

# Understanding Video Ads on Social Media Platforms

by

Yiqun Cao

B.S., Carnegie Mellon University (2016)

Submitted to the Department of Management  
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Management

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2022

© Massachusetts Institute of Technology 2022. All rights reserved.

Author .....  
Department of Management  
August 4, 2022

Certified by.....  
Sinan Aral  
David Austin Professor of Management  
Professor of Information Technology and Marketing  
Thesis Supervisor

Accepted by .....  
Catherine Tucker  
Sloan Distinguished Professor of Management  
Professor of Marketing  
Chair, MIT Sloan PhD Program



# Understanding Video Ads on Social Media Platforms

by

Yiqun Cao

Submitted to the Department of Management  
on August 4, 2022, in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy in Management

## Abstract

This dissertation consists of three chapters that explore the effects of creative elements in video ads on social media platforms and investigate how these ads impact consumer behavior in both normal times and times of crisis.

The first chapter performs a comprehensive exploration of video ads on social media platforms using large-scale observational data. Understanding what makes video advertising effective in boosting performance is a crucial task, given firms' heavy investment and its ubiquitous presence in our daily life. It is also a demanding task due to the limited availability of video ad data and the complexity of video features. We first conduct unsupervised clustering to provide a taxonomy of video ad features to understand the commonalities and differences that designers favor when creating video ads. We find videos with much speech tend to have a lower presence of text and a higher presence of people. Second, we perform a feature importance ranking after constructing meaningful and representative creative features. We run multilevel linear models of selected video and campaign features on ad performance outcomes to gain insights into the effectiveness of video elements. We observe that text and its early appearance deteriorate view-related outcome metrics, whereas the presence of people and their early appearance improve view-related metrics. Third, we explore the heterogeneous effects of basic video elements on advertising performance across different platforms, industries, and campaign objectives.

The second chapter follows up on the results from the previous chapter and investigates how algorithmic optimization interacts with whether the firm features the product early or late in the ad. Advertising algorithms seem sophisticated in achieving a single specific objective, such as views and conversions for digital video ads. However, the algorithm's focus on only one goal may present problems for advertisers if they hope the ad could achieve multiple goals such as building awareness, raising interest, and boosting conversions. Using a field experiment, we find that ad algorithms can effectively achieve the prescribed video objective - for example: changing the campaign objective from view to click would significantly improve the probability of clicks while decreasing the short-duration video view probability. The mechanism is that the target audience will change as the campaign goal changes, because opti-

mization algorithms always try to find the "right" audience to maximize the specified performance metrics. We find that the algorithm's single-minded pursuit of a specified objective can be moderated by how quickly content about the product is revealed. Our results suggest that in advertising markets where algorithms are programmed to narrowly fulfill one objective, advertisers need to tailor content to engage users in a way that helps them achieve multiple objectives.

The third chapter investigates how people's response to digital ads changed with fluctuations in the COVID pandemic situation. As new variants keep coming, the world has experienced multiple rounds of COVID-19 virus hits. With the high mutation rate of the virus, people may have to live with COVID-19 and its variants for quite a while. The main research question addresses how people's responses to online ads, in the form of views and conversions when viewing them, change with the ups and downs of the pandemic's severity, specifically its perceived severity as reflected by stay-at-home behavior. We use a difference-in-differences identification strategy and a fixed-effects model. The main results show that ad conversions increase as the pandemic situation becomes more severe. We find the effect stems not just from people replacing offline shopping with online shopping interchangeably, but also from the psychological impact of COVID. Since people are likely to coexist with COVID-19 for the foreseeable future, our research helps firms and businesses better understand consumers' behavior and better adjust to future changes in the pandemic situation through their marketing practices.

Thesis Supervisor: Sinan Aral  
Title: David Austin Professor of Management  
Professor of Information Technology and Marketing

## Acknowledgments

Foremost, I would like to express my sincere gratitude to my advisor Prof. Sinan Aral for his continuous guidance, tremendous support, and amazing resources during my Ph.D. study and research. I also want to thank him for being very understanding. His insights and encouragement keep me motivated.

I want to thank my committee. Prof. Catherine Tucker provided me with tremendous research guidance. I am extremely grateful for her kindness, care and support. Prof. Dean Eckles is always ready to help me with insightful research suggestions and highly constructive feedback. I truly appreciate that.

I have had two fulfilling research internship experiences. My sincere thanks go to the entire VidMob team for providing me with an internship, data, experiments, and multiple research opportunities. I also want to thank Snap for the research opportunity and the incredible internship experience.

I have had a great time at MIT interacting with brilliant faculty members, peers, and friends inside and outside the Marketing Department. I deeply appreciate all the academic and non-academic advice, feedback, and help from them. I would also like to thank the Sloan Ph.D. program office for its continuous help and care.

Last but not least, I want to thank my dear parents, Jiandong Cao and Hong Yu, for their unconditional love and support, and my husband, Xiao Zhang, for his incredible support and encouragement all the time.



# Contents

<b>1 Performance Effects of Video Ad Features: Evidence from a Large-scale Video Platform</b>	<b>17</b>
1.1 Introduction . . . . .	17
1.2 Literature Review . . . . .	20
1.3 Data . . . . .	22
1.3.1 Data Sampling Criteria . . . . .	22
1.3.2 Video Tag Data Description . . . . .	23
1.3.3 Performance Metrics & Other Campaign-Level Data . . . . .	24
1.4 Unsupervised Feature Exploration . . . . .	27
1.5 Feature Importance Analysis . . . . .	28
1.5.1 Feature Construction . . . . .	29
1.5.2 Group Feature Importance Methods . . . . .	30
1.5.3 Group Feature Importance Results . . . . .	31
1.6 Regression Analysis . . . . .	33
1.7 Heterogeneous Effects . . . . .	37
1.7.1 Industry Category . . . . .	37
1.7.2 Media Platform . . . . .	39
1.8 Conclusion . . . . .	39
1.9 Appendix . . . . .	41
1.9.1 Figures . . . . .	41
1.9.2 Tables . . . . .	59

<b>2</b>	<b>Do Algorithms Help Firms Achieve Their Targeting Objectives?</b>	<b>65</b>
2.1	Introduction . . . . .	65
2.2	Literature Review . . . . .	68
2.3	Methodology . . . . .	69
2.3.1	Treatment . . . . .	70
2.3.2	Data Sampling and Identification . . . . .	71
2.3.3	Data Summary . . . . .	72
2.3.4	Data Analysis . . . . .	73
2.4	Main Results . . . . .	75
2.4.1	Possible Mechanism . . . . .	76
2.5	Follow-up Experiments: Age Appeal as an Additional Treatment Variation . . . . .	80
2.6	Robustness Checks . . . . .	81
2.6.1	Additional Advertising Outcomes . . . . .	81
2.6.2	Accounting for Time Variation . . . . .	82
2.6.3	Adding the View-through Campaign Objective . . . . .	82
2.7	Audience & Campaign Characteristics Driving the Differences . . . . .	83
2.7.1	Organization Type . . . . .	83
2.7.2	Age . . . . .	85
2.7.3	Viewing Time . . . . .	85
2.8	Conclusion . . . . .	86
2.9	Appendix . . . . .	89
2.9.1	Figures . . . . .	89
2.9.2	Tables . . . . .	100
<b>3</b>	<b>How Does the Severity of the COVID-19 Pandemic Affect Digital Advertising?</b>	<b>107</b>
3.1	Introduction . . . . .	107
3.2	Literature Review . . . . .	109
3.3	Data Description . . . . .	110



3.3.1	Pinterest Ad Data . . . . .	111
3.3.2	Spectus Mobility Data . . . . .	114
3.3.3	COVID Data . . . . .	115
3.4	Methodology . . . . .	115
3.5	Results & Discussion . . . . .	117
3.5.1	Further Discussion . . . . .	120
3.5.2	Robustness Checks . . . . .	121
3.5.3	Generalizability – Alternative Social Media Platforms . . . . .	123
3.6	Heterogeneous Results . . . . .	125
3.6.1	Effects by Industry . . . . .	125
3.6.2	Effects by Campaign Objective . . . . .	126
3.7	Conclusion . . . . .	127
3.8	Appendix . . . . .	129
3.8.1	Figures . . . . .	129
3.8.2	Tables . . . . .	143



# List of Figures

1-1	Example of Amazon Rekognition bounding box . . . . .	41
1-2	Barplots of proportion by campaign objective, platform and industry	42
1-3	Boxplots of outcome metrics by campaign objective, platform and industry . . . . .	43
1-4	Unsupervised k-means clustering of creative elements' relative duration	44
1-5	Basic video creative element group importance . . . . .	45
1-6	Random forest group importance of face and person video elements on predicting advertising performance . . . . .	46
1-7	Lasso group importance of face and person video elements on predicting advertising performance . . . . .	47
1-8	Fixed-effects model coefficients of specific video elements on advertising performance . . . . .	48
1-9	Fixed-effects model coefficients of basic video features by industry - part 1 . . . . .	49
1-10	Fixed-effects model coefficients of basic video features by industry - part 2 . . . . .	50
1-11	Fixed-effects model coefficients of basic video features by industry - part 3 . . . . .	51
1-12	Fixed-effects model coefficients of basic video features by media platform	52
1-13	Elbow method to determine the optimal number of clusters . . . . .	53
1-14	Robustness check on k-means clustering of creative elements with 4 clusters . . . . .	54

1-15	Robustness check on k-means clustering of creative elements with 6 clusters . . . . .	55
1-16	Random forest sub-group breakdown of basic video elements . . . . .	56
1-17	Lasso sub-group breakdown of basic video elements . . . . .	57
1-18	Coefficients of basic video features by campaign objective . . . . .	58
2-1	Screenshot of Facebook Experiment Platform Setup . . . . .	89
2-2	Screenshot of Published Facebook Ad . . . . .	90
2-3	Predicted Probabilities with 95% C.I. in Different Treatment Conditions	91
2-4	Age Distribution of Campaign Audience and Page Followers . . . . .	92
2-5	Comparison in Audience Age Distribution Between Younger vs. Older Content Tailoring . . . . .	93
2-6	Predicted Probabilities of % Video Plays with 95% C.I. in Different Treatment Conditions . . . . .	94
2-7	Predicted Probabilities with 95% C.I. in Different Treatment Conditions with View-through Optimization as an Additional Campaign Objective Treatment Condition . . . . .	95
2-8	Predicted Probabilities with 95% C.I. in Different Treatment Conditions by Local vs. National Audience Appeal . . . . .	96
2-9	Predicted Probabilities with 95% C.I. in Different Treatment Conditions by Age . . . . .	97
2-10	Hourly 3s View and Click Rates for Three Days (A & B) and Ten Days (C & D) . . . . .	98
2-11	Predicted Probability Comparison Between First Three Days and Remaining Seven Days . . . . .	99
3-1	Summary of the campaign industry and objective percent . . . . .	129
3-2	Time series plots of ad expenditure and impressions . . . . .	130
3-3	Region distribution of Pinterest ad impressions . . . . .	131
3-4	Stay-at-home time series data of example states . . . . .	132
3-5	COVID-related data over time of NY, CA and TX . . . . .	133

3-6	Correlation between stay-at-home and death rate in MA . . . . .	134
3-7	Correlation between stay-at-home and death rate in CA . . . . .	135
3-8	Effects of stay-at-home on ad effectiveness by industry categories . . .	136
3-9	Moderating effects of person creative by industry categories . . . . .	137
3-10	Effects of stay-at-home on ad effectiveness by industry categories . . .	138
3-11	Stay-at-home time series data of all states . . . . .	139
3-12	COVID-19 fatality time series data of all states . . . . .	140
3-13	Vaccination rate time series data of all states . . . . .	141
3-14	Aggregated 2s view rate time series data of example states – parallel trend . . . . .	142



# List of Tables

1.1	Table of basic video creative elements summary statistics . . . . .	59
1.2	Table of basic campaign feature and outcome metrics summary statistics	60
1.3	Fixed-Effects Model with Basic Video Elements . . . . .	61
1.4	Fixed-Effects Model with Basic Video Elements - Relative Appearance Duration . . . . .	62
1.5	Fixed-Effects Model with Basic Video Elements - Start Time . . . . .	63
1.6	Fixed-Effects Model with Basic Video Elements - Location . . . . .	64
2.1	Summary Statistics of Outcome Metrics: . . . . .	100
2.2	Table of Main Binomial Regression Results . . . . .	101
2.3	Marginal Effects of Changing One Treatment Variable While Keeping the Other Treatment Variable Constant . . . . .	102
2.4	Chi-square Test Results of Demographic Age and Gender Distributions Between Older vs. Younger Appeal for Each Treatment Group . . . . .	103
2.5	Table of Marginal Effects When Changing from Early Informational Content to Late by Organization’s Audience Appeal . . . . .	104
2.6	Binomial Regression Coefficients of 3-second Video Plays and Link Clicks	105
3.1	Summary of campaign metrics . . . . .	143
3.2	Table of ad campaign performance metric summary by industry . . . . .	144
3.3	Table of ad campaign performance metric summary by campaign ob- jective . . . . .	145
3.4	Table of main ad creative presence summary by industry . . . . .	146
3.5	Fixed effect model of the COVID indicators on stay-at-home . . . . .	147

3.6	Main results for 2s view outcome . . . . .	148
3.7	Main results for conversion outcome . . . . .	149
3.8	Main results for firm-defined conversion outcome . . . . .	150
3.9	Table of 2s views – segmenting the pandemic duration . . . . .	151
3.10	Table of conversions – segmenting the pandemic duration . . . . .	152
3.11	Table of firm-defined conversions – segmenting the pandemic duration	153
3.12	Table of alternative independent variables on ad conversions . . . . .	154
3.13	Table of conversions – Facebook platform . . . . .	155
3.14	Table of alternative independent variables on 2s views . . . . .	156
3.15	Robustness check – using DMA as location unit instead of state . . .	157
3.16	Robustness Check – additional video feature covariates . . . . .	158
3.17	Table of 3s views – Facebook platform . . . . .	159



# Chapter 1

## Performance Effects of Video Ad Features: Evidence from a Large-scale Video Platform

### 1.1 Introduction

Advertising in video format has become a prevalent practice on social media platforms. Video-based advertising, containing audio and dynamic visual elements, has been found to perform better in drawing attention [12, 30] than text- or image-based advertising. Video ad expenditure in the US has increased from \$8.92 billion in 2017 to \$10.18 billion in 2019, and spending is expected to reach \$12.66 billion by 2024 [61]. It is also projected that, by the end of 2022, 82% of consumer internet traffic will be from online videos [51]. Due to its growing dominance and proven effectiveness in advertising industries, businesses are eager to understand how different video features affect people's willingness to watch and purchase.

Understanding video is a challenging task due to its high dimensionality. For example, machine learning papers tend to use pixel color values to define the raw input features in video analytics. However, the use of pixels sacrifices interpretability to a large extent. Now, some tools could make use of machine learning techniques

to extract some interpretable features. Take Amazon Rekognition as an example. It provides moment-to-moment video descriptions including the presence of people and facial expressions, gaze directions, the approximation to celebrity, approximate age, gender, text, and speech, with a detailed and comprehensive account such as the area, the positions, the confidence, and the value of these features at each 0.1-second interval of the video. The descriptive data again provides an extremely rich set of video information for each video ad asset. The sheer amount of information can be overwhelming for drawing meaningful insights because of the high dimensionality. Our paper would like to simplify, construct and comprehend the rich set of video creative elements given the data provided by the aforementioned recognition system.

We are not aware of any existing publication that has given such an extensive and interpretable high-level overview of video ad features and their possible relationships with ad performance. We want to make the first attempt using large-scale empirical data. We collaborated with a video creative ad firm that serves ads on multiple online social media platforms, including Facebook, Google, Twitter, Snapchat, Pinterest, and LinkedIn. We sampled some of their video ads that ran in 2019. Overall, we analyzed 70,006 different video ads, which ran 90,015 ad campaigns on the six major social media platforms. Based on monthly active users, five of the six platforms are ranked in the 15 most popular social networking sites or apps worldwide as of October 2020 [16]. The impressions (i.e., the total number of times the video ads were shown to consumers in these campaigns) add up to around 1 trillion. The advertised products or services span 22 industries, including technology, entertainment, financial, food, retail, transportation, nonprofit, and health, covering almost every aspect of our daily lives.

The first part of the paper makes an initial exploration of the basic features provided in the data to gain a better understanding of video creative characterization and correlation. In particular, we use k-means unsupervised clustering to categorize videos purely based on their creative design, regardless of their advertising performance. We would like to understand any common patterns when creative designers come to create these videos. Furthermore, we provide a broad classification of videos

based on the commonly used creative elements.

The second part of the paper serves to characterize specific features, explore their importance individually and collectively, and understand the association between features and ad performance. We first choose two different methods to conduct our feature selection. We use lasso to look predominately at the importance of the regularized linear relationship. We also use random forest to better account for the non-linear relationship between our features. Since each creative element has multiple features relating to its confidence, size, positions, and variations over time, we mainly use group lasso and group random forest to determine the collective importance of each creative element, which could also provide some insights for our analysis. Overall, we find that campaign-related features, industry, and video-summarizing features are most important for determining outcomes. Among video creative elements, we find basic creative elements such as text, people, and speech are most important. Comparatively, face and celebrity are less critical. Among specific creative elements, females tend to be more important than males. The middle gaze is generally more important than other gaze directions. There is much more variation in emotion importance across different outcomes. Based on importance analysis insights, we apply fixed-effects models to assess the effects of the video creative elements on advertising performance. We start with the primary elements and the most basic features, i.e., indicators of the presence of each element. Some interesting findings are that speech and the presence of people are preferred, whereas text is not favored when it comes to view-related metrics. Meanwhile, longer duration and the earlier appearance of faces and celebrities is favored. The effect of a logo being present is not significant. We also further investigate the effects of specific video elements, including facial emotion, gaze direction, gender, and age group. We find the presence of females and young adults is consistently favored.

