awareness) we consider the predicted impact of a change in the planning mix. These predictions are not intended to be causal estimates, but instead they indicate potential opportunities for improving lifts based on typical campaign performance in our data. The possible changes we consider all take the form of “acting like” a particular archetype. In concrete terms, “acting like” an archetype means selecting channels with the same probabilities as that archetype (i.e. to act like Archetype d , choose channel c with probability w_{cd}), and then conditional on channels selected, allocate spend like the average member of that archetype in the dataset.

To implement the simulation, we first use the model estimates for channel usage probabilities by archetype (shown in Table 4). We next compute the average spend in the dataset by archetype, after each campaign is assigned to the likelihood for which it has highest membership probability. The spend shares by archetype are shown in Table 7.

Table 7: Average Spend Shares by Archetype

Archetype →	1	2	3	4	5	6	7
Share: Cinema		.01	.01	.01	.03	.04	.01
Share: Magazines	.06		.01	.02	.01	.03	
Share: POS	.02	.02	.02	.06			.04
Share: Newspapers	.02	.01	.02		.13	.04	.01
Share: Radio	.04	.02	.03	.02	.06	.03	.01
Share: Youtube	.04		.10				.02
Share: Online Video		.11		.03	.02		.08
Share: Facebook	.11	.03	.06	.05	.02	.11	.07
Share: Online Display	.01	.13	.08	.08	.06		.01
Share: Outdoor		.05	.08	.35	.17	.23	.26
Share: TV	.69	.62	.59	.38	.49	.53	.49

Notes: shares below 0.01 have been rounded to zero.

To generate a campaign profile that approximates an archetype, we follow the steps below. We use S to stand for simulated quantities. Consider Archetype d :

1. Draw channel usage indicators. For each c , draw u_c^S from a Bernoulli distribution with parameter \hat{w}_{cd} , which is Archetype d 's estimated usage probability of channel c .
2. From mean spend shares shown in Table 7 select values s_c^* for channels where $u_c^S = 1$.
3. To ensure that the shares sum to one, normalize: $s_c^S = s_c^* / \sum_c s_c^*$.

We then simulate from the posterior predictive distribution of the model to obtain expected lifts. We run the above steps for each archetype. Finally, we compare the expected lifts to average observed lifts. Table 8 shows the expected lift, in absolute terms and as a percentage improvement, for the best archetype for each outcome.

Table 8: Model Recommendations

Outcome	Best Archetype	Channel to Upweight	Improvement	
			Absolute	Percentage
Motivation	4	POS	+0.018	+83%
Association	3	YouTube	+0.025	+96%
Unaided Awareness	7	Online Video	+0.009	+52%
Aided Awareness	4	POS	+0.095	+750%

Notes: Absolute improvement is the mean simulated lift minus the mean observed lift in the dataset. Percentage improvement is absolute improvement divided by mean observed lift in the dataset. Channel to upweight is the channel which had the largest increase in probability of usage, as a percent of average observed usage probabilities.

Given the large amounts of money to achieve lifts, these results appear to be extremely economically meaningful. One note is in order about the improvement in Aided Awareness: Figures 8 and 9 show that while Archetype 4 performs much better on *average*, it also comes with very large risk. There is also a very high likelihood of zero lifts in the data (see the hurdle intercept in Table 5), which suggests a relatively lower amount of informative data and thus noisier estimates, but also suggests a very low baseline level of Aided Awareness in

the dataset upon which to improve. Together these factors suggest that the improvement is credible but should be interpreted with some caution and with proper context.

Rough Financial Impact Figures

A dataset comparable to ours in richness and scope, but which also includes measurements of sales outcomes, is very unlikely to exist. Given the ability of this dataset to allow us to answer questions on relationships between media usage and brand metric lifts, the lack of financial outcomes seems a more-than-acceptable trade-off. However, despite lacking sales data linked to each campaign, we can still produce some rough back-of-the-envelope estimates of financial impact from the prescriptions in Table 8 by relying on previous research. The best match for this purpose is [Srinivasan, Vanhuele, and Pauwels \(2010\)](#), wherein the authors produce sales elasticities with respect to mindset metrics similar to those in our data.