The third and last goal of the paper is to explore heterogeneity based on different platforms and industry categories. The intended audience, the placement of ad videos, and the interface of each social media platform can vary a lot. For example, Facebook commonly presents sponsored ads in the form of a news feed. People see them as they

scroll down the screen. YouTube presents video ads within organic YouTube videos. The audience age distribution of each platform also differs. For example, LinkedIn is primarily for working professionals, and we find that text has a positive effect on LinkedIn. Industry categories also bring about distinctive styles of video ads, as the nature of the products and the appropriate audience groups again differ drastically. For example, we find celebrity has a negative effect on education industry ads, whereas text has a positive effect in the professional services industry.

A necessary clarification is that our paper does not intend to make causal claims about the effect of each creative element of video ads. Instead, the primary purpose of this paper is to provide a comprehensive characterization of video creatives, illustrate the importance and possible effectiveness of each creative element, and explore the heterogeneous effects by differing campaign setup and ad nature. Our research tries to help better guide the research in video ad creatives with associational evidence and large-scale generalizable data. We hope our extensive video feature exploration inspires subsequent research with a more causal focus.

## 1.2 Literature Review

One stream of the literature relates to the characterization of advertising creatives. Pictures, text, brand, and headline are four major regions of interest that were found to draw people’s attention in traditional print ads [70]. In some eye-tracking lab settings, it has been found that the brand element receives the most eye fixations per unit of size, followed by the text and then the pictorial [69]. The advertising literature stresses the significance of both visual and verbal characteristics in dual role theory [56]. Recently emerging studies on video advertising also present multiple ways of characterizing features. In an attempt to automatically code political campaign advertisement videos, features are categorized as image summarization, text, face, speech, and music [32]. The existing ad design creative literature provides us with the foundation of feature construction. Meanwhile, our paper still makes a unique contribution to video advertising by covering a large variety of products and services.

Following feature characterization, the directly related stream of literature assesses the effectiveness of features in advertising. The brand is one type of advertising element that exists in every ad and has been studied widely. Research into TV commercials has found that the presence of a brand-differentiating message in a commercial causes a significant decrease in zapping probabilities [60]. Multiple research studies have proved that faces are also particularly effective in attracting people’s attention [37, 28, 40, 6]. Static or dynamic graphics are another typical ad element. In a comedy movie clip, happy, happier, happiest is the best order for scenes. It is found that a positive trend in happiness results in higher watching intentions and more significant success at the box office [41]. Such pictorial elements frequently involve humans. A direct mail field experiment in South Africa conducted by a consumer lender found that including a photo of an attractive woman increased loan demand by about the same amount as reducing the interest rate by 25% [8]. In TV commercials, there is also a positive association between auditory energy levels in ads and ad-watching behavior [72]. Moreover, background music has been found to affect audience moods and purchase intentions [3]. In the most recent format – skippable video advertising – high-arousal stimuli increase ad effectiveness in ad-congruent contexts, and product involvement affects the intrusiveness of high- and low-arousal skippable ads [7]. It is also found that entertaining ad content mitigates the negative impact of the transition from unskippable to skippable ads, given that consumers have a negative prior. Meanwhile, more informational ad content is preferred if the consumer has a negative prior about the advertiser’s product, whereas less information is preferred with a positive prior [20]. In influencers’ videos, when products are advertised in the most engaging part, sales improve [71]. Overall, many video features have their own literature. Despite much research devoted to studying the effectiveness of particular advertising or video features, to our knowledge, our paper is the first to provide a comprehensive overview of the common video ad creatives and their importance in determining ad performance. We hope to fill this gap by running a large-scale feature importance ranking and regression analysis.

The main contribution of our research lies in a thorough exploration of the cre-

atives’ importance and effectiveness for ad performance. We have found literature that provides an overview of the effectiveness of video creative elements in comedy movie trailers via an online experiment [41]. However, the video marketing practice of comedy trailers can be greatly different from video ads in other industries, and both the target audience and people’s reactions might also differ in experiments versus when they use social media. There is plenty of literature that provides insights on print ads [69, 52, 53, 70]. However, the design of print ads is still fundamentally different from video ads, and reactions gathered from eye-tracking are better at diagnosing people’s attention levels than understanding people’s lower-funnel conversion intentions. So far, we have not seen any literature that provides such a broad overview of video ad creatives. This seems to be an open area to which we can contribute.

## 1.3 Data

Our dataset is obtained from an anonymous advertising firm that specializes in video ad creation. The firm provided us with de-identified video data processed by Amazon Rekognition, which produces various video tag features, along with the videos’ corresponding performance metrics on Facebook, a widely used social media platform.

### 1.3.1 Data Sampling Criteria

The anonymous ad analytics firm provided us with advertising campaign data from Jan 1, 2019, to Dec 31, 2019. When sampling the ad campaigns, we required the video duration to be at least 3 seconds and no more than 120 seconds. We specified the minimum 3s criterion because one of the main performance metrics we focus on in our analysis is the 3-second play rate. For those video ads lasting less than 3s, the 3s view rate would have a different interpretation and thus would no longer be accurate. We limited the video ads to a maximum two-minute duration because 1) more than 98% of the video data in the firm’s 2019 database was for videos lasting  $\leq 120$ s, and 2) longer-duration video ads could be as long as an hour, for which completely different video creative strategies might be needed compared to short video ads. Moreover, we

only included video ads with at least 1\$ spent and 50 impressions. We believe any ad campaign with fewer than 50 impressions or less than 1\$ spent is more likely to be a trial run than a formal advertising campaign.

### 1.3.2 Video Tag Data Description

The descriptive video data is the output of Amazon Rekognition tag data, which describes everything that can be detected by computer vision in the video. The main topics of the tag data are person, face, text, speech, and logo. Apart from speech, all the other basic features have bounding box data to proxy their sizes. Fig. 1-1 gives an example of a bounding box produced by the Amazon Rekognition system. Whenever its computer vision algorithms detect a face appearing in the video, it draws a tight rectangular shape based on the outer edges of the face. The bounding box area gives an approximate percentage for the face size compared to the entire video screen size. Similarly, Amazon Rekognition also detects any person, text, and logo appearing in the video and produces bounding boxes for them in each frame. Apart from bounding box data for the person and text categories, we also have the number of people and the speed of text and speech in each frame.

Furthermore, the person and face categories have a rich set of sub-category features, including the predicted gender, age range, gaze direction, and facial expression of the detected face, as well as how likely it is that the detected person is a celebrity. For each of these sub-category features, we also have their appearance timing and confidence scores. In the same example in Fig. 1-1, we see that Amazon Rekognition gives a confidence score of 100, indicating that the detected person is 100% likely to be Jeff Bezos.

Table 1.1 shows the summary statistics of the basic video features. The first column is the proportion of videos in the sample in which each feature was detected. The rest of the columns are the summary statistics including minimum, 1st quantile, median, mean, 3rd quantile and maximum of the duration of appearance of each feature relative to the entire video duration. We find that 90% of the videos have at least some text in them, and the average amount of time that text appears for in

each video is around 74%. This makes sense because, in many short online video ads, the main messages are delivered in the form of text. Moreover, people also appear in many videos (around 78%), and the average duration for which they appear in each video is also greater than half of the time. In contrast, speech, celebrities and logos appear less commonly. However, for features like celebrity and logo, we are still not quite sure whether that is due to algorithmic detection limitations.

### 1.3.3 Performance Metrics & Other Campaign-Level Data

In addition to video creative data, we also have Facebook aggregated performance metrics and associated campaign-level determinants (shown in Table 1.2 and Fig. 1-2).

Table 1.2 gives the summary statistics of the key campaign factors – impressions and campaign expenditure – as well as the key outcome metrics we consider: 2s or 3s video play rate, through-play rate, click rate and conversion rate. All of these data are left-skewed with mean values higher than median. Meanwhile, we notice that both click rates and conversion rates are really low, with mean values equal to 0.01, and median values of 0.01 and 0.00.

The video duration in Table 1.2 represents how long each video was in seconds, and the campaign duration refers to how many full days each campaign ran for. Impression refers to the number of times this video was shown to users within this campaign. Impressions may include multiple views of ads by the same people. Spending refers to the total amount of money the campaign cost.

Two- or three-second video plays and through-play reflect people’s interest in the video ads. Two-second or three-second video plays count the number of instances when the user viewed the video ad continuously for at least 2 or 3 seconds. Whether the data is for 2s video play or 3s video play depends on what was available for each platform and what our collaboration firm was able to pull. We have 3s view data for Facebook and Twitter, and 2s view data for Snapchat, Pinterest and LinkedIn. We do not have YouTube’s 2s or 3s view data, so we do not include it when analyzing the 2s or 3s view rates. Through-plays count the number of instances when the user watched



through the entire video. Clicks refer to clicks on any area of the ad that links to destinations or experiences for the ad. In the case of Snapchat, it refers to the swipe action that serves a similar purpose to clicks on other platforms. Our definition of conversion, however, is less accurate than the other three. It refers to the maximum of any one of the lower-funnel actions such as app installs or activation, checkouts, add-to-carts and purchases. These are actions that could bring direct revenue to the firms. The reason it's less accurate is that not all campaigns set up conversion tracking. Some campaigns may attract purchases but report 0 conversion due to no tracking. The second reason is that even with tracking, actions like purchases may still not be accurately tracked because users might not make the purchase in the same session, or may visit offline stores. Third, the conversion action for different products or services in different industries can vary a lot and requires different levels of effort. For example, an in-store purchase is much more time-consuming and expensive than activating a free app. Fourth, on some platforms, conversion is defined by the advertiser, so the criteria for conversion might change. All the outcome metrics are in the form of a "rate". The rate is calculated with respect to the corresponding impression value. For example, the 3s play rate of a video is calculated by the number of times a video ad has been viewed for more than three seconds (i.e., 3s plays) divided by the number of impressions.

Finally, in Fig. 1-2 we present the percentage of objectives that the ad campaigns chose, the platforms they ran on and the industries they belonged to. For industries, we only present those segments representing more than 1% in our data sample. Around two-thirds of the ad campaigns adopted either conversion or app installs as the campaign objective. These are relatively lower-funnel campaign objectives compared to others like brand awareness, reach and video views. Different campaign objectives might produce very different ad performance outcomes.

In Fig. 1-3, we use a boxplot to illustrate the differences in objective on the top-left. The four different colors represent the four different performance metrics we measure. The campaign objectives are arranged in the order of the corresponding marketing funnel, from high to low. Although different platforms might have different

names for campaign objectives, and categorize them more specifically, we re-classify each platform's campaign objectives into five main categories to get a more intuitive comparison. In general, the outcomes are consistent with the funnel stage of the campaign objectives. For example, the higher-funnel campaign objectives of video views and brand awareness have very similar distributions for all the outcomes. The lower-funnel campaign objectives tend to have higher lower-funnel metric measurements. Each campaign objective did well in optimizing its own goal. For example, the video view objective has the highest 2s/3s view rate and through-play rate, the click objective has the highest click rate and the conversion objective has the highest conversion rate. However, the difference in conversion rate might also be due to the limited ability to accurately represent and track lower-funnel actions. However, if we compare the same outcome among ads with different funnel-stage campaign objectives, the distributions can differ significantly. The reason for such differences is the different targeting strategies as brought by the corresponding optimization objective and subsequently very different audience groups.

The plot on the right of Fig. 1-3 illustrates how each metric differs by platform. The metrics are all in log form, to better enable comparison, as their original distributions are highly skewed. We do not have a two- or three-second view rate for YouTube, so we only report three metrics for the platform. YouTube seems to have the highest through-play rate compared to all the other platforms. However, its click rate is not significantly higher compared to others. Pinterest has the highest two- or three-second view rate.

In the bottom plot of Fig. 1-3, we use a boxplot to illustrate industry differences. We select several industry categories that represent at least 1% of the entire data. The apps & websites and automotive industries have the highest short-duration view rates. There is more variation in click rate and conversion rate compared to view rate. As explained earlier, we cannot be sure whether the difference in conversion rate is due to differences in industry or differences in conversion metric measurement for each industry. Overall, the common pattern for each platform is two-/three-second view rate > through-play rate > click rate > conversion rate, regardless of campaign

objectives, media platforms or industries.

## 1.4 Unsupervised Feature Exploration

After getting an overall picture of the basic video features and campaign features, we wanted to learn more about the design patterns and possible classification of these videos. At this stage, we were just interested in video classification without consideration of their campaign setup and performance. We chose k-means to perform the unsupervised video classification because k-means is a common and interpretable practice for unsupervised clustering.

When running k-means, we only included the video creative elements' relative duration (the appearance duration of a creative element relative to the entire video duration). We excluded the campaign variables because our main purpose in this unsupervised clustering was to get a picture of video clustering and adding campaign variables could have interfered severely with the clustering process. We used relative appearance duration for the basic video creative elements including people, face, celebrity, text, speech and logo, the relative appearance duration of varying demographics including gender and age, and the relative duration of varying gaze directions. We included smile as the only emotion variable because the detection of other emotions tends to have very low accuracy with low confidence. Adding emotion might have added too much noise into the clustering process. These variables were easy to understand and were standardized with a fixed range from 0 to 1. Since there was no outcome variable for cross-validation, we adopted the elbow method [66] and determined the optimal number of clusters for k-means clustering to be five, as shown in Fig. 1-13.

In Fig. 1-4, we present the clustering results of all these variables. Each color indicates a different cluster, and there are five clusters. Each point represents the mean of the relative duration for a creative element in the corresponding cluster. The size of the point can be seen as a proxy for the within-cluster standard deviation value. The bottom plot summarizes the clustering of the basic creative elements,

the middle plot summarizes the people’s age and gender elements, and the top plot summarizes the gaze directions and smile elements. We observe that logo appearance duration is generally low, and text appearance duration is generally high regardless of clusters. In terms of speech duration, the cluster with the highest relative speech duration (cluster 2 - yellowish-green color) has the lowest text duration, but also a relatively high person and face appearance duration. This implies that video ads need to have more speech- and person-related elements to compensate for the choice of less text. Finally, the top three clusters of person appearance duration are also the top three in face and celebrity appearance duration. This is intuitive as these three elements tend to be very correlated. Meanwhile, the top cluster for people, face and celebrity (cluster 5) is also the top cluster for early adult, female, smile and middle gaze, implying a common design pattern adopted with the appearance of people in video ads, i.e., young female adults looking straight at the audience with a smile. In addition, we also observe the generally minimal presence of senior people and upward gaze directions in video ads. These observations are robust to initial seeding and small variation in cluster numbers as we tried multiple random seedings and different numbers of clusters, including four and six clusters as shown in Figs. 1-14 and 1-15.

## 1.5 Feature Importance Analysis

The previous unsupervised clustering section provides a high-level view of the common creative strategies designers would like to adopt. Naturally, the next step is to explore whether these strategies are important and effective in achieving their advertising goals. In this section, we conduct feature importance analysis by first constructing the features that we use in our prediction model. We summarize the features from the output of the Amazon Rekognition system so that the number becomes manageable and the features are interpretable. Second, we further group features of the same creative element category or sub-category and perform group lasso and random forest to rank group feature importance.

### 1.5.1 Feature Construction

We construct most features based on the algorithmically detected timing, size, location and confidence, as well as features specific to each video element. In general, for the basic video element categories like text, person, face, celebrity and logo, we include the first appearance (start time), last appearance (end time), appearance duration (element duration), average size, and variation of size (standard deviation) throughout a video. In terms of the location, we divide the screen into  $3 \times 3$  zones, namely the top-left, top-center, top-right, left, center, right, bottom-left, bottom-center, and bottom-right zones, and calculate the percentage of time for which the specified video element appears in each zone. Depending on the size of the bounding box, an element can appear in more than one zone or even all nine zones at a given time. Meanwhile, some basic video elements have additional features. Text has the additional features of speed and average visibility, logo has average clarity, and celebrity has average confidence. For the person element, we also include the number of people. Moreover, speech is also a basic element without visuals, so it has fewer features: duration, start time, end time and speed.

For the more specific video elements like emotion, gaze direction, age group and gender, we include the first appearance (start time), last appearance (end time), appearance duration, average confidence and standard deviation of confidence. We do not include size because they either overlap or are highly correlated with the face and person sizes. In addition, we also include the total video duration and number of labels in the video.

The grouping of video creative features is fairly intuitive. For basic video elements like text, speech, person, face, celebrity and logo, we include all their corresponding subfeatures like location, size and timing as well as element-specific sub-features as one group. They give an overall description of these elements' appearance. For the specific video elements, we group all confidence and timing features based on their sub-elements like the exact gaze direction or emotion. Finally, we only include an "other features" group for video features. This group includes features that are

difficult to categorize into video creative element groups. Specifically, they are the total duration of the video, the total number of labels detected in the video by Amazon Rekognition, and the industry to which the advertised product/service belongs. The reason we include industry as part of video features is that different categories of products or services also significantly affect the overall creative design in video ads.

In addition to video features, we also constructed or included all the necessary campaign-related features. In terms of campaign features, we constructed the campaign setup group with variables like campaign schedule, campaign objective, impressions, expenditure, and the campaign audience group, like the percentage of female, male and different age groups whenever available.