The metric from our dataset which best matches the metrics from [Srinivasan, Vanhuele, and Pauwels \(2010\)](#) is Motivation, which is a combination of consideration, purchase intent, and brand liking. The corresponding metrics reported in [Srinivasan, Vanhuele, and Pauwels \(2010\)](#) are brand liking and consideration. We note that purchase intent was also in the dataset in [Srinivasan, Vanhuele, and Pauwels \(2010\)](#) but was not used according to that paper because it was so highly correlated with consideration. Our approach for the purposes of this thought experiment are to use the elasticity of sales with respect to consideration as a lower estimate and the the elasticity with respect to liking as a higher estimate. The rationale is that our metric is some combination of the two things, and so the elasticity would naturally fall somewhere between them. The reference unit we use is millions of dollars, meaning our goal is to calculate the incremental revenue per \$1M dollars of revenue that could be expected by increasing Motivation by .018 (first row of Table 8). Because the results from [Srinivasan, Vanhuele, and Pauwels \(2010\)](#) are elasticities, we need to assume a starting level of Motivation. In our dataset, Motivation is measured on a scale between 0 and 1, as discussed, where 1 represents 100% of respondents giving answers indicating favorable

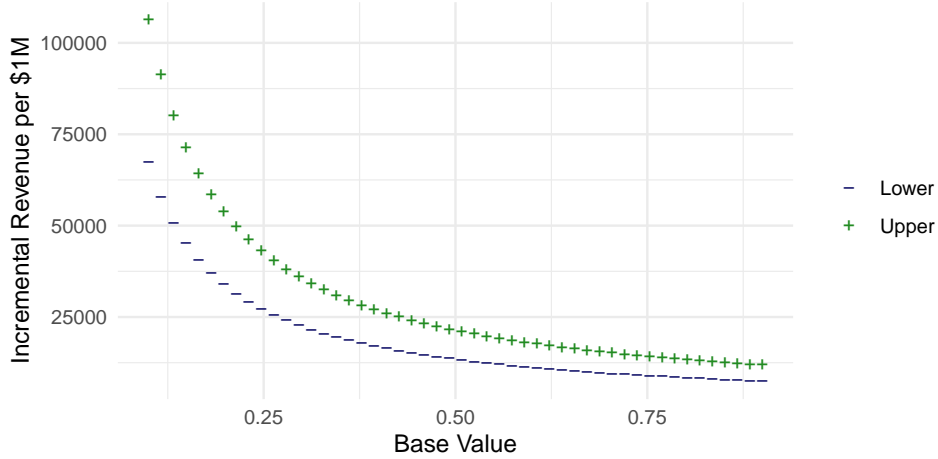


Figure 10: Financial Impact through Motivation

attitudes toward the brand. To produce rough figures for financial impact then, we use a set of equally spaced grid points between 0.1 and 0.9, which represent different starting levels of Motivation. For each grid point, we use the following formula to get estimates for incremental revenue on \$1M (ΔR):

$$\Delta R = \frac{\varepsilon \times 0.018 \times 1,000,000}{\text{grid point}}, \quad (10)$$

where ε is a cumulative elasticity figure from [Srinivasan, Vanhuele, and Pauwels \(2010\)](#) (either 0.374 if associated with consideration for the lower estimate, or 0.59 if associated with liking, for the upper estimate). Figure 10 shows the results. Elasticity naturally captures diminishing returns: if Motivation levels are very low, the incremental impact is much larger than if Motivation levels are already very high. The highest estimated impact is approximately \$100,000 additional dollars per \$1M in current revenue, and the lowest is approximately \$7,500. Adopting a particular media selection archetype does not require changing spend levels but only re-allocating spend, which means that even if realized incremental revenue were on the very low end of this range, the intervention would be worthwhile.