### 1.5.2 Group Feature Importance Methods

We mainly use two very different types of feature selection techniques. The first is group lasso [67, 73, 46], which solves the following optimization problem

$$\min_{\beta \in R^p} \left( \|y - \sum_{l=1}^L X_l \beta_l\|_2^2 + \lambda \sum_{l=1}^L \sqrt{p_l} \|\beta_l\|_2 \right)$$

where  $l$  indicates group,  $p_l$  is the number of predictors in group  $l$ , and  $\|\cdot\|_2$  is the Euclidean norm. We apply the randomized lasso based on the stability selection method [47]. Based on stability selection, the feature importance is calculated by random bootstrapping the sample many times, each time multiplying each variable by a random multiplier between 0 and 1, selecting the optimal lambda using cross-validation, running group lasso and then counting the number of times each feature gets selected by lasso regression. The importance is obtained from the number of selected times divided by the total number of times re-sampled.

The second approach to assess importance is to use random forest [9], which can explore non-linear relationships. It is a type of ensemble learning and consists of a number of decision trees. The random forest feature importance is produced based on permutation importance [9]. To assess the group importance, we use a group variable importance measure that extends the permutation importance to groups of variables

[27]. Let  $X^T = (X_1, X_2, \dots, X_p)$  be a random vector, and  $J = (j_1, \dots, j_k)$  be a  $k$ -tuple of increasing indices in  $1, \dots, p$  with  $k \leq p$ . The random forest group importance of a sub-vector  $X_J = (X_{j_1}, X_{j_2}, \dots, X_{j_k})^T$  is defined as

$$l(X_J) := E \left[ (Y - f(X_{(J)}))^2 \right] - E[(Y - f(X))^2]$$

where

$$X_{(J)} = (X_1, \dots, X'_{j_1}, X_{j_1+1}, \dots, X'_{j_2}, X_{j_2+1}, \dots, X'_{j_l}, X_{j_l+1}, \dots, X_p)^T$$

such that  $X'_j = (X'_{j_1}, X'_{j_2}, \dots, X'_{j_k})^T$  is a replicate of  $X_J$ , which is independent of  $X_J$ ,  $Y$  and all other predictors. An essential clarification about our feature selection approach is that a feature being important does not necessarily mean it will contribute positively to the outcome. The feature might also make the ad performance significantly worse.

### 1.5.3 Group Feature Importance Results

Fig. 1-5 illustrates the group importance of basic video creative elements and campaign groups when predicting each outcome metric (in log form) respectively using lasso (left figure) and random forest (right figure). The feature importance when predicting each performance metric is displayed in different colors. The round point shape refers to video features, and the triangles refer to campaign-related features. The size of a point indicates the number of individual features falling into each feature group. The higher the importance score, the more important the group of features is in predicting each advertising outcome. In general, we observe that the basic element appearance groups like text, speech, logo, person and celebrity have lower importance than campaign and industry groups regardless of outcome metrics. Conversion has the highest feature importance among the four outcome metrics, regardless of feature groups. Moreover, among all the basic video creative elements, text appearance in general has the highest importance, followed by person appearance. Celebrity appearance generally has low importance. This might be explained by the low celebrity

detection accuracy of the Amazon Rekognition system.

We also look at the importance breakdown of these basic video creative elements in Figs. 1-16 and 1-17. Instead of grouping all the time, location, size and any other relevant features together into one group for each element, we make more, smaller groups to look at how each type of feature might have different importance levels in determining the outcome metrics. The type time includes the start time, end time and appearance duration relative to the entire video duration. The type size includes average size and variation in size. The type location includes the percentage of appearance in all nine zones on the screen. The type confidence generally includes average confidence and variation. For text and logo, the additional average clarity also falls into the confidence type. The type speed refers to the total words per second. Figs. 1-16 and 1-17 show the importance of these subgroups of basic video creative elements using random forest and lasso respectively. The size of the points in the plots represents the number of features, and the color represents different basic creative elements. In both figures, we observe that text confidence and speed are the two most important features when predicting all four outcome metrics. Among logo subgroups, logo time is the most important. As for celebrity, the location subgroup is generally less important. Other than that, there are not many consistent patterns.

Despite the lower importance of basic creative elements compared to campaign groups in Fig. 1-5, they are still comparatively more important than the specific video creative elements in Figs. 1-7 and 1-6 in terms of importance magnitude. There are more variations in these specific creative elements. Figs. 1-7 and 1-6 show the importance of these sub-elements using both feature importance methods. Each group has five features: two describe the average confidence and variation in confidence, and the other three describe the start time, end time, and appearance duration. Each color represents a different category, including emotion, gaze direction, age group and gender. Across both figures and the four outcome metrics, we observe some common patterns that the presence of females is generally more important than that of males. Middle gaze direction is more important than other gaze directions. In Fig. 1-7, we observe that either smile or happy emotion is often one of the most important



emotion groups when using lasso for prediction. We observe that young adult is consistently the most important age group for all metric predictions in Fig. 1-6 using random forest, and for both view-related metric predictions in Fig. 1-7 using lasso. Meanwhile, the senior age group is the least important age group for all outcomes when using random forest. This might be because senior people are generally not targeted by online ads.

To sum up, we find that campaign-related feature groups, and video summary features like media duration, number of objects and industry, have higher overall importance than video creative elements. Basic video elements such as text and person have higher importance than specific video elements that describe the people- and face-related elements such as age, gender, emotion and gaze direction. Among the basic video elements, text is the most important, followed by person. Among specific elements, the middle gaze direction is commonly the most important of all the gaze directions. Female is more important than male. Young adults are more important than other age groups in predicting view metrics. In terms of subgroups such as the time, confidence and location feature groups, we do not observe very consistent patterns. Lastly, the importance we discuss here refers to how much influence the group has in affecting advertising performance, but the impact is not guaranteed to be positive. In the next section, we try to assess whether the effect is positive or negative and exactly how large it is.

## 1.6 Regression Analysis

We conduct the regression analysis in this section, aiming to learn about the effects of video features on ad performance. We adopt the fixed-effects model specification that takes fixed effects on partner ID, industry, week, platform and campaign objective combination.

$$y_{ijklt} = \alpha_j + \delta_k + \gamma_l + \delta_t + \beta^T x_{ijklt} + \epsilon_{ijklt}$$

The partner ID  $j$  represents each different advertiser account. Each advertiser account might run multiple advertising campaigns for the same or similar products or services

on more than one platform. There is a commonality for ad campaigns run by the same advertiser due to similar products/services and similar campaign strategies. Similarly, the products or services falling into the same industry  $k$  might also share common traits. We include time fixed effect  $t$  indicating the week in the year 2019 to account for the seasonal effect of advertising campaigns. Meanwhile, the same ad might also be run multiple times with different campaign objectives on different media platforms  $l$ . We believe the combination of campaign objective and platform will also make a difference in intercepts due to varying algorithmic recommendation effectiveness on different platforms under different objectives.

Given we have constructed hundreds of features with a lot of collinearity between them, the first step is to get an overall and simplified picture of the most basic video creative elements. With reference to insights from the feature selection section, we decided to include only the simplest variables, dummy indicators of whether text, person, face, celebrity, logo and speech elements were present in each video in our regression model, plus the other campaign and video features including video duration, number of labels, ad industry, campaign setup, campaign audience and objective. Table 1.3 shows the main effects we are interested in, i.e., whether the presence of these basic video features is significantly associated with ad performance. Each column presents the regression coefficients with respect to different outcome metrics. As defined in the Data section, we used the 3s view rate, through-play rate, click rate and conversion rate, all in log forms, to measure each ad’s performance. These four metrics approximately cover users’ response to the ad from high-funnel lead generation, i.e., the 3s view rate in Model (1) and through-play rate in Model (2), to the mid-funnel lead-nurture click rate in Model (3), down to low-funnel sales (i.e., conversions) in Model (4). We were able to gain a clearer picture of the entire consumer journey based on these four outcome metrics. All the outcome metrics are in log form when running regressions because their original values tend to be skewed. Moreover, we assume heteroskedasticity and hence use cluster-robust standard errors.

In Table 1.3, we could interpret the first value in Model (1) ( $-0.058$ ) as the presence of text in a video ad being associated with around a 5.6% decrease in the

2-second/3-second video play rates, and this effect is statistically significant. The presence of speech is significantly positively associated with an increase in 2s/3s play rates. The presence of people is also positively associated with 2s/3s play rate metrics with statistical significance. Both the face and celebrity indicators are a subset of the person indicator. There has to be a person for a face or a celebrity to be detected. We observe that given there are people in the video, face and celebrity are in general no longer significant. The presence of a logo has no significant effects either. Looking at the table in general, we observe that the presence of creative elements only has a significant effect on the highest-funnel metric of 2s/3s video plays and not much effect on others. Besides the basic video creative elements, we include the general video features and campaign features as controls. In Table 1.3, we observe that many of them have a significant effect on the advertising results. Video duration has a significantly negative effect on through-plays but a significantly positive effect on the click rate. The adverse effect on through-plays might be because the long video duration is more likely to make people lose patience to finish watching the entire video. Meanwhile, a longer video might be better at providing more information and enticing people to learn more, leading to a positive effect on clicks. High impressions are associated with lower performance metrics, whereas higher spending is associated with better performance.

In Table 1.4, we include more video creative features in the fixed-effects model. On top of the basic video creative element appearance indicators and all the other control variables, we further include the relative appearance duration of each basic video element. The first value in Model (2) ( $-0.105$ ) shows that a 10% increase in the appearance duration of people relative to the video duration when text is present in the video is associated with around a 1.0% significant decrease in the through-play rate. We observe that increased duration of face and celebrity produces significantly positive effects for the click rate and through-play rate, respectively. A 10% speech duration increase is also associated with a 1.1% increase in the through-play and conversion rates. There is not much of a significant effect observed for person and logo duration.

In Table 1.5, we replace the duration variables with the start time variables in fixed-effects models. We assess the effects of the start time of the basic video elements, if these elements appear in the video. The first value in Model (1) (0.272) is interpreted as a 10% increase in text start time in videos where people are present, i.e., a 10% delay in text start time relative to the video duration brings a 3.13% increase in the 2s/3s play rates. Meanwhile, the earlier appearance of people is strongly favorable as a 10% delay in appearance start time is accompanied by a 1.01% decrease in the 2s/3s play rates and a 0.76% decrease in the through-play rate. In addition, the early appearance of celebrities is also preferred as it is associated with better click rates. Lastly, the effects of a delay in speech start time are mixed with reduced through-play rates and increased click rates.

Next, we explore the more specific elements like emotion and gaze direction by including either the relative duration or the average confidence in the basic fixed-effects models. Based on the feature importance results, we found that timing features are predominantly important for gaze direction, gender and age group. We used the percentage of duration compared to the entire video duration as these elements' variables. For the emotion groups, given the feature importance results do not give a definite answer and we found confidence has a higher variance with certain emotions having low accuracy, we used the average confidence variable for the emotion elements.

Fig. 1-8 shows the coefficients of these detailed creative element groups when included in addition to the basic fixed-effects model. The top two plots show the effects of gender and age. Female appearance has a significantly positive effect on the 2s/3s play rates and click rate. In terms of age groups, we observe that a longer appearance of teenagers contributes to higher 2s/3s play rates. Moreover, a longer duration of mature adults contributes to higher 2s/3s play rates and click rate. Lastly, a longer appearance of seniors has a significantly negative effect on the 2s/3s video play rates. These effects might be due to the nature of the advertised products' or services' target audience. The plot in the middle shows the coefficients of gaze direction. We observe that the down-right, mid-left, middle and mid-right gaze directions all lead to higher 2s/3s play rates. The middle and up-right gaze directions are associated

with higher click rates. Moreover, the down gaze direction is negatively associated with the through-play rate and the up-left gaze direction is negatively associated with the click rate. The last plot illustrates the effects of emotions' average confidence. Higher confidence in calmness and surprise are associated with negative click rates, while higher confidence in surprise is associated with a positive through-play rate. Due to the limitations in emotion detection, and the fact the results do not produce obvious patterns, we do not attempt to draw too many implications.

Overall, the results imply that we should consider having longer and earlier appearances of people to retain the audience's attention. Celebrity is encouraged, and its earlier appearance also brings better results. The effect of speech is uncertain as it produces opposite effects on different metrics. In consideration of better 2s/3s play rates, we should avoid placing text at the bottom or early on, or placing people on the very right or left sides. The results from the specific video elements indicate female is favored. Among all the gaze directions, middle gaze directions generally produce positive contributions to 2s/3s views.

## 1.7 Heterogeneous Effects

We understand that different platforms have varying interface designs, algorithmic targeting capabilities, and user bases that produce drastically different results. Moreover, different industry categories advertise for distinct products or services that result in differing ad designs. In the data exploration section in Fig. 1-3, we find that the distributions of outcome metrics vary by these categories. In this section, we'd like to explore how the effects of video creative features might differ by industry and media platform. In Fig. 1-18, we also explore how the effects of video creatives differ by campaign objectives.

### 1.7.1 Industry Category

In Fig. 1-9, we compared the effects of text and speech in video ads for each industry category in our sample. The appearance of text has a negative effect on one or two

outcome metrics in many industry categories including education, financial services, health and wellness, home improvement, luxury goods, retail, restaurants, technology, transportation and travel (in the top plot of Fig. 1-9). Among them, the luxury goods industry incurs the biggest negative effects in both magnitude and number of outcome metrics, i.e., three out of four are positive. The luxury goods industry tends to be highly priced, and the ad visuals and product designs play a powerful role in determining people's willingness to watch or learn more. On the contrary, professional services is the only industry in which text contributes to significant positive effects on three of the four metrics. This might be related to the fact that professional services is a knowledge-intensive industry segment that focuses less on appearance design and more on information. In the bottom plot of Fig. 1-9, we observe positive speech effects in multiple industries such as alcohol, automotive, education, government/nonprofit, real estate, and transportation.

Fig. 1-10 illustrates the effect of the appearance of a logo or person in video ads for each industry category. In the top plot, we observe logos contribute to better advertising outcomes in the luxury industry. This confirms the importance of logos and branding in luxury goods advertising. For the real estate industry, the logo is negatively associated with both video play metrics but positively associated with traffic. The bottom plot presents the effects of person in each industry. In most industry segments, the effects are either insignificant or positive. The one noticeable exception is the restaurant industry, which sees negative effects on conversion rates. In particular, the effect of person on the conversion rate in the restaurant industry is significantly worse than most of the other industry segments. This might be due to the nature of the restaurant industry, which people expect to focus on food.

Given the appearance of people, in Fig. 1-11 we observe the effects of the appearance of face and celebrity on different outcome metrics in each industry. The top plot shows either positive or insignificant effects of face in most industry segments. A few significant but mixed effects are also observed for the appearance of celebrity. Specifically, positive effects are generally observed in video play rates, e.g., in the luxury goods, telecommunications and transportation industry segments. We observe a

negative effect on click rate and conversion rate for education and technology. These are the two industries in which celebrity branding does not play an important role.

### 1.7.2 Media Platform

Fig. 1-12 illustrates the effects of basic video elements for each media platform: Facebook, LinkedIn, Pinterest, Snapchat, Twitter and YouTube. Each has four outcome metrics except for YouTube, which does not provide either 2s or 3s video play metrics. The effects of text are either insignificant or significantly negative on most platforms except for YouTube, where text has a positive association with conversions. The effects of speech on video play rates are commonly positive or insignificant. On LinkedIn, the speech effects on both play rates are much higher than the effects on click and conversion rates. The effects of logo on LinkedIn are significantly positive for all metric measurements. This might be because of the professional nature of LinkedIn. Logo effects on conversion rates on Snapchat are significantly negative. Finally, given the appearance of a person, the effect of possible celebrity is positive on the LinkedIn conversion rate and the Snapchat through-play rate.

## 1.8 Conclusion

There are some strategic implications based on our feature importance and fixed-effects model analysis. First, in terms of the basic creative elements, we find text to be the most important predictor of outcome metrics. The presence of text and its early appearance deteriorate the view-related outcome metrics, but not necessarily lower-funnel metrics. Advertisers do not have to be overly paranoid about the use of text. Unlike text, the presence of people and their early appearance contribute to better outcomes. It might be advisable to include people, faces and celebrities in ads to draw attention and persuade the audience at an earlier time. The effects of the presence of speech vary a lot by outcomes and hence it is difficult to draw convincing insights. For specific creative elements, we find having people gazing in the middle has the highest importance among all gaze directions, and it might

lead to better outcomes. The presence of females is more attractive than males. Emotion elements seem to have higher importance among specific creative elements. However, the effects do not produce consistent patterns and it is difficult to draw meaningful insights. Our heterogeneous analysis also generates some useful insights. In the luxury goods industry, the use of text is strongly discouraged, whereas text can play a positive role in video ads in the professional services industry. In ads for the transportation industry, the use of speech and faces, and particularly celebrities, are favored. In the education industry, the use of faces is favored but celebrities are not favored.

Admittedly, the major limitation of our research is the lack of robust causality. Despite some effort to address the endogeneity in video ad design by using fixed effects, the identification is still flawed. We also recognize that the effects of video features tend to be difficult to isolate, so perfect identification can be extremely hard to achieve. Our next step would be to run A/B experiments for better identification to validate the results we obtained in this research analysis. We are hoping this chapter will serve as an exploration that provides a comprehensive overview of the commonly defined and easily extracted video creative elements, including their importance and effectiveness.