Pairwise Interaction Models

To help see why modeling channel complementarity is useful, as well as to understand the benefit of using Latent Class as a first stage, we can compare our results to the most common alternative: pairwise effects, without applying Latent Class first. The results for our data are shown in Web Appendix D.

Pairwise interaction models for planning ten channels show 56 rows of data, including estimates of interaction terms. However, this output has severe limitations. First, the coefficients are mostly untrustworthy due to the number of parameters relative to the number of data points. Second, there is no principled way to determine which, if any, coefficients to focus attention on for managerial insights. Third, even accepting the complications arising from this model, it still only captures pairwise interactions. Lastly, by construction, this model collapses the channel selection problem into either a narrower cross-effects individual efficiency construct, or a coarser and less informative "synergy" construct (and all channels in the 'on' position); risking an analysis that optimizes to individual channel effectiveness rather than campaign level outcome. Altogether, this comparison demonstrates why our approach is more statistically and practically tractable and valid. More significantly, the comparisons show that our approach is both uniquely able to address the conceptual question of channel selection and complementarity, while also being more insightful to managers.

MOVING IMC FROM CONCEPT TO THEORY

While not testing a specific theory, our results are very useful in the context of the original constructs inside of the IMC concept. Furthermore we see our results as a motivation for a more comprehensive theory of marketing integrated communications that provides testable hypotheses, predictive power, and valuable mechanic explanation for the behavior and interplay of media.

First, we were able to identify and demonstrate the existence of a significant and prob-

lematic managerial lay theory. Consistent with broad managerial guidance, we find that managers seek to optimize the reach of their campaigns, or in other words, they seek to maximize the size and exposure of the audience (given their brand targeting parameters). Furthermore, the argument arising from this theory is that audience reach, and this significant exposure is sufficient, such that all other outcomes for the brand should naturally arise (reach efficiency leading to awareness, consideration and ultimately purchase). To be clear, by construction, this lay theory posits that all channels are undifferentiated outside of their audience composition and ability to deliver against this audience.

However, we find the opposite outcome with respect to brand consequences. Empirically, this optimization in reach has failed to translate into optimized brand outcomes. Therefore, we can point to a (at minimum) 2-dimensional heterogeneity of channels: (1) audience composition, and the channel's ability to deliver messages to this audience, and (2) a channel specific functional advantage, effect, or impact it has on its audience that arises from the channel itself and not from creative or content choices.

Secondly, we find that different brand outcome goals (Aided Awareness, Unaided Awareness, Association, and Motivation) yield different best archetypes, or combinations of selected channels. This highlights that not only does the functional-level heterogeneity exist, but that channel function carries different relevance or value for different goals. And, as this differential combinatorial value of channels by outcome is the suggested required element of complementarity, we find that our results are a strong indication that complementarity is an important and separate construct, which should be maintained as separate component inside of the IMC concept.

What is complementarity, and how might it arise? For exposition, let's consider a stylized construction of a single objective, such as brand salience. We can then imagine that building salience requires two sub-processes. We can leave them unnamed and assume they are latent, but maintain that they are distinct and necessary components of creating salience. In this stylized discrete world, let's assume that radio advertising acts on sub-process 1 exclusively,

while television and outdoor advertising both act on sub-process 2 alone. If this were the case, radio and television would be complementary investments. This tiny campaign with two channels would address each sub-process of salience, and together they would perfectly make up for each others' deficiency (no channel handles all sub-processes alone). Similarly, and alternatively, radio and outdoor would also form complementary investments, with radio handling sub-process 1 and outdoor handling sub-process 2. However, television and outdoor are not complementary channels with respect to salience, as each would be acting on sub-process 2, and this campaign construction would leave sub-process 1 devoid of activity. Complementarity deals directly with the question of which channels should be used, based on their function with respect to an objective.

What then of cross-effects? Still in our stylized world described above, both radio+television and radio+outdoor represent complementary selections. However, we might imagine that — for whatever reason — the joint presence of television ads actually improves the effectiveness with which radio acts on sub-process 1. Alternatively, we might imagine that the combination with outdoor actually harms the effectiveness of radio acting on its own separate sub-process. These positive or negative effectiveness gains are the channel cross-effect. They do not alter where the channel is active, but the strength with which it acts. As such, cross-effects will impact spend, inasmuch as we can reduce spend given gains in effectiveness and efficiency due to a supporting channel, or we can compensate from loss of efficiency from a negative cross-effect by pumping ever more money into a channel. But, there is not any amount of investment that will cause a channel to switch functionality and create complementarity. These are likely an emergent characteristic from channel design interacting with consumer utilization behaviors and patterns.

Therefore, complementarity addresses the channel selection problem, while cross-effects, on its own, provides an understanding of the effectiveness of a channel in the context of a group of channels working together. This is not to undermine the value of cross-effects within the IMC concept, but rather to flag the importance for all constructs, and the need

to understand where and how they act.

Further towards our understanding of marketing integration, we identify a fascinating boundary condition in product category. As noted before, of the three IMC constructs, one (consistency) exists at a content level, while the other two (complementarity and cross effects) exist at the channel level. However, we were able to identify one product category (CPG) where complementarity — a channel level effect — collapses and a singular combinatorial archetype becomes a dominant strategy (archetype 6), regardless of brand advertising goal.

Critically, this points towards important and nuanced considerations when explaining the interaction of media strategies, brand contexts, and brand goals, particularly given desires to simplify the IMC concept. Whereas, given the explosion of media options and the critical dynamic multi-dimensionality of media value (audience, function and interactions), we might instead benefit from an opposite move towards a homegrown IMC theory, providing a complete set and defined constructs, their compositions, relational structure and applicable boundary conditions. Ultimately, this shift would yield us much great explanatory and predictive power for the complex media environment, as well as a course-correction and better guidance for practice.

DISCUSSION

Our study demonstrates that there is a significant value in considering the complementarity of media channels in campaigns, beyond pairwise interactions. The analysis takes some initial steps in safeguarding the necessity of complementarity inside IMC theory and literature and in answering the call by [Batra and Keller \(2016\)](#) for better managerial guidance on the usage and combination of media channels.

One major implication is that there is no universal silver-bullet campaign strategy. Across all campaign archetypes identified and possible outcomes, no single combination of channels offers the best possible result. We show that the concept of “best” can even differ by risk appetite, such that some might aim for highest average outcome lifts, while others might

prefer a more conservative lower downside risk option. Different campaign goals and industry contexts require different media combinations.

That does not mean all channel combinations have a ‘best’ purpose. While we find no *dominant* strategies, we do find *dominated* ones. Archetype 2 (i.e., “Generic Digital”) might serve as the best example. As a reminder, this archetype combines TV with online video and display, likely pointing to a common practice of replicating the TV asset on online video and banners/display. Despite being the observed format of 23% of campaigns, it essentially never offers best performance, regardless of campaign goal, outperforming only Archetype 1 for Aided Awareness (but losing to 3, 4 and 5).

The complementarity of channels in this study raises another important consideration. Channels do in fact have channel-specific strengths and weaknesses that are mostly invisible if we are restricted to the consideration of a singular outcome or dimension of channel performance. This means there are reasons to use (or not use) a channel beyond reach, impressions, or frequency. The coverage of a target market (audience composition) is not the only contribution of a channel. One might excel in awareness, while others drive motivation. Channels are not perfectly fungible, and managers should consider the unique individual functionality of channels relative to other available channels, along with the potential for complementarity in the media plan. The implication is that the value of a media channel is not a function of simple supply constraints (i.e., “how many views can I buy?”), but also of the channel’s positioning and differentiation (i.e., “what do I accomplish with a view?”)