# 1.9 Appendix

## 1.9.1 Figures

Figure 1-1: Example of Amazon Rekognition bounding box

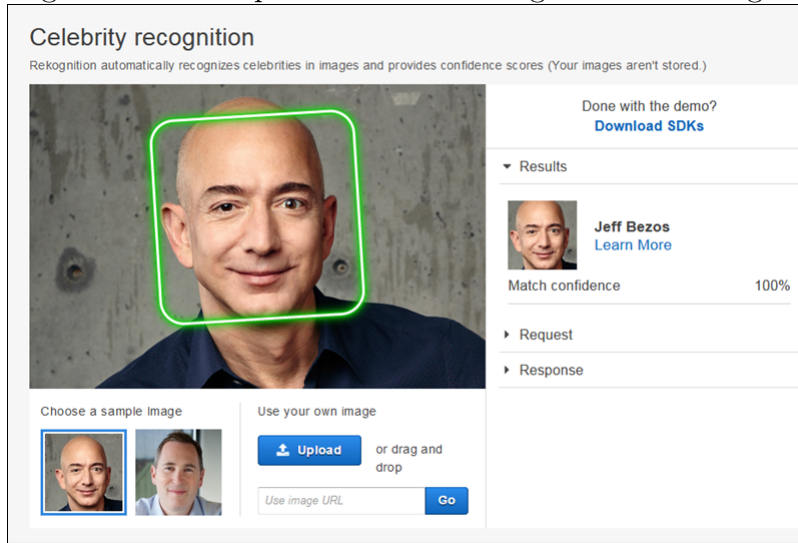
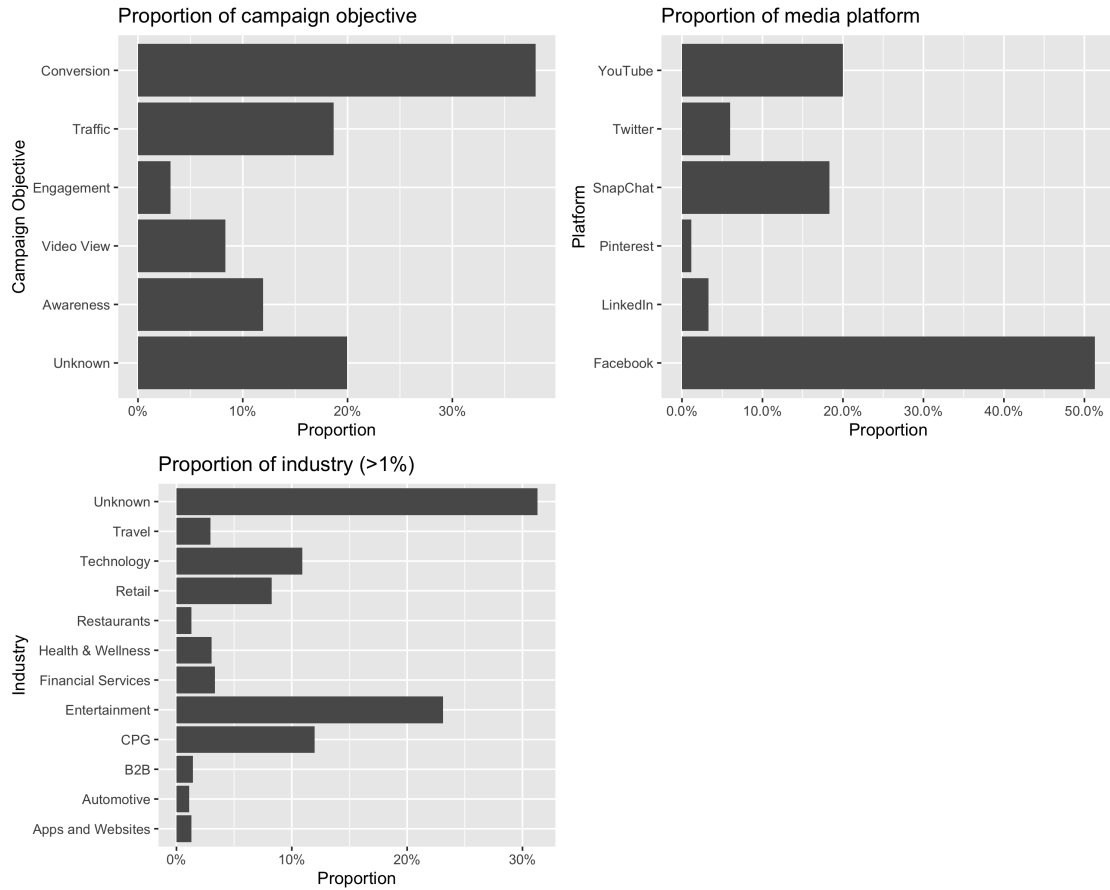
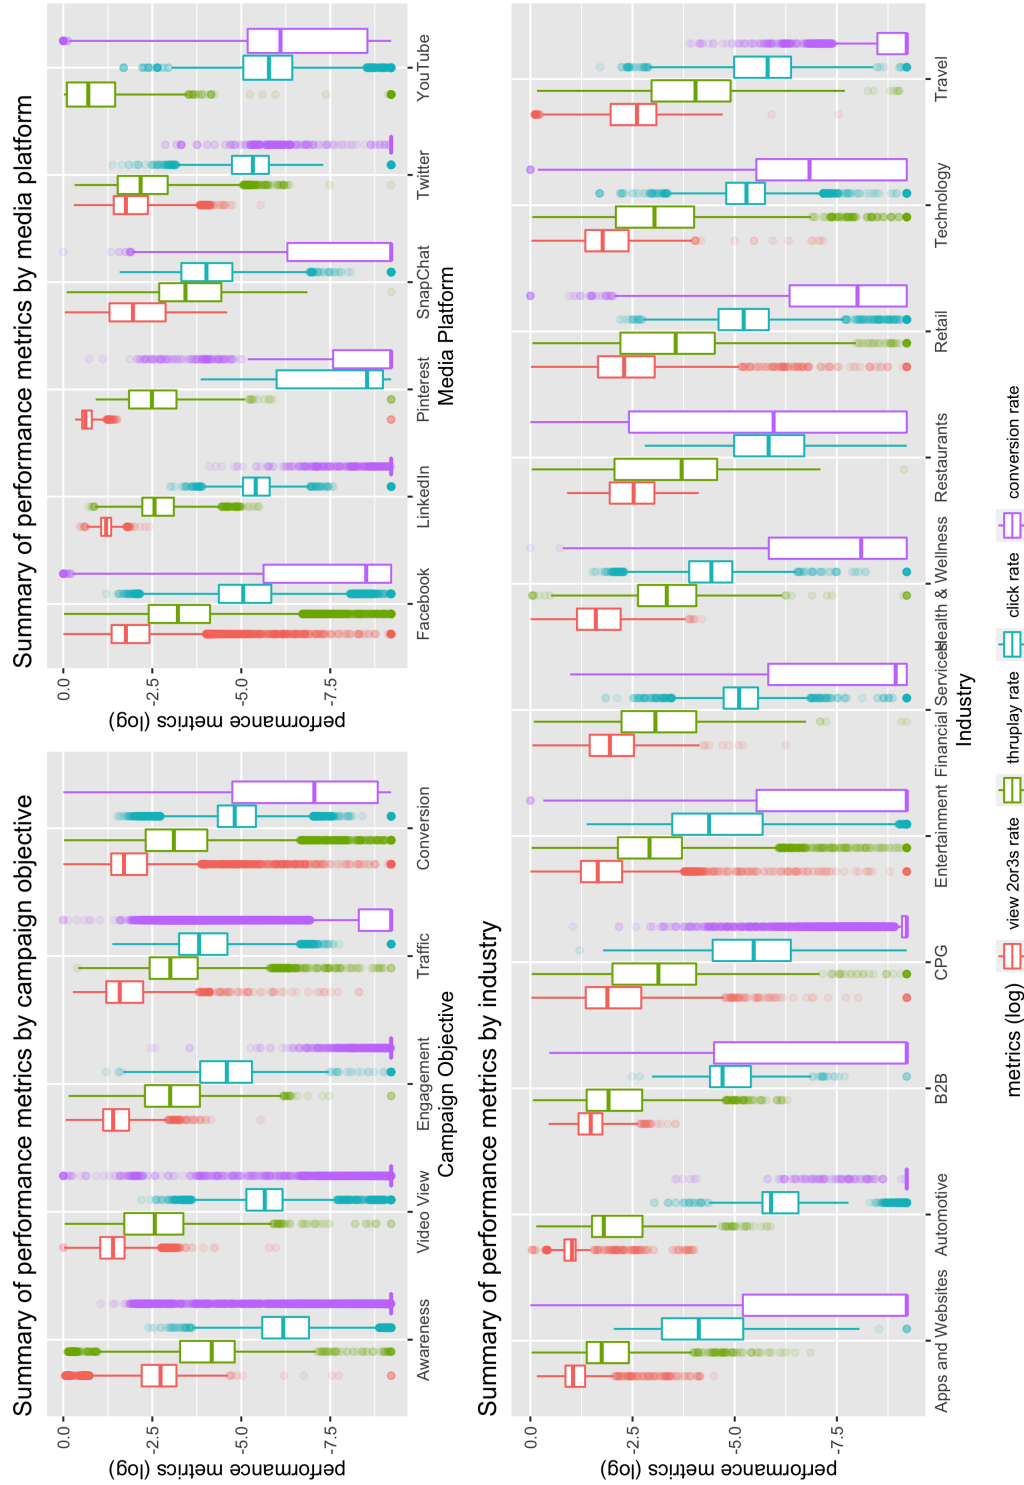


Figure 1-2: Barplots of proportion by campaign objective, platform and industry




Notes: The top left plot shows the proportion of each campaign objective in the data. The campaign objectives are synthesized and categorized into five major categories. The top right plot shows the proportion of each social media platform in our data. The bottom plot shows the proportion of industries that have proportion larger than 1%.

Figure 1-3: Boxplots of outcome metrics by campaign objective, platform and industry



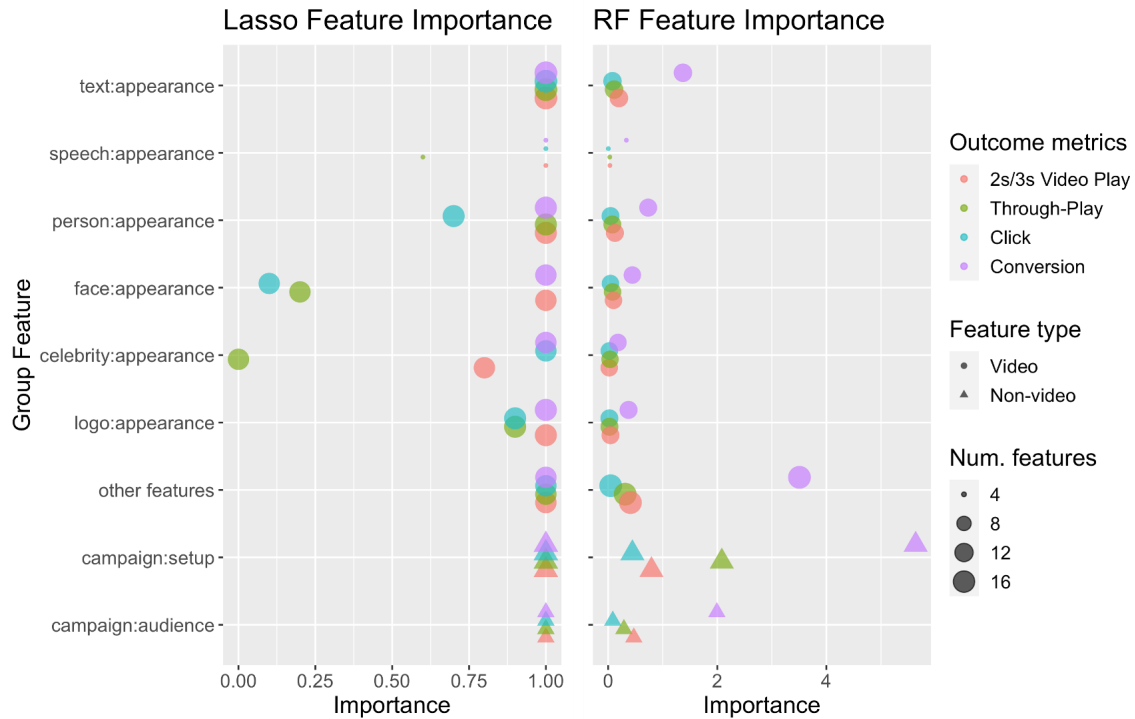
Notes: The top left boxplot summarizes the metrics of each campaign objective in the data. The campaign objectives are categorized into five major categories. The top right plot summarizes the metrics on each social media platform in our data. The bottom boxplot summarizes the metrics in industries that have proportion larger than 1%.

Figure 1-4: Unsupervised k-means clustering of creative elements' relative duration




Notes: Relative duration is obtained by the appearance duration of a given creative element in a given video ad divided by the total duration of the video ad. The x-axis shows the appearance elements we use in k-means. They include the basic elements (bottom plot), the age and gender of people appearing in the video (middle plot), and the gaze directions and whether people smile in the video (top plot). Each color represents a different cluster. Each point represents the mean of the relative duration for a creative element in the corresponding cluster. The size of the point represents the within-cluster standard deviation.

Figure 1-5: Basic video creative element group importance



Notes: The importance of each feature group on predicting each advertising outcome metrics (in log). The four colors represent four outcome metrics. The left plot shows the importance of video creative element groups using group lasso with stability selection. The right plot shows the importance of creative elements groups using random forest with group permutation importance. The size of points represents the number of features within each group. The shape of points indicates whether the group belong to campaign logistics (triangle) or video categories (circle).

Figure 1-6: Random forest group importance of face and person video elements on predicting advertising performance



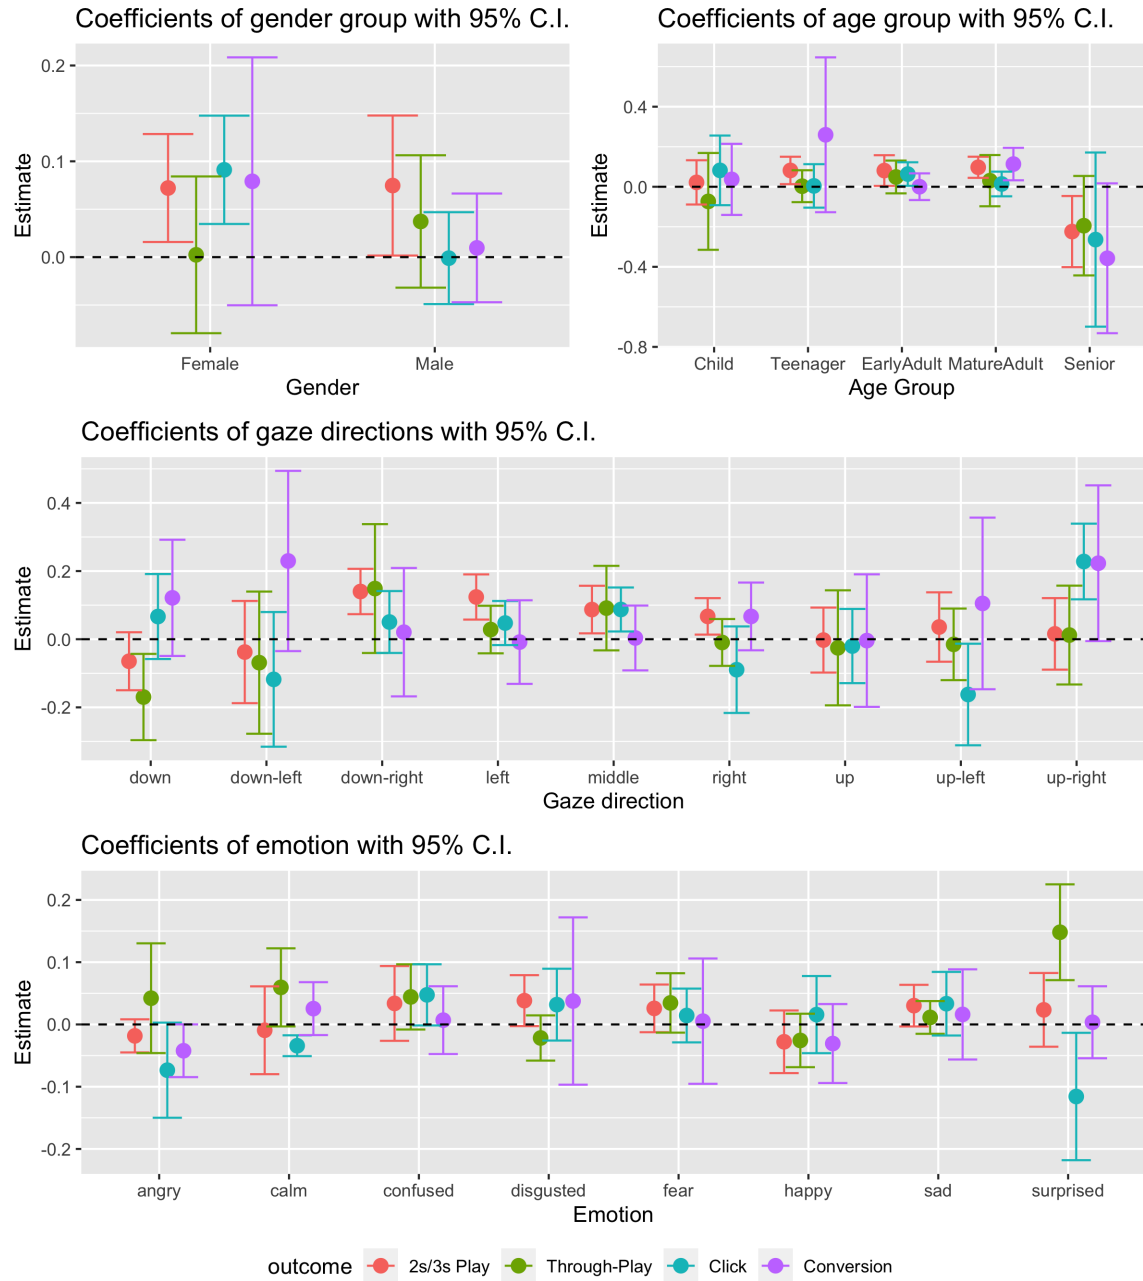
Notes: The figure presents the random forest group permutation importance scores of the specific video creative elements including gender, age, gaze directions and emotions of detected people and face in the video. The four plots correspond to four outcome metrics. The color represents the creative element group category. Each group include five features: relative start time, end time, duration, confidence mean and standard deviation.

Figure 1-7: Lasso group importance of face and person video elements on predicting advertising performance



Notes: The figure presents the randomized lasso group importance of the specific video creative elements including gender, age, gaze directions and emotions of detected people and face in the video. The four plots correspond to four outcome metrics. The color represents the creative element group category. Each group include five features: relative start time, end time, duration, confidence mean and standard deviation.

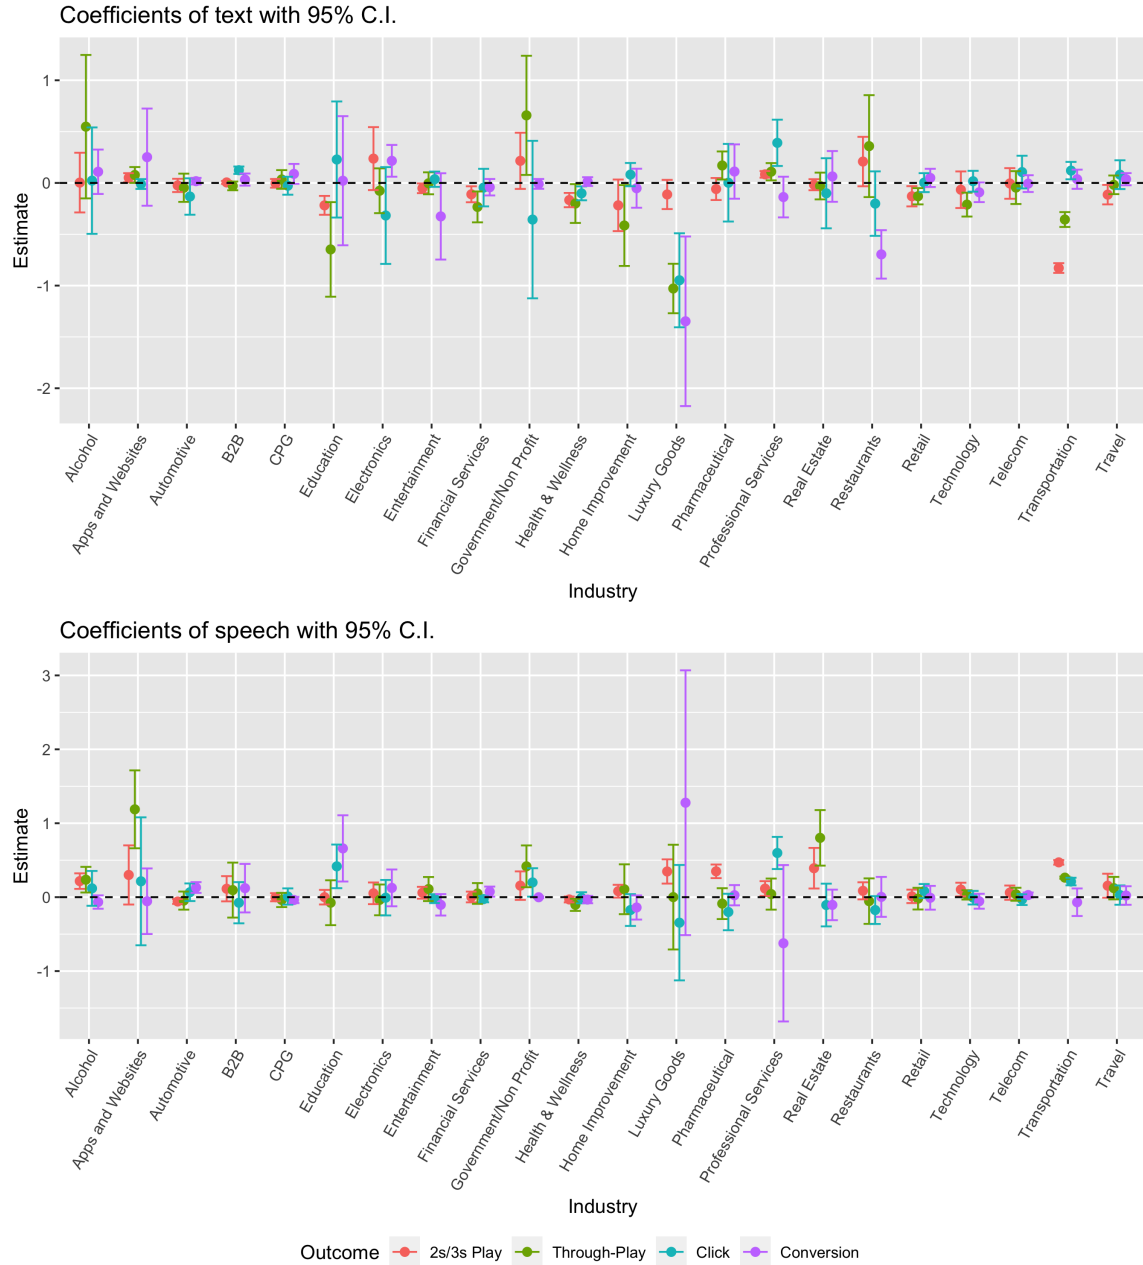
Figure 1-8: Fixed-effects model coefficients of specific video elements on advertising performance



Notes: Each point represents the fixed effects regression coefficient estimates of gender, age, gaze and emotion creative elements, with 95% confidence intervals. The fixed effects include week, advertiser, industry, platform and campaign objective. The outcome variables are log 2s/3s video play rate, log through-play rate, log click rate and log conversion rate in four different colors. The control variables, apart from the variables in the plots, include the basic creative elements' appearance indicators, campaign setup and campaign audience demographics.


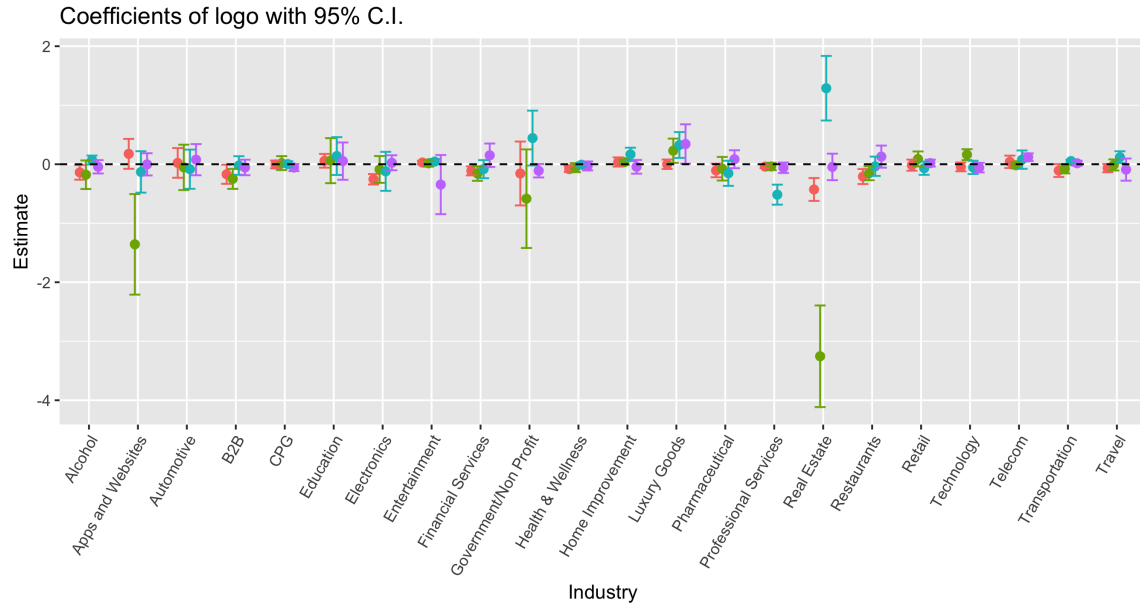


Figure 1-9: Fixed-effects model coefficients of basic video features by industry - part 1



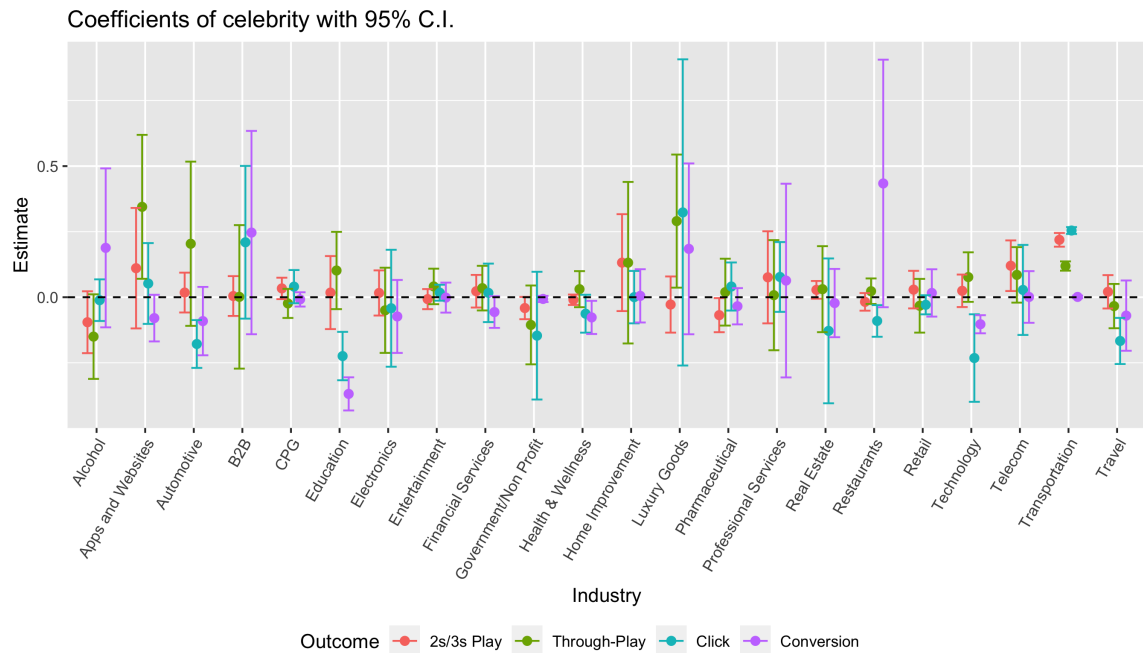

Notes: Each point represents the fixed effects regression coefficient estimates of text and speech indicators, with 95% confidence intervals, for each industry segment. The fixed effects include week, advertiser, platform and campaign objective. The outcome variables are log 2s/3s video play rate, log through-play rate, log click rate and log conversion rate in four different colors. The independent variables include the basic creative elements' appearance indicators, campaign setup and campaign audience demographics.

Figure 1-10: Fixed-effects model coefficients of basic video features by industry - part 2



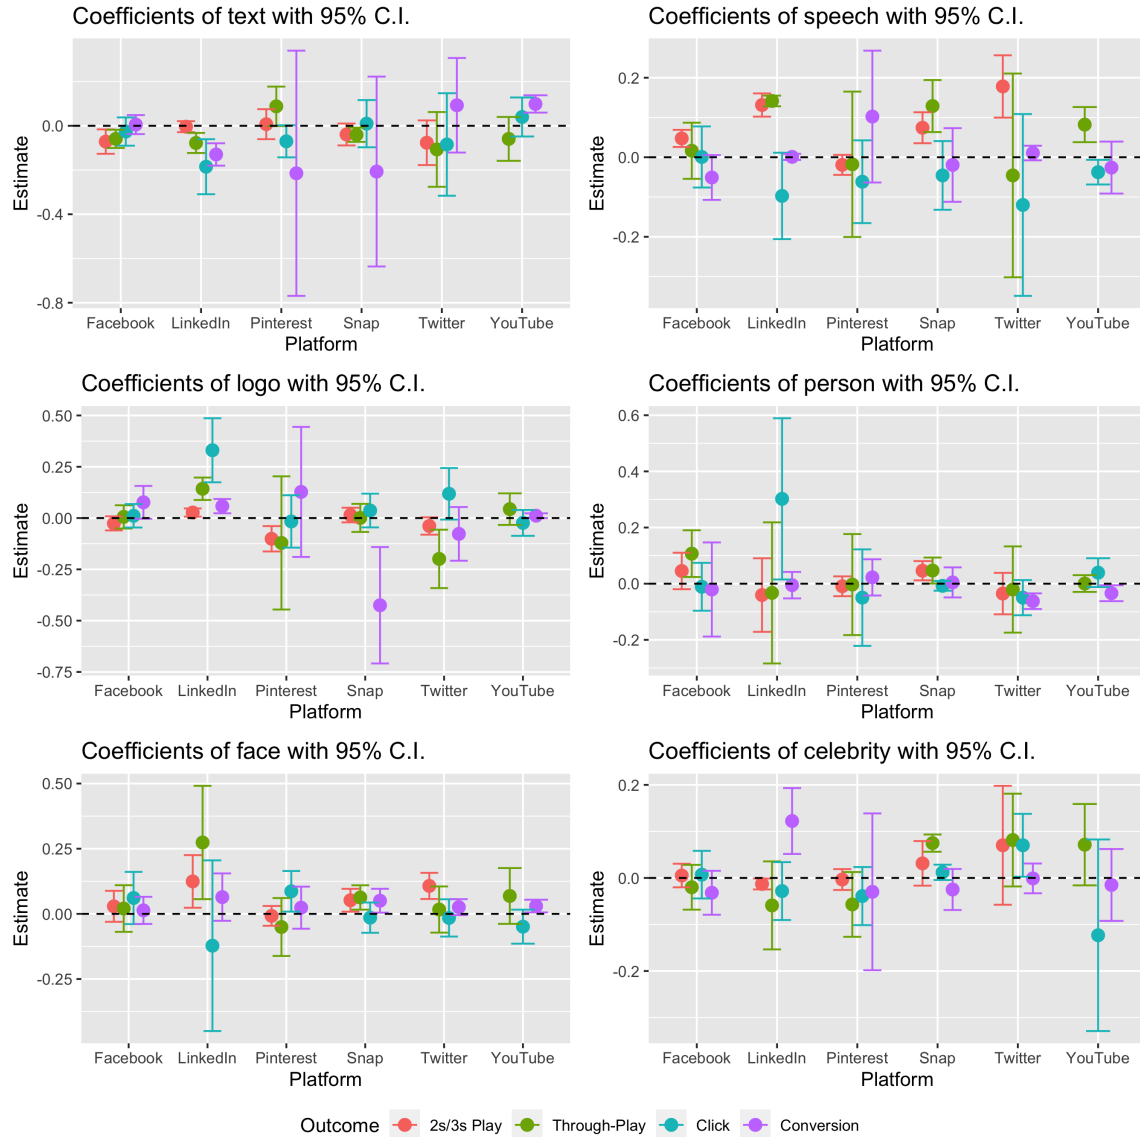
Notes: Each point represents the fixed effects regression coefficient estimates of logo and person indicators, with 95% confidence intervals, for each industry segment.

Figure 1-11: Fixed-effects model coefficients of basic video features by industry - part 3




Notes: Each point represents the fixed effects regression coefficient estimates of face and celebrity indicators, with 95% confidence intervals, for each industry segment.

Figure 1-12: Fixed-effects model coefficients of basic video features by media platform




Notes: Each point represents the fixed effects regression coefficient estimates of basic creative elements' appearance indicators, with 95% confidence intervals, for each media platform. The fixed effects include week, advertiser, industry and campaign objective. The outcome variables are log 2s/3s video play rate, log through-play rate, log click rate and log conversion rate in four different colors. The independent variables include the basic creative elements' appearance indicators, campaign setup and campaign audience demographics.

Figure 1-13: Elbow method to determine the optimal number of clusters




Notes: the curve knee, i.e., the optimal number of clusters, is determined to be 5.

Figure 1-14: Robustness check on k-means clustering of creative elements with 4 clusters




Notes: Number of clusters = Number of optimal clusters - 1.

Figure 1-15: Robustness check on k-means clustering of creative elements with 6 clusters



Notes: Number of clusters = Number of optimal clusters + 1.

Figure 1-16: Random forest sub-group breakdown of basic video elements



Notes: The figure presents the group permutation importance of the subdivision of basic video creative elements. The four plots correspond to four outcome metrics. The color categorizes the basic creative element group. The size of points represent the number of features within each sub-group.