We demonstrate the importance of considering higher-order (i.e., more than pairwise) relationships in media planning and in research. Archetype 6 (Television + Facebook + Outdoor) provides a good example, as the performance of any sub-component dyad suffers by the removal of the third component. Despite comprising only 8% of our observed campaigns, this archetype is the closest to a silver-bullet, offering the best outcomes regardless of campaign goal, provided the brand was in the CPG category. However, for all other categories, archetype 6 was *never* ideal. These results highlight that the division between

traditional and digital channels may be a false dichotomy. While practice and academia might usually consider these to be in opposition, or two separate but internally cohesive groups (Naik and Peters 2009), the reality appears to be more nuanced and consideration should be given to the functional complementarity of channels regardless of their categorization. Again, the triad of TV, Facebook, and outdoor as a high performance strategy necessitates the use of billboards; hardly fitting with the narrative that advertising might be strictly a digital focus in the future.

Table 8 highlights goal specific characteristics. Across awareness (aided or unaided), association, and motivation, there are a variety of archetypes to follow, channels to invest in, and a range of resulting improvements. This indicates that there are natural differences in our ability to influence different aspects of the brand, or to comparatively call the customer to action successfully, which further reinforces the need to respect channel heterogeneity.

Implications for Practice

The present study provides new insights for practitioners who are concerned with multimedia campaign construction, demonstrable advertising effectiveness, and the balance between brand building versus consumer activation (Binet and Field 2013). We begin by showing common patterns in campaign structures via the Latent Class approach, as shown in Table 4. Current industry discussion on media selection centers on the immediate and growing investment in digital channels, which carries the risk of inflating expectations of digital channel performance and applicability. This increases the risk of over-indexing on these channels as Adidas claims to have done in late 2019 (Vizzard 2019). We then provide a view on how some of the largest global brands have executed their campaigns, providing a counterweight to that discussion and a reference based in recent historical evidence where traditional channels are still very much an important feature of integrated campaign design.

We show that the GRP-only paradigm of media selection and purchase is problematic in that it reduces media plan construction down to a too-simple market coverage problem. If the

GRP-only hypothesis were true, we would be able to identify one single campaign structure to dominate all others (i.e. the combination of channels providing the single largest possible market coverage). This practice is unfortunately an increasingly dominant unspoken strategy for media planners and reinforced by the pricing structure of channels, as they are presented on the basis of GRPs, impressions, or opportunities to view, depending on the context. This is a supply dominant logic for platforms and channels that ignores the needs and motivations of brand owners. One might ask, for example, why buy into cinema advertising when I can have more impressions on Facebook? The answer is: because they are functionally different, and a cinema exposure will elicit a different response when compared to a Facebook exposure. The reality of the dominance of reach-buying is clear from Table 2, as is the extent of potential for improvement on lifts.

The ease of falling into this framing of advertising is fairly straightforward to show if we consider the discussion around the health of channels. For example, in popular and industry press, Facebook has historically been evaluated primarily by the number of users and their demographics. As another example, the US bill to ban TikTok elicited an immediate advertiser discussion on diversifying media portfolios in order to reach heavy TikTok users (Barnett 2024). In these and most other instances, the question of media allocation is reduced to audience construction and efficiency of reach, but the functional benefit of a channel-specific exposure is never addressed. As we have shown, there is a significant amount of additional advertising power available to brands if they consider channel function as well as audience.

Furthermore, we provide specific guidance on channel performance and channel complementarity for four possible brand campaign goals: Unaided Awareness, Aided Awareness, brand Association, and Motivation. This is particularly important, as our data shows that even among the largest firms and brands present in our dataset, a number of campaigns are deploying strictly worse-performing combinations of channels. Instead, by following the goal-specific structures shown in our study, brands should be able to realize both greater

and more consistent campaign lifts across the whole range of the purchasing funnel.

Implications for Research

This work provides a number of contributions to the literature on advertising and integrated marketing communications. First, by leveraging the largest and broadest view to date on campaign effectiveness, we produce evidence of channel-specific, as well as structure and goal dependent campaign effectiveness. We provide an initial empirical response to [Batra and Keller \(2016\)](#), based on the meta-analytic view of over one thousand campaign lift studies.