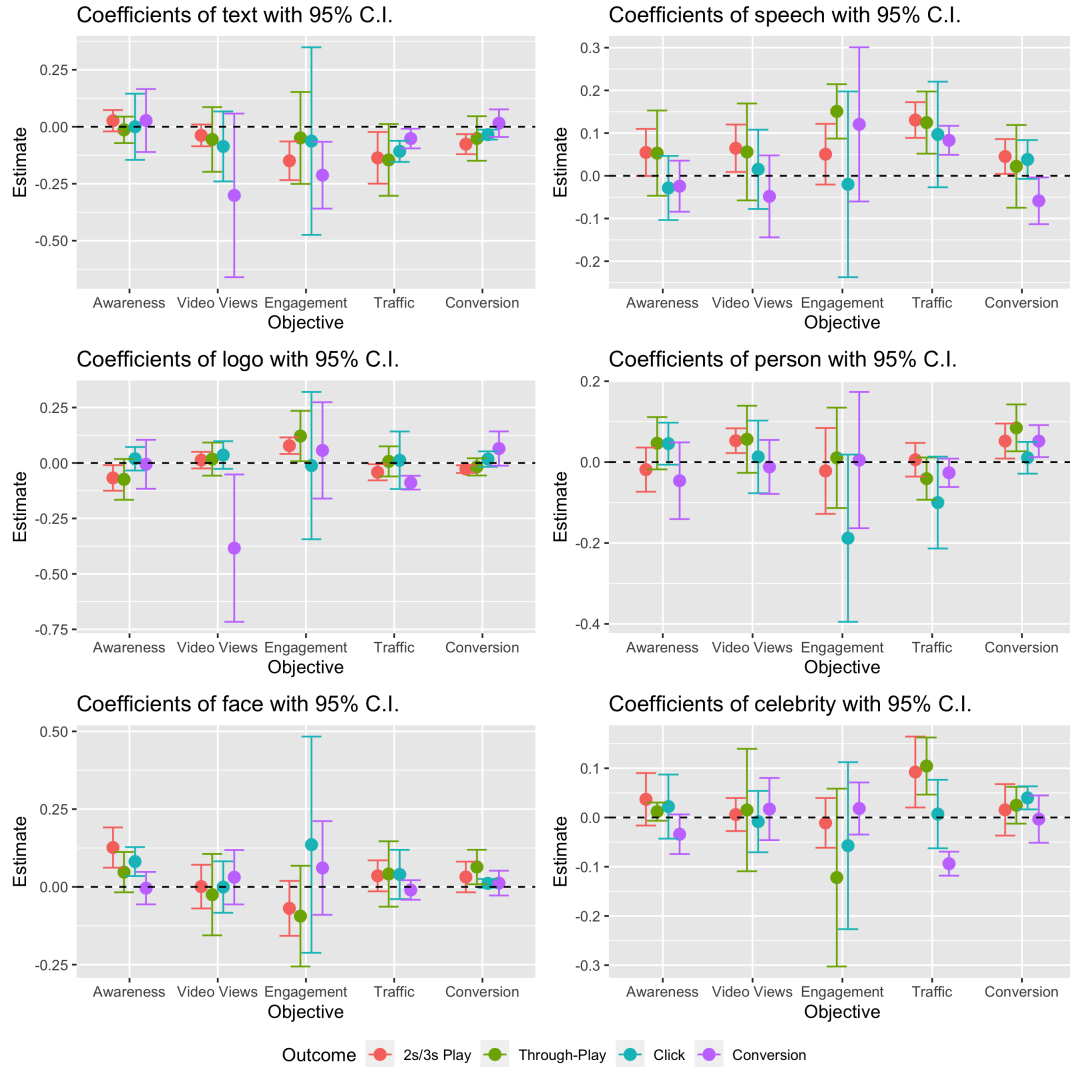


Figure 1-17: Lasso sub-group breakdown of basic video elements



Notes: The figure presents the randomized lasso group importance of the subdivision of basic video creative elements. The four plots correspond to four outcome metrics. The color categorizes the basic creative element group. The size of points represent the number of features within each sub-group.

Figure 1-18: Coefficients of basic video features by campaign objective



Notes: Each point represents the fixed effects regression coefficient estimates of basic creative elements' appearance indicators, with 95% confidence intervals, for each campaign objective. The fixed effects include week, advertiser, industry and platform. The outcome variables are log 2s/3s video play rate, log through-play rate, log click rate and log conversion rate in four different colors. The independent variables include the basic creative elements' appearance indicators, campaign setup and campaign audience demographics.

## 1.9.2 Tables

Table 1.1: Table of basic video creative elements summary statistics

Creative Elements	Prop.	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
text	0.90	0.00	0.65	0.89	0.74	0.97	1.00
face	0.67	0.00	0.00	0.30	0.37	0.70	1.00
celebrity	0.37	0.00	0.00	0.00	0.18	0.29	1.00
person	0.78	0.00	0.00	0.56	0.49	0.86	1.00
speech	0.35	0.00	0.00	0.00	0.20	0.33	1.00
logo	0.30	0.00	0.00	0.00	0.09	0.03	1.00

*Note:* The first proportion column refers to the proportion of video that has the specific creative element detected, calculated by the number of videos that has appearance of the specific creative element divided by the total number of videos in the data. The statistics in the rest of the columns represents the relative duration – duration of the element appearance in a video divided by the total duration of the video.

Table 1.2: Table of basic campaign feature and outcome metrics summary statistics

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
media duration (sec)	3.00	6.14	10.20	19.36	15.34	120.00
campaign duration (day)	1.00	5.00	16.00	43.71	44.00	365.00
impressions	50	9.5k	93.0k	12.4m	962.7k	7.3b
2s/3s play rate	0.00	0.08	0.19	0.23	0.30	1.00
thru-play rate	0.00	0.02	0.06	0.17	0.16	0.99
click rate	0.00	0.00	0.01	0.01	0.01	0.30
conversion rate	0.00	0.00	0.00	0.01	0.00	1.00

*Note:* The top three rows summarizes the important campaign predictors. The bottom four rows summarizes the four main outcome metrics that we use in our model.

Table 1.3: Fixed-Effects Model with Basic Video Elements

	<i>Dependent variable: (log)</i>			
	2s/3s Play Rate	Thru-Play Rate	Click Rate	Conversion Rate
	(1)	(2)	(3)	(4)
text	-0.058** (0.018)	-0.087 (0.051)	0.027 (0.020)	0.032 (0.050)
person	0.045*** (0.012)	0.057 (0.031)	0.037 (0.020)	-0.035 (0.027)
person:face	0.037 (0.019)	0.047 (0.027)	-0.036 (0.025)	0.027 (0.014)
person:celebrity	0.018 (0.013)	0.030 (0.023)	-0.053 (0.043)	-0.025 (0.024)
speech	0.063*** (0.017)	0.049 (0.029)	-0.017 (0.022)	-0.051 (0.036)
logo	-0.016 (0.016)	0.039 (0.030)	0.009 (0.018)	-0.070 (0.073)
log(video duration)	0.043 (0.032)	-0.783*** (0.141)	0.190** (0.067)	0.004 (0.032)
log(impressions)	-0.238*** (0.035)	-0.159 (0.091)	0.027 (0.065)	-0.297* (0.117)
log(spend)	0.251*** (0.035)	0.223** (0.069)	0.021 (0.076)	0.396** (0.121)
Audience controls	Yes	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Yes
Observations	317,520	650,353	650,353	650,353
R <sup>2</sup>	0.676	0.790	0.605	0.806
Adjusted R <sup>2</sup>	0.675	0.789	0.604	0.806

*Note:* Robust standard errors clustered at industry, platform and campaign objective levels. The audience controls include the percentage of audience of each gender and age categories for Facebook ads and gender for Twitter ads. We do not have audience demographics data for other platforms. The fixed effects include week, advertiser, industry, platform and campaign objective. The outcome variables are log 2s/3s video play rate in column (1), log through-play rate in column (2), log click rate in column (3) and log conversion rate in column (4). The control variables, apart from the variables in the table, include campaign audience gender and age groups. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 1.4: Fixed-Effects Model with Basic Video Elements - Relative Appearance Duration

	<i>Dependent variable: (log)</i>			
	2s/3s Play Rate	Thru-Play Rate	Click Rate	Conv. Rate
	(1)	(2)	(3)	(4)
text duration	-0.074 (0.037)	-0.105* (0.048)	-0.004 (0.046)	-0.054 (0.027)
person duration	0.035 (0.022)	0.005 (0.047)	-0.065 (0.041)	0.017 (0.076)
face duration	0.056 (0.033)	0.010 (0.029)	0.076* (0.036)	0.030 (0.035)
celebrity duration	0.069 (0.043)	0.128* (0.049)	0.092 (0.052)	-0.027 (0.084)
speech duration	0.058 (0.042)	0.101*** (0.026)	-0.066 (0.049)	0.103* (0.039)
logo duration	-0.033 (0.039)	-0.031 (0.052)	-0.041 (0.049)	-0.110 (0.130)
Video Feature Ctrls	Yes	Yes	Yes	Yes
Campaign Controls	Yes	Yes	Yes	Yes
Audience Controls	Yes	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes	Yes
Observations	317,520	650,353	650,353	650,353
R <sup>2</sup>	0.677	0.790	0.605	0.806
Adjusted R <sup>2</sup>	0.676	0.789	0.604	0.806

*Note:* Robust standard errors clustered at industry, platform and campaign objective level. The basic video feature controls refer to basic creative elements' presence indicators, the media duration and log number of labels. The campaign setup controls refer to the campaign log expenditure and impressions. The audience controls refer to audience demographic distributions. The fixed effects include week, advertiser, industry, platform and campaign objective. The outcome variables are log 2s/3s video play rate in column (1), log through-play rate in column (2), log click rate in column (3) and log conversion rate in column (4). The control variables, apart from the variables presented in the table, include the basic creative elements' appearance indicators, campaign setup and campaign audience demographics. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 1.5: Fixed-Effects Model with Basic Video Elements - Start Time

	<i>Dependent variable:</i>			
	2s/3s Play Rate	Thru-Play Rate	Click Rate	Conv. Rate
	(1)	(2)	(3)	(4)
text start time	0.272*** (0.068)	0.028 (0.089)	0.010 (0.067)	0.119 (0.072)
person start time	-0.106** (0.032)	-0.079* (0.036)	0.072 (0.042)	-0.050 (0.046)
face start time	-0.051* (0.024)	0.003 (0.026)	0.015 (0.029)	0.007 (0.025)
celebrity start time	-0.058 (0.031)	-0.061 (0.063)	-0.125* (0.048)	-0.014 (0.071)
speech start time	-0.008 (0.044)	-0.080* (0.029)	0.169* (0.068)	-0.090 (0.062)
logo start time	0.009 (0.033)	0.036 (0.051)	0.014 (0.033)	0.168 (0.087)
Video Feature Ctrls	Yes	Yes	Yes	Yes
Campaign Controls	Yes	Yes	Yes	Yes
Audience Controls	Yes	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes	Yes
Observations	317,520	650,353	650,353	650,353
R <sup>2</sup>	0.677	0.790	0.605	0.806
Adjusted R <sup>2</sup>	0.676	0.789	0.604	0.806

*Note:* Robust standard errors clustered at industry, platform and campaign objective level. The control variables, apart from the variables presented in the table, include the basic creative elements' appearance indicators, campaign setup and campaign audience demographics. The fixed effects include week, advertiser, industry, platform and campaign objective. The outcome variables are log 2s/3s video play rate in column (1), log through-play rate in column (2), log click rate in column (3) and log conversion rate in column (4). \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 1.6: Fixed-Effects Model with Basic Video Elements - Location

	<i>Dependent variable: (log)</i>			
	2s/3s Play Rate	Thru-Play Rate	Click Rate	Conv. Rate
	(1)	(2)	(3)	(4)
text:center	-0.053 (0.044)	-0.078 (0.042)	-0.090* (0.041)	0.016 (0.036)
text:top	-0.049 (0.037)	-0.060 (0.048)	0.111 (0.062)	-0.041 (0.108)
text:bottom	-0.079* (0.030)	-0.048* (0.023)	0.146 (0.074)	-0.020 (0.034)
text:left	-0.071* (0.030)	0.006 (0.035)	0.082 (0.060)	-0.094 (0.047)
text:right	0.005 (0.047)	-0.094* (0.046)	-0.004 (0.048)	0.115 (0.065)
person:center	-0.094* (0.043)	-0.033 (0.059)	-0.058 (0.046)	-0.040 (0.036)
person:top	0.050* (0.023)	0.059 (0.041)	0.069 (0.034)	0.043 (0.066)
person:bottom	0.126** (0.042)	0.014 (0.075)	0.055 (0.031)	-0.018 (0.046)
person:left	-0.095** (0.034)	-0.021 (0.029)	-0.018 (0.019)	-0.016 (0.048)
person:right	-0.057 (0.032)	-0.022 (0.026)	0.021 (0.032)	-0.059 (0.038)
logo:top	-0.001 (0.026)	-0.006 (0.022)	0.025 (0.030)	-0.085** (0.028)
All Controls	Yes	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes	Yes
Observations	317,520	650,353	650,353	650,353
R <sup>2</sup>	0.678	0.790	0.606	0.806
Adjusted R <sup>2</sup>	0.676	0.789	0.605	0.806

*Note:* Robust standard errors clustered at industry, platform and campaign objective level. The control variables, apart from the variables presented in the table, include the basic creative elements' appearance indicators, campaign setup and campaign audience demographics. The fixed effects include week, advertiser, industry, platform and campaign objective. The outcome variables are log 2s/3s video play rate in column (1), log through-play rate in column (2), log click rate in column (3) and log conversion rate in column (4). \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .



# Chapter 2

## Do Algorithms Help Firms Achieve Their Targeting Objectives?

### 2.1 Introduction

In the pre-digital era, it used to be the case that advertising would try to elicit the desired response by optimizing ad content. For example, an ad designed to elicit persistent interest would be designed differently from an ad designed to try to get consumers to take immediate action. However, in the digital age, advertisers can now use digital algorithms and specify to the advertising provider that their ads should elicit a particular action. This chapter studies how this transformation affects how advertisers should approach ad design.

Compared to television advertising, digital advertising, and especially skippable video ads, allows viewers to easily and quickly skip ads and avoid the delay of seeing the organic content they desire [20]. Digital advertising also makes it easier to track the audience's reactions and the extent to which they skip ads or only watch the beginning, download apps, or click. Behavioral data has helped digital advertising platforms train machine learning algorithms to target an audience more precisely to obtain the desired action. Such data were previously unavailable for mass-media advertising.

Due to the ability to use digital data and machine learning to achieve concrete

goals, digital ad platforms now offer very specific campaign objectives that advertisers must choose between. For example, on platforms such as Facebook and Amazon, advertisers must specify only one campaign objective for each campaign, from upper-funnel goals such as reaching a larger potential audience to lower-funnel goals such as increasing traffic or even conversion. Evidently, the ability to specify that an advertiser hopes for awareness or conversions and for an audience that is geared towards awareness or conversions is a huge benefit to advertising brought on by the digital age and machine learning. However, some advertisers may worry that the specified actions will come at the expense of genuine engagement with their ad, for example. Such a concern is especially valid given the multitask literature suggesting that it is often impractical to achieve multiple goals, if an agent, in our case an algorithm, is rewarded for achieving a single goal [29].

We investigate this using field experiment data from Facebook, where multiple non-profits ran ads using the same video ad content. Our field tests deployed a  $2 \times 2$  design. The two dimensions of variations are higher-funnel (2s continuous video view) vs. lower-funnel (link click) optimization objectives and whether the product is featured before vs. after 3 seconds. We find evidence that ad algorithms are very effective at achieving specified objectives. For example, an ad optimized for video views will receive many more views than an ad optimized for clicks, even though the ad content is identical. Similarly, if an ad objective is specified as clicks, we find it receives significantly more clicks than the ad under the view objective even if the ad content is the same. We present evidence that the effectiveness of algorithms at delivering their objective comes from their ability to identify an audience that is more likely to click or more likely to view. Therefore, the algorithm automatically distorts the audience that the ad receives to achieve the goal.

We then investigate how advertisers should respond to this algorithmic audience distortion if they want to achieve both clicks and views by varying the second treatment in our field experiment, i.e., when the key information is revealed. In a digital world with increasingly fragmented attention, the point at which to emphasize a product-based marketing message within a video seems to be a factor that could po-

tentially moderate audience distortion due to algorithmic optimization. On the one hand, placing product information at the beginning of a skippable advertising video could help ensure more people get to know the product before they skip the ad. On the other hand, revealing the product later might be better for attracting people's interest and engagement. Our results show that the 3s view probability of videos with late informational content tends to be higher than early information when the campaign is specified to optimize clicks.

Our proposed mechanism for the observed effects concerns the different levels of distortion in audience targeting as brought by optimization due to different campaign objectives. Unfortunately, the right audience based on the platforms' algorithms does not always align well with the right audience for advertising firms. We believe there are two types of online audiences. We refer to one type as information seekers and the other as content consumers. We define information seekers as people who genuinely want to obtain more information about the product and are also much more likely to be converted to true potential customers. In contrast, content consumers mostly watch video ads for entertainment purposes and have little interest in the advertised product or service. We believe that a higher-funnel campaign objective creates more audience distortion as the ad is delivered to a high proportion of content consumers, and the views can be optimized when delivered to more content consumers. This is confirmed by the contrast between the high 3s view rates and the extremely low click rates of the view campaigns in the first set of results. Moreover, information seekers are more patient in waiting for key ad messages. That explains the 3s video play probability for late informational content being higher than early information under lower-funnel click objectives, as illustrated in the second set of results.

The experiment findings raise a managerial question of whether achieving one high-performance metric necessarily signals the advertisement's success. Digital platforms often do an excellent job of achieving the one optimization goal specified in the advertising campaign setup, thanks to advanced algorithms. However, some algorithmic optimizations may boost a single specific metric at the expense of missing true potential customers. Moreover, it calls into question the current practice that

the platforms make advertisers specify just one optimization objective while setting up any advertising campaign. The main contribution of this chapter lies in identifying the audience distortion in digital marketing, especially when specified with a high-funnel campaign objective. We point out that deviation from the actual intended audience is due to the platform’s algorithmic optimization. We also provide empirical evidence on whether and how the timing of informational content in video ads exerts an impact on mitigating audience distortion. Overall, we hope that our research alerts advertisers to bias brought by algorithmic optimization practices and the cost of overly relying on a platform’s algorithmic targeting ability. We encourage firms to be more strategic when publishing ad campaigns and designing video ads in consideration of reaching true potential customers.

## 2.2 Literature Review

This chapter contributes to three streams of academic literature.

The first stream is the literature discussing the bias in algorithmic optimization used by advertising platforms in the auction process. The main obstacle with Facebook experimentation’s random assignment is its auto-bidding algorithms [1, 21]. The platform’s sophisticated optimization algorithms during ad auctions deliver the ad to people who are most likely to generate profitable actions for the platform. Although some attempts have been made to resolve such ad delivery bias – e.g., ghost ads that run simulation auctions to identify the would-be-exposed consumers in the control group [35] – these proposed methods are not currently adopted in industry practices. Our results build on the findings of algorithmic bias through a comparison of audience differences under varying campaign objectives.

The second stream focuses on algorithmic fairness. Research shows that social fairness might sometimes be compromised due to algorithmic bias [2, 17, 45, 39]. One example is that the optimization algorithm is found to discriminate against women when delivering gender-neutral, job-promoting ads because women are more expensive to target [39]. Our research contributes to this topic by suggesting there might also

be business costs to all firms that publish ads on these digital platforms, in addition to the social cost that previous research illustrates.

The final stream is the timing of informational content in video ads. The informational content that we focus on in this chapter is the key product benefits. Such informational content generally discloses some main product attributes, including pricing and promotion, which could serve as a competitive advantage [5, 4, 44]. Theory suggests that a high-quality firm might choose to produce video ads with no or minimal informational content in order to invite the consumer to engage in search [44]. In the case of skippable video ads, multiple research projects also explore people’s skipping behavior and how it relates to content informativeness [63, 14, 20, 15]. This paper on skippable ads [20] is closely relevant to our research. It explores and finds that more informational ad content improves conversion whenever the consumer has a negative prior about the advertiser’s product. In contrast, less information is preferred when consumers have a positive prior. Overall, there is not much empirical work assessing the effects of the timing of informational content in skippable ads causally. Our research wants to fill this gap by providing field experiment evidence.

## 2.3 Methodology

We explored our question of interest by running field experiments on a digital advertising platform – Facebook. We chose Facebook because it has sophisticated targeting capability and a mature advertising A/B testing experimentation platform, and it provides relatively more detailed performance metrics. Our main experiment aimed to assess the effects of varying campaign optimization objectives and the timing of informational content. We recruited 17 organizations and created videos that fit our experimental design based on their video and image materials.

This section first describes the treatment conditions we have in the experiment, then we proceed to data sampling and discussion of the identification. Next, we give an overview of our data by providing some summary statistics, including the key outcome metrics. Lastly, we describe the model we use to analyze the collected data.

### 2.3.1 Treatment

We adopted a  $2 \times 2$  experimental design for each organization's Facebook ad campaign. Facebook provides a well-established A/B testing platform to experiment with the campaign objective and different creative designs, as shown in Figure 2-1.

The first treatment variation was campaign optimization objectives. Whenever advertisers want to set up an online campaign, the platform will ask them to select a campaign objective that they would like to optimize for. Campaign objectives define what advertisers want to accomplish with a specific Facebook ad campaign and then help Facebook optimize their ads according to the selected objective. The campaign objectives determine the audience to whom the ad is delivered and how the platform charges. There are two different campaign optimization objectives that we consider, based on where the objective lies in the marketing funnel. From top to bottom, marketing funnels cover the process of a consumer from getting to know the product to eventually making the purchase decision. The higher-funnel optimization objective is video views, which optimizes a 2-second view-through. A 2-second view-through means at least two seconds of video play if the video is 2 seconds or longer. Fulfillment of this objective signals consumers' possible awareness and interest in the product/services. The lower-funnel objective is to optimize link clicks. With this objective, Facebook tries to deliver ads to people who are most likely to click on the link and load the website. Link click optimization is a relatively lower-funnel campaign objective compared to video view optimization as it involves more active consideration about purchasing the advertised product/service.

The other treatment variation lay in the timing of informational content in a video ad. Informational content is defined as the text messaging that reveals the key product benefit or the organization's mission in our context. We varied the timing of the first appearance of any informational content for each organization by creating two versions of the video ads. The early informational content group referred to video ads with informational content appearing within the first 2 seconds; the late informational content group referred to video ads with informational content

appearing after 3s of the video. For each organization, we tried to make the creative design in the two versions as similar as possible, with the only difference being the timing of the informational content.

In short-duration video ads, informational content is typically presented in text messages and speech. Given that around 85% of users watch videos on Facebook with the sound off [49], we used text messages to disclose the informational ad content in our experiments.

The experiments spanned March 2020 to November 2020. For each organization, the campaign lasted three days. None of the organizations had overlapping campaign periods. Each treatment condition for each organization had the same budget allocation, i.e., \$5 per treatment per day. For all the video ads in our experiment, we included a "Learn More" button at the bottom of the ad, as shown in Figure 2-2. We defined the target audience as 18+ in age and located within the US for all organizations.

### **2.3.2 Data Sampling and Identification**

We collected data by recruiting mostly nonprofit organizations, creating video ads that fit our experiment design, and running Facebook ad campaigns as A/B experiments. In total, we recruited 17 organizations and hence performed 17 rounds of A/B experiments on Facebook. Among them, 16 were small- or medium-sized nonprofit organizations, and one was a small local restaurant. These organizations had diverse geographical distribution, not only confined domestically but also reaching out to developing countries like India, Uganda, Cameroon, Kenya, etc. The service targets varied from endangered animals to disadvantaged people such as minorities, refugees, the poor, the disabled, the elderly, and children with special needs. The service missions covered many areas such as environmental protection, job opportunities, healthcare, education, and affordable food. When creating and publishing Facebook ad campaigns, we required the video ads to be skippable at any time of the video and between 10s and 30s in duration. Ideally, to gain more statistical power and generalized insights, we might want to run as many campaigns on different ad-

vertising products or services as possible. Given budget constraints, we managed to run 17 sets of experiments.

The main obstacle with such random assignment, particularly for creative design testing, is Facebook’s auto-bidding algorithms [1]. The fact that people who are shown the ad are those who are most likely to generate profitable actions for the platform might also apply in the case of changing the ad’s creative design. For example, suppose we set up A/B testing to test the two creatives with the same campaign objective, one with early informational content and the other with late informational content. Although users are randomly separated into two groups, not everyone in each group will see the video ad. The first one will be shown to some of group A’s users, specifically picking those who are most likely to prefer early informational content, and the second will be shown to some of group B’s users, picking those who are most likely to prefer late informational content. Therefore, the main limitation of our research is that we could not determine whether the effects we observed for informational content timing mainly came from the counterfactual difference or the difference in audience. To address this issue, we will be running a follow-up study that changes another more detectable aspect of the creative ad design – age appeal. We find no substantial evidence suggesting significant audience distortion due to creative design change. We will discuss this follow-up experiment in more detail in later sections.

### **2.3.3 Data Summary**

We consider several different metrics for our main experiment. All the metrics are aggregated by counts due to Facebook’s de-identification practices. The four primary metrics that we will be using to calculate outcome measures are impressions, 3s video plays, through-plays, and link clicks. Impressions refer to the number of times a video is shown to an audience in the campaign. We use it in the calculation of our outcome measure. Overall, for the main analysis, there are a total of 377,848 impressions as shown in Table 2.1. Among them, 74.0% fall under view objective campaigns and 26.0% fall under click objective campaigns. Recall that both view and click objective



campaigns are allocated the same budget. With the same amount of cost, view objective campaigns deliver to many more people. Based on the summary statistics, advertisers might have a false impression that view optimization is a better campaign objective as it is more cost-efficient in delivering the ad to an audience. However, our analysis shows that does not necessarily mean that the view objective is a better choice.

In addition to impressions, we also collected view time-related metrics such as 3-second video plays and through-plays. A three-second play refers to watching at least the first 3s of the video ad before skipping it. A through-play indicates either watching through the entire video, if the video ad is less than 15s in duration, or watching for at least 15 seconds, if the video ad is more than 15s in duration. In our later analysis and results, we might interchangeably refer to it as either through-play or 15-second video play. Finally, we collected a link click metric, which counts link clicks on any ad area that links to destinations or experiences for the ad. In our case, it refers to the "Learn More" button as shown in the example in Figure 2-2. We plan to use this metric to explore people's true interest in different versions of ads in the main study. The three metrics reflect people's different levels of engagement and interest in the video ads. The total counts of the three metrics are presented in Table 2.1's "Count" column. The rate metrics, as shown in the last column, are obtained by dividing the count by the total number of impressions. For example, the 3-sec video play rate is the number of times a video ad has been viewed for more than three seconds divided by the total impressions. It also serves as a proxy of the probability of a user viewing a video ad for at least 3s. We observe that the rates decrease exponentially along the marketing funnel from the 3s video play rate to the through-play rate and from the through-play rate to the link click rate.

### **2.3.4 Data Analysis**

Facebook only provides advertisers with aggregate data that counts the number of actions and impressions that ads receive. Advertisers on Facebook do not have access to individual-level data in order to protect Facebook users' privacy. To address this,

in our empirical specification, we used a binomial regression model with a logit link as our primary empirical specification. The finest aggregation level available is by campaign  $g$ , day  $t$ , and demographic group  $d$  – age and gender. For each aggregate unit, we obtained the count of impressions  $r$  and the count of positive advertising actions  $s$ .

In this regression analysis, the actions we study include a 3-sec video play, a (15s) through-play, and a link click. We use  $g$  to denote each organization we create videos based on and use  $i$  to indicate each aggregate unit.  $E_{ig}$  is an indicator of whether an aggregate unit of users  $i$  under organization campaign  $g$  is shown a video ad with early informational content within the first 2s.  $V_{ig}$  is an indicator of whether a unit is shown a video with the campaign objective that optimizes for clicks ( $V_{ig} = 0$ ) or video views ( $V_{ig} = 1$ ). In this specification, we used video campaigns with late informational content and optimizing clicks as the reference group, i.e., when both  $V_{ig}$  and  $E_{ig}$  are 0. To account for the endogenous difference between each organization and the nature of their advertising content, we added the organization-specific fixed effect. Since each organization’s campaigns were run during non-overlapping dates, the organizational fixed effects also account for the time fixed effect. The likelihood of having  $s$  positive engagement actions with  $r$  impressions for each aggregation unit  $i$  in organization campaign  $g$  is modeled as:

$$F(\beta'x_{ig})^s(1 - F(\beta'x_{ig}))^{r-s}$$

$$F(z) = \frac{\exp(z)}{1 + \exp(z)}$$

where  $x_{ig} = [1, V_{ig}]$  in the first set of analyses, in which we mainly look at the effects of changing campaign objectives, and  $x_{ig} = [1, V_{ig}, E_{ig}, V_{ig} \times E_{ig}]$  in the second set of analyses, in which we go a step further to look at interactions between campaign objectives and the timing of informational content. We hypothesize that ads under the view objective have much higher 3s video play rates and lower click rates compared to ads under the click objective. This hypothesis might seem counter-intuitive at

first glance because people generally expect more views will lead to more interest and hence more clicks and conversion. The main reason for such observations is the different types of audiences under the two campaign objectives. There are possibly many more true potential customers under the click objective compared to the view objective campaigns.

## 2.4 Main Results

We applied the binomial model specification to our data. We obtained the regression results in Table 2.2, with 3s video plays as the outcome in columns (1) and (2), through-plays in columns (3) and (4), and link clicks in columns (5) and (6). Columns (1), (3), and (5) correspond to the regression model with the campaign objective as the only treatment variable. We observe that the view objective has a significant positive impact on 3s video plays compared to the click objective in column (1), whereas it also has significant negative impacts on the through-plays and link clicks in columns (3) and (5). The main reason for the opposite directions of impact is more audience distortion under the view objective, and we will discuss this in detail in Section 2.4.1.

To alleviate audience distortion, one possible strategy is to change the creative design. Our approach is to vary the informational timing, and we incorporate this into the model. Columns (2), (4), and (6) correspond to the regression models with campaign objective, information timing, and their interaction term as treatment variables. The coefficients in the binomial regressions refer to the odds ratio, and the effects are not easy to interpret with interaction terms. Therefore, apart from the regression coefficient values, we present the predicted probabilities with 95% confidence intervals in Figure 2-3 and marginal predicted probabilities in Table 2.3 to illustrate our results.

Table 2.3 presents the marginal effects of changing the campaign objective from clicks to the view optimization objective under each information timing condition in columns (1) and (2), and the marginal effects of changing the timing of informational

content on predicted probabilities in columns (3) and (4). For example, the value of 0.025 with 0.008 standard error in the top row, column (3), can be interpreted as the increase in predicted probabilities when the informational content timing is changed from early to late under the click objective. In column (3), we observe the predicted probability of watching a video ad with late informational content for at least 3 seconds under the click objective is significantly higher than the same video ad with early informational content, by 0.025, and the through-play probability of late informational content under the click objective is also significantly higher than that of the early information setting. Columns (1) and (2) are also consistent with the observations in Table 2.2, with a significant positive effect on 3s play and a negative effect on link clicks when changing to the view objective. Late information videos have a significant negative marginal effect on the through-play probability, whereas early information videos have an insignificant effect.

We plot the predicted probabilities of each outcome action for all the treatment conditions in Figure 2-3. The x-axis shows the four treatment conditions. The y-axis represents the corresponding outcome metrics: A – 3s video play probability, B – 15s video play probability or through-play probability (we use the two terms interchangeably), and C – link click probability. As expected, we observe a significant gap in outcomes between the two campaign objectives. We also observe that late information has significantly higher 3s video play and through-play probabilities than early information under the click objective. Lastly, we find that the "Click Objective + Late Information" combination produces significantly higher 15-second view-through predicted probabilities than all the other three treatment conditions in Figure 2-3B, and the other three have similar predicted probabilities.

### 2.4.1 Possible Mechanism

The significant difference in click and view-related outcomes, as observed in the main results, is because changing the campaign objective will change the optimization objective of the platform's internal algorithms, resulting in a change in the audience group that a video ad is delivered to. Algorithms do not understand advertising

firms' fundamental business goal of increasing overall sales. Instead, their mission is to maximize the single action as specified in the campaign setup. When trying to achieve this single objective, algorithms do not care about or may even compromise the other objectives that the firm also essentially cares about. The distortion in audience groups is evident if we compare the three plots in Figure 2-3. The 3s view probability under the click objective is much lower than that under the view objective, whereas the pattern is reversed for the link click probability. If the audience distribution is the same in both the click and view objective campaigns, we should expect the patterns to remain the same for the 3s view, 15s view, and link click outcomes.

We believe there are two types of ad audiences on Facebook: information seekers and content consumers. Information seekers want more information about the product. They are also much more likely to take lower-funnel actions such as clicks and purchases. Content consumers mostly watch video ads for entertainment. It is rare for them to take any lower-funnel actions. Based on our definition, there is a much higher proportion of true potential customers among information seekers than content consumers. There are several clarification points. First, our classification refers to each time an ad was shown to an individual. The same person can be an information seeker for one ad and a content consumer for another. Moreover, the information seeker audience proportion will be relatively higher in the click objective condition than those in the view objective condition, but it does not mean information seekers necessarily make up the majority of the audience in the click objective condition. Lastly, information seekers are not necessarily always true potential customers. It is possible that a person is interested in charitable causes in general. After seeing the ad by this organization, they decide that they are not interested in this particular cause. In this case, they are an information seeker but not a potential customer. Similarly, content consumers might still become customers but the probability is much lower. We observe that ads under the view campaign objective tend to have higher 3s view probabilities because they are shown to a much higher proportion of content consumers than under the click objective. We also observe that ads under the click objective tend to have much higher click probabilities, similarly because these ads

are delivered to a relatively higher proportion of information seekers as found by the optimization algorithms.

In Figure 2-4, we present the age distributions of the campaign audience under the view objective (A) and the click objective (B), and the age distribution of these organizations' Facebook page followers (C), which we believe is a proxy for the true underlying audience. We observe a high proportion of young people among page followers, whereas the proportion of young people under the view objective is extremely low. Under the click objective, the age groups are more evenly distributed. These observations make the click objective audience distribution look more similar to the page follower distribution than the view objective audience distribution.

To get a more quantitative comparison between the three age distributions, we also calculate the KL divergences between the audience age distributions under the two campaign objectives and the page followers' distribution, respectively. KL divergence is formally defined as  $D_{KL}(p||q) = \sum_{i=1}^N p(x_i) \log \left( \frac{p(x_i)}{q(x_i)} \right)$  and we use the Facebook page followers' age distribution as  $p$  as we consider it to be a proxy for the true audience distribution, and we take the click or view campaign audience age distributions as  $q$ . The KL divergence of the page followers' distribution and audience under the click objective is 0.246, whereas the KL divergence of the page followers' distribution and audience under the view objective is 0.828. Recall that the lower the KL divergence value, the better we have matched true distribution  $p$  with approximate distribution  $q$ . The lower KL divergence value between page followers and the click campaign audience further confirms that there are more true audiences in click objective campaigns than view objective campaigns.