Our approach can also serve as a foundation for researchers to operationalize high level constructs in IMC. Simultaneously, via better measurement and utilization of all high-level IMC constructs, we were able to generate predictions on advertising campaign construction and outcomes. However, this opportunity opens a set of valid questions around the nature of channel complementarity itself, designed-for versus emergent characteristics, dynamic as well as competitive contexts, and beyond. There is a need to connect micro (consumer level, processing) effects to macro (channel and firm level, econometric) effects ([Batra and Keller 2016](#)). As such, this study might be an impetus and the first step in transitioning IMC from a concept towards a formalized and structured theory, which would enable an increasingly cogent set of future analyses on the nature of media evolution and utilization, including cross-country variation and deeper competitive analysis.

Our results point to functional heterogeneity in channels, beyond audience composition and size. Not only does this bring to question the strict digital versus traditional media split we've already discussed, but also highlights a significant risk by researchers assuming that all channels within a single domain are equivalent. For example, collapsing a variety of channels into a collective "Social Media" label should be done cautiously, as it would imply that Facebook, Twitter, Instagram, Snapchat, TikTok, and beyond are functionally equivalent. In broadest terms, our results suggest that assumption is wrong.

Relatedly, by showing that media channels are indeed functionally heterogeneous, an immediate set of questions arise around the identification of divergent functions, as well as the causes for these differences. And while we cannot exhaustively provide channel functionality categorization here, we do propose that these differences should arise from an interplay of channel design (e.g. feed construction, velocity of content delivery, aspect ratio, etc.) and user behavior (e.g. navigation, rationale for use, usage occasion). Future research should be able to address precise channel function attribution, sharpening advice to advertisers as well as channel owners.

Lastly, we shed light on the under-researched area of brand-building (Unaided Awareness, Aided Awareness, Brand Association, and Motivation). While the call for sales as a critical outcome variable (Sethuraman, Tellis, and Briesch 2011) has been thoroughly accepted (see Table ?? in Web Appendix F for sales focused studies), the exclusive dedication to sales leaves a significant gap in understanding brand building. This work helps to fill this gap, and supports the view that campaign effectiveness is best served by considering brand metrics in addition to sales.

Note that while our analysis covers many channels, it is not exhaustive. Due to lack of data, we are missing search advertising and some new and popular additions (e.g. TikTok). This means caution is warranted when it comes to making precise recommendations suggested from this study. Instead, we focus on the larger, more general findings, such as the fact that no single plan is best for all outcomes and contexts. These insights will almost certainly remain relevant with the addition of more channels, since the answers will grow more complex, not simpler. Furthermore, businesses can likely improve ROI at a broader level by examining their historical media mix data and moving toward archetypes that are better for their contexts and goals.

REFERENCES

- Barnett, Kendra “Advertisers on alert as bill to ban TikTok or force divestment advances in US House,” <https://www.thedrum.com/news/2024/03/07/advertisers-alert-bill-ban-tiktok-or-force-divestment-advances-us-house> (2024).
- Batra, Rajeev and Kevin Lane Keller (2016), “Integrating marketing communications: New findings, new lessons, and new ideas,” *Journal of Marketing*, 80 (6), 122–145.
- Berman, Ron (2018), “Beyond the last touch: Attribution in online advertising,” *Marketing Science*, 37 (5), 771–792.
- Betancourt, Michael “Identifying Bayesian Mixture Models,” (2017) https://mc-stan.org/users/documentation/case-studies/identifying_mixture_models.html.
- Binet, Les and Peter Field (2013), *The Long and the Short of It: Balancing Short and Long-Term Marketing Strategies* Institute of Practitioners in Advertising.
- Borden, Neil H (1964), “The concept of the marketing mix,” *Journal of advertising research*, 4 (2), 2–7.
- Carpenter, Bob, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell (2017), “Stan: A probabilistic programming language,” *Journal of statistical software*, 76 (1).
- Chang, Yuhmiin and Esther Thorson (2004), “Television and web advertising synergies,” *Journal of advertising*, 33 (2), 75–84.
- Culliton, James William et al. (1948), *Management of marketing costs* Division of Research, Graduate School of Business Administration, Harvard Univ.
- Danaher, Peter J and Tracey S Dagger (2013), “Comparing the relative effectiveness of advertising channels: A case study of a multimedia blitz campaign,” *Journal of Marketing Research*, 50 (4), 517–534.
- Danaher, Peter J, Tracey S Danaher, Michael Stanley Smith, and Ruben Loaiza-Maya (2020), “Advertising effectiveness for multiple retailer-brands in a multimedia and multichannel environment,” *Journal of Marketing Research*, 57 (3), 445–467.
- Deighton, John (1996), *Features of good integration: two cases and some generalizations* Hillsdale, NJ: Erlbaum.
- Dijkstra, Majorie, Heidi EJJM Buijtsels, and W Fred Van Raaij (2005), “Separate and joint effects of medium type on consumer responses: a comparison of television, print, and the Internet,” *Journal of Business Research*, 58 (3), 377–386.
- Duncan, Tom and Sandra E Moriarty (1998), “A communication-based marketing model for managing relationships,” *Journal of marketing*, 62 (2), 1–13.
- Edell, Julie A and Kevin Lane Keller (1989), “The information processing of coordinated media campaigns,” *Journal of Marketing Research*, 26 (2), 149–163.
- Fischer, Marc (2019), “Practice Prize Paper—Managing Advertising Campaigns for New Product Launches: An Application at Mercedes-Benz,” *Marketing Science*, 38 (2), 343–359.
- Fischer, Marc, Sönke Albers, Nils Wagner, and Monika Frie (2011), “Practice prize winner—dynamic marketing budget allocation across countries, products, and marketing activities,” *Marketing Science*, 30 (4), 568–585.
- Fossen, Beth L and David A Schweidel (2019), “Social TV, advertising, and sales: Are social shows good for advertisers?,” *Marketing science*, 38 (2), 274–295.

- Gatignon, Hubert and Dominique M Hanssens (1987), “Modeling marketing interactions with application to salesforce effectiveness,” *Journal of Marketing Research*, 24 (3), 247–257.
- Gopalakrishna, Srinath and Rabikar Chatterjee (1992), “A communications response model for a mature industrial product: Application and implications,” *Journal of Marketing Research*, 29 (2), 189–200.
- Haugtvedt, Curtis P., David W. Schumann, Wendy L. Schneier, and Wendy L. Warren (1994), “Advertising Repetition and Variation Strategies: Implications for Understanding Attitude Strength,” *Journal of Consumer Research*, 21 (1), 176–189.
- Hoffman, Matthew D and Andrew Gelman (2014), “The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo,” *J. Mach. Learn. Res.*, 15 (1), 1593–1623.
- Houston, Franklin S (1986), “The marketing concept: what it is and what it is not,” *Journal of marketing*, 50 (2), 81–87.
- IMARC “Radio Advertising Market: Global Industry Trends, Share, Size, Growth, Opportunity and Forecast 2023-2028,” (2023) <https://www.imarcgroup.com/radio-advertising-market>.
- Inc., Alphabet “Form 10-k 2022,” (2022) <https://www.sec.gov/Archives/edgar/data/1652044/000165204423000016/goog-20221231.htm>.
- Joo, Mingyu, Kenneth C Wilbur, Bo Cowgill, and Yi Zhu (2014), “Television advertising and online search,” *Management Science*, 60 (1), 56–73.
- Keller, Kevin Lane (1993), “Conceptualizing, measuring, and managing customer-based brand equity,” *Journal of marketing*, 57 (1), 1–22.
- Kolsarici, Ceren and Demetrios Vakratsas (2018), “Synergistic, antagonistic, and asymmetric media interactions,” *Journal of Advertising*, 47 (3), 282–300.
- Krasnikov, Alexander, Satish Jayachandran, and V. Kumar (2009), “The Impact of Customer Relationship Management Implementation on Cost and Profit Efficiencies: Evidence from the U.S. Commercial Banking Industry,” *Journal of Marketing*, 73, 61–76.
- Linzer, Drew A. and Jeffrey B. Lewis (2011), “poLCA: An R Package for Polytomous Variable Latent Class Analysis,” *Journal of Statistical Software*, 42 (10), 1–29 <https://www.jstatsoft.org/v42/i10/>.
- MarketWatch “Outdoor Advertising Market Size 2023 with Regional Scope, Report Coverage, Historic Data and Outlook 2030,” (2023) <https://www.marketwatch.com/press-release/outdoor-advertising-market-size-2023-with-regional-scope-report-coverage-historic-data-and-outlook-2030-2023-06-06>.
- McAlister, Leigh, Raji Srinivasan, Niket Jindal, and Albert A Cannella (2016), “Advertising effectiveness: the moderating effect of firm strategy,” *Journal of Marketing Research*, 53 (2), 207–224.
- Naik, Prasad A and Kay Peters (2009), “A hierarchical marketing communications model of online and offline media synergies,” *Journal of Interactive Marketing*, 23 (4), 288–299.
- Naik, Prasad A and Kalyan Raman (2003), “Understanding the impact of synergy in multimedia communications,” *Journal of marketing research*, 40 (4), 375–388.
- Schweidel, David A and Wendy W Moe (2014), “Listening in on social media: A joint model of sentiment and venue format choice,” *Journal of marketing research*, 51 (4), 387–402.
- Sethuraman, Raj, Gerard J Tellis, and Richard A Briesch (2011), “How well does advertising work? Generalizations from meta-analysis of brand advertising elasticities,” *Journal of Marketing Research*, 48 (3), 457–471.

- Srinivasan, Shuba, Oliver J Rutz, and Koen Pauwels (2016), "Paths to and off purchase: quantifying the impact of traditional marketing and online consumer activity," *Journal of the Academy of Marketing Science*, 44 (4), 440–453.
- Srinivasan, Shuba, Marc Vanhuele, and Koen Pauwels (2010), "Mind-set metrics in market response models: An integrative approach," *Journal of Marketing Research*, 47 (4), 672–684.
- Statista "Advertising - Worldwide," (2024) <https://www.statista.com/outlook/amo/advertising/worldwide>.
- Stephen, Andrew T and Jeff Galak (2012), "The effects of traditional and social earned media on sales: A study of a microlending marketplace," *Journal of marketing research*, 49 (5), 624–639.
- Tobitt, Charlotte "Print ad revenue halves in six years as three tech giants make USD400bn annually," (2023) <https://pressgazette.co.uk/marketing/global-print-advertising-market-halves-in-six-years-but-publishers-struggling-to-compete-with-online-oligopoly/>.
- Tuckman, Howard P and Cyril F Chang (1991), "A methodology for measuring the financial vulnerability of charitable nonprofit organizations," *Nonprofit and voluntary sector quarterly*, 20 (4), 445–460.
- Vakratsas, Demetrios and Tim Ambler (1999), "How Advertising Works: What Do We Really Know?," *Journal of Marketing*, 63, 26–43.
- Vehtari, Aki, Andrew Gelman, and Jonah Gabry (2017), "Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC," *Statistics and computing*, 27, 1413–1432.
- Vehtari, Aki, Daniel Simpson, Andrew Gelman, Yuling Yao, and Jonah Gabry (2022), "Pareto smoothed importance sampling," *arXiv preprint arXiv:1507.02646*.
- Venkatesan, Rajkumar and Vita Kumar (2004), "A customer lifetime value framework for customer selection and resource allocation strategy," *Journal of marketing*, 68 (4), 106–125.
- Vizzard, Sarah "Adidas: We over-invested in digital advertising," <https://www.marketingweek.com/adidas-marketing-effectiveness> (2019).
- Wedel, Michel and PK Kannan (2016), "Marketing analytics for data-rich environments," *Journal of Marketing*, 80 (6), 97–121.