An alternative explanation is that, under the click objective, instead of more information seekers there are just more bad consumers who simply click without watching the video. However, we would like to argue against this based on two factors. The first is that the through-play probability under the click objective is significantly higher than that under the view objective. If there are genuinely many bad consumers who blindly click without watching then we should expect the opposite pattern for through-play probability. However, we obtain the average video play duration is 3.95s

under the click objective and 3.34s under the view objective. This finding further confirms that the audience under the click objective on average watches video ads for a duration similar to the audience under the 2s continuous view objective, if not longer.

Our research also presents an empirical setting for a multitask principal-agent problem. Firms or advertisers generally have various goals in mind when they advertise products or services on digital platforms, including high-funnel goals such as raising awareness through more video views and mid-funnel goals such as raising interest or consideration through more traffic. Meanwhile, the ultimate goal of the firm's advertising effort is to increase sales and profits. Therefore, we could view this as a multitask principal-agent case with the firm or advertiser being the principal and the advertising platform being the agent. The basic idea of the multitask [29, 19] is that when multiple tasks are critical for the principal but only one task is easily measurable and compensated for, naturally the agent will focus their efforts on that one single task and disregard other tasks that are equally important to the principal but are provided with no or extremely low incentive. Similarly, in this case, since the firm is only allowed to specify one optimization objective, the platform will abandon all efforts toward other objectives, including the firm's ultimate sales maximization goal.

Next, we move on to see whether changing the creative design strategy, and specifically the timing of the informational content, might help alleviate the audience distortion or provide any insights. Comparing the timing of the informational content, we observe that late informational content has a significantly higher 3s view rate than early informational content when the campaign objective is to optimize for link clicks in Figure 2-3A. The gap is because information seekers care more about the key information and are more patient in waiting for key messages. This does not necessarily mean that information seekers prefer to watch videos with late informational content, as the link click probabilities are very close in the two treatment conditions. Another interesting observation is that the late information content videos also have a significantly higher 15s view rate than early information content videos under the

click objective. This observation further confirms that information seekers tend to be more patient with key ad messages.

## 2.5 Follow-up Experiments: Age Appeal as an Additional Treatment Variation

Based on the heterogeneous analysis in the next section, we find that one way for algorithms to achieve campaign objective optimization is to deliver ads to older people. We explore whether tailoring content to a younger population might alleviate this particular optimization bias in the follow-up experiment. We employed the  $2 \times 2 \times 2$  experimental design on four organizations. For each organization, we ran the campaign with eight treatment conditions. We kept the original  $2 \times 2$  treatment variation in our design, and the additional treatment variable was whether the video creative included more scenes of younger vs. older people. The key idea was to compare the audience age distribution between the video that included more scenes of young people and the video that included more scenes of older people under the same campaign objective and timing of product information. We hypothesize that, in general, the appearance of older people appeals to an older population and the appearance of younger people appeals to a younger population.

Moreover, we understand that the creative strategy variation might also potentially cause audience distribution differences due to the platform's algorithmic optimization. Specifically, the platform may deliver the early info version to people who favor early informational content, and the late info version to people who favor late information. Therefore, in our follow-up study, we conducted another set of experiments in which we changed another creative design aspect – age appeal – which could cause a more measurable difference in audience distribution. By measuring how much distortion was created in this measurable aspect, we could get a sense of whether the informational content timing treatment posed a significant risk to our identification.

Figure 2-5 shows the audience age distributions between video creative designs



with older people and younger people. The four plots correspond to four different treatment conditions. In Table 2.4, we conducted a  $\chi^2$  test to learn whether the demographic distribution was independent of the variation in age appeal creative design. Based on Figure 2-5 and Table 2.4, we observe that there is generally not much difference between age distributions for different age appeal creative designs, except for the treatment with the view objective and early informational content which has a significant  $\chi^2$  value. However, even for the view objective and early info treatment, we only observe a slight difference with younger creative design delivered to slightly older people.

We would like to investigate whether tailoring content to a younger vs. older population might change the overall audience distribution by assessing whether the age distribution becomes younger when tailoring content to young people. However, based on the current results in Figure 2-5 and Table 2.4, we do not observe any such pattern suggesting a large and consistent shift in audience demographic distribution. Moreover, we also intend to evaluate how large the optimization bias is with variation in creative design. Based on our results, it seems the bias stemming just from creative design change is far smaller, if not nonexistent, than from campaign objective change.

## 2.6 Robustness Checks

There are several potential doubts about our research in terms of its generalizability, which we would like to address in the robustness checks.

### 2.6.1 Additional Advertising Outcomes

Apart from the main outcome metrics we have highlighted – i.e., 3s plays, through-plays, and link clicks – we further explored how the predicted probabilities of other view-related metrics differ for the four treatment conditions in Figure 2-6. The figure presents five additional outcomes: 25%, 50%, 75%, 95% and 100% video plays. The x-axis shows the predicted probabilities of the five video play duration outcomes, and the four different colors show the predicted probabilities of the four treatment

conditions. The pattern of 25% video plays closely follows what we observe for 3s video plays. The large gap between the two campaign objectives for 3s video plays and 25% video plays decreases immediately when it comes to 50% or longer-duration video plays. The pattern of 75% video plays also resembles what we observe for through-plays, and it persists for 95% and 100% video plays.

## 2.6.2 Accounting for Time Variation

All the campaigns took place on different and non-overlapping dates, ranging from March to November 2020. Each campaign lasted for three days, starting at 12 am Pacific Time and ending at 11.59 pm on the third day. We added organization-specific fixed effects in all of our models to account for the variation in results due to different times of year. For a robustness check, we also added two additional time-related fixed effects: the seven days of the week and dates. Table 2.6 presents the binomial regression coefficients on link clicks and 3-second video plays. The difference in regression coefficients is negligibly small when adding the additional time fixed effects, and hence our results are robust to variation in campaign timings.

## 2.6.3 Adding the View-through Campaign Objective

Another concern is that the bias brought by the 2-second continuous view campaign objective is an extreme case. People may intuitively think maybe a 2s view is too short, and if we optimize for through-plays instead, more information seekers will be targeted even under the view objective. We therefore include an additional campaign objective – 15s view-through to form a  $3 \times 2$  experimental design for three organizations.

As shown in Figure 2-7, although the view-through campaign objective achieves the highest 3s video plays and 15-second video plays, the click rate is not significantly different from the 2s continuous view objective. This observation suggests that view-through objective campaigns are most likely to target a different kind of content consumer who watches videos for a longer duration.

## 2.7 Audience & Campaign Characteristics Driving the Differences

Although Facebook does not provide individual-level data for user privacy reasons, we could still obtain grouped data characterized by demographics, organizational type, and time breakdowns. In this subsection, we investigate how the observed audience and campaign characteristics could drive ad performance differences. Specifically, we explore what types of organizations naturally attract more content consumers, which age groups are more patient, and how advertising results vary with time of day and campaign duration.

### 2.7.1 Organization Type

This subsection further elaborates on higher audience distortion in 3s view rates across view objective campaigns. Although there is no significant difference in the 3s play probability between early information videos and late information videos under the view objective, the effect difference could be more prominent when we classify organization types into local appeal vs. national appeal.

The organizations with local appeal mainly focus on people in limited local areas in the US. The true audience is generally confined to one state, most likely just a few counties. These organizations provide education, food, grocery, medical, or community services to local areas. For example, Daily Table is a nonprofit organization that provides affordable groceries to those in need near the Boston area in Massachusetts. We requested the audience distributions from each local organization's Facebook page, and six out of eight organizations provided us with the audience location distributions. 40% of followers of all these local organizations were located within the same state in which the organization provided services, among Facebook users who had revealed their city location. The organizations with national appeal fell under two subtypes. The first subtype refers to the nonprofit organizations that work on international causes. Action for Cheetahs in Kenya is a Kenya-based nonprofit or-

ganization whose mission is to protect cheetahs in Kenya. Another subtype refers to national nonprofit organizations that are active in multiple varying locations within the US. For example, Consult Your Community helps to provide consulting services to small businesses across the entire US, not limited to any single region. Based on the characteristics of organizations' audience appeal, we expect organizations with a national appeal to have more interested audiences as their missions can be relevant to a broader audience base.

Figure 2-8 presents the predicted probabilities of each outcome action for both types of organizations under the four treatment conditions. Each plot corresponds to an outcome action. The comparison between the two campaign objectives is evident based on the plots, and they are consistent with what we observe for the main effects. Moreover, under the view objective, the 3s video play and through-play probabilities for local organizations are far higher than those for organizations with national appeal. When comparing the early vs. late information timings, we also calculate the marginal effects of changing from early to late information timing given the same campaign objective in Table 2.5 to get a more accurate comparison. One interesting observation in Figure 2-8 and Table 2.5 is that, under the view objective, the 3s video play probability for early information video ads is now significantly higher than late information video ads for local organizations but significantly lower for national organizations. Such an observation is consistent with our expectation that there is a comparatively higher proportion of information seekers for ads with national appeal, under the view objective. It is supported by the observation that national appeal organizations have much lower view rates but similar click rates. Moreover, the fact that the late information has a higher 3s video play probability for national organizations further confirms our hypothesis, as it is consistent with the "more patient in waiting for the key message" characteristic we described for information seekers in our proposed mechanism.

## 2.7.2 Age

We are interested in exploring how the effects might differ across age groups. The age groups are defined based on the groupings that Facebook provides: 18 – 24, 25 – 34, 35 – 44, 45 – 54, 55 – 64 and 65+. Figure 2-9ABC presents the probabilities of 3s plays, through-plays, and link clicks, respectively. Each age group is represented by a different color, and 95% confidence intervals are also included.

Figure 2-9 shows the predicted probabilities of the four treatment conditions for each age group. Each plot corresponds to one outcome measure. In Figure 2-9AC, we observe that people under age 45 generally have a significantly lower probability of a 3s view or link click compared to people over 55 years old under the click objective. Across all three outcomes, the common pattern of click objective campaigns regardless of information timing is that, as people age, there is an increasingly higher probability of outcome actions. However, when we look at the view objective results, we observe no consistent or significant pattern between different age groups, confirming that the audiences under the two different campaign objectives are different.

## 2.7.3 Viewing Time

Figure 2-10 presents the hourly change in terms of click rates and 3s view rates. Figure 2-10AB shows the aggregated data of all organizations. The vertical blue separators in these two plots indicate the end of a day with the cutoff at midnight.

In Figure 2-10AB, the click objective campaigns deliver ads at a more reasonable time. If we look at the confidence bands across time, click objective videos have more variations in the confidence interval width than view objective videos. Specifically, click objective videos have wider confidence intervals during late night and early morning and narrower intervals during the daytime, suggesting that the click objective campaign is more selective of the audience by changing its frequency of delivery based on time. In contrast, the confidence bands of the view objective treatment conditions do not vary a lot. The 3-second view rates of both the view objective and click objective videos in Figure 2-10A are converging as click objective videos improve their

3-second video play rates over time. This observation implies that the optimization algorithm probably self-improves during the campaign, especially for click objective campaigns.

Almost all the campaigns ran for three consecutive days except for one campaign, which we ran for ten days because we wanted to know how the pattern might change when the campaign lasted for a longer duration. The variations across the 10-day period for this particular organization are shown in Figure 2-10CD. The vertical black line indicates the end of the first three days. Based on Figure 2-10C, it seems the algorithms keep adjusting. The 3s view rate gradually converges to a similar level for all four conditions.

Figure 2-11 compares the aggregated 3s view and click probabilities in the first three days and the remaining seven days for this particular organization. There is a significant increase in 3s views under the click objective, and the increase is particularly drastic for early information videos. The results further confirm the observations in Figure 2-10CD that the ad delivery algorithms self-adjust over time, and hence the pattern might change. In general, ad delivery is better at optimizing for campaign objectives as campaigns last for a longer time. This observation is expected as algorithms now have more accurate training data. However, since this is only one organization, the results might not be conclusive or generalizable enough.

## 2.8 Conclusion

Through a field experiment conducted as multiple sets of Facebook ad campaigns, we find that algorithms are incredibly apt at achieving pre-defined targeting objectives. However, to achieve this objective, the algorithm might deliver the ads to essentially different audiences in a way that the advertiser does not anticipate. With higher-funnel objectives such as view optimization, more content consumers will be targeted. In contrast, lower-funnel objectives such as click optimization target more information seekers and these campaigns tend to be more expensive. The audiences under the two types of campaign objectives are different, based on the contrasting

view and click behavior that the audience under the view vs. click campaigns exert. Moreover, we find that changing video creative design strategies, in particular the timing of informational content from early to late, might improve views under a traffic optimization campaign objective while maintaining clicks. Given the algorithmic tradeoff in audience selection, a comparatively better approach to achieve both clicks and views is to choose the click objective but present product information at a later time.

The main managerial insights lie in the campaign publishing strategies for firms that want to reach their true potential customers. Advertisers might prefer a high-funnel campaign objective such as view optimization because it is usually cheaper and can achieve much higher impressions and view rates than other objectives. Moreover, advertisers might also associate a high view rate with a higher conversion rate. However, they might not be aware of the highly different audience groups that the ad could reach under different campaign objectives. The most direct solution to this problem is to specify a low-funnel objective, which helps the ad to reach a larger group of information seekers, including more potential customers. In addition, the multitask principal-agent problem arises when the firm sometimes still wants to have decent impressions and video views for other marketing purposes, or when the low-funnel metrics are difficult to track and measure. One possible solution is that platforms could offer firms options to specify multiple objectives. Another possible solution is to offer a voluntary bonus payment structure instead of a simple piece-rate contract. Research has provided empirical evidence that this might work when concerns about fairness and reciprocity come into play [23]. The two solutions provide platforms with either explicit or implicit incentives to implement better multitask learning algorithms to account for multiple outcomes [11, 57, 74].

In terms of creative video design, many advertisers worry that a large proportion of the audience might not even know what the video is advertising if the key ad message does not appear at the start of the video, and the skip rate will be high at the start. Consequently, advertisers are paying for no gain. We hope to alleviate such concerns as there is no significant difference between early or later informational

content for lower-funnel metrics. We even find that later informational content could slightly boost initial views for traffic campaigns, although there is not much effect on clicks.

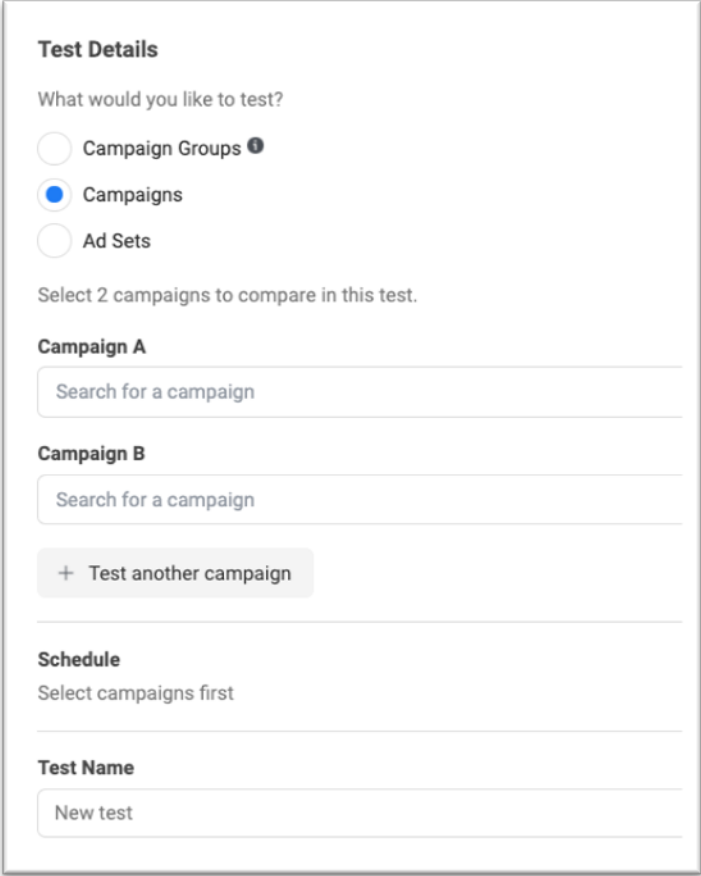
There are limitations to our research. First, we focus on a single platform – Facebook – and other platforms may have different advertiser objectives limiting the generalizability of our results. Second, our experiment was conducted with small organizations, so it may not generalize to larger brands. Third, as we do not have Facebook’s exact algorithm, our explanations are inferential rather than based on being able to explore the actual code of the algorithm. Notwithstanding these limitations, however, we believe this chapter is a useful first step in trying to understand some of the consequences of algorithm-driven advertising impressions.



## 2.9 Appendix

### 2.9.1 Figures

Figure 2-1: Screenshot of Facebook Experiment Platform Setup



**Test Details**

What would you like to test?

Campaign Groups <sup>1</sup>

Campaigns

Ad Sets

Select 2 campaigns to compare in this test.

**Campaign A**

Search for a campaign

**Campaign B**

Search for a campaign

+ Test another campaign

---

**Schedule**

Select campaigns first


---

**Test Name**

New test

Figure 2-2: Screenshot of Published Facebook Ad

Refresh Report

 **Sankofatech**  
Sponsored · 

...

 **19% of students graduating with computer science degrees were black**

SANKOFATECH.ORG  
**Become-a-Volunteer** [LEARN MORE](#)


 Like  Comment  Share

Figure 2-3: Predicted Probabilities with 95% C.I. in Different Treatment Conditions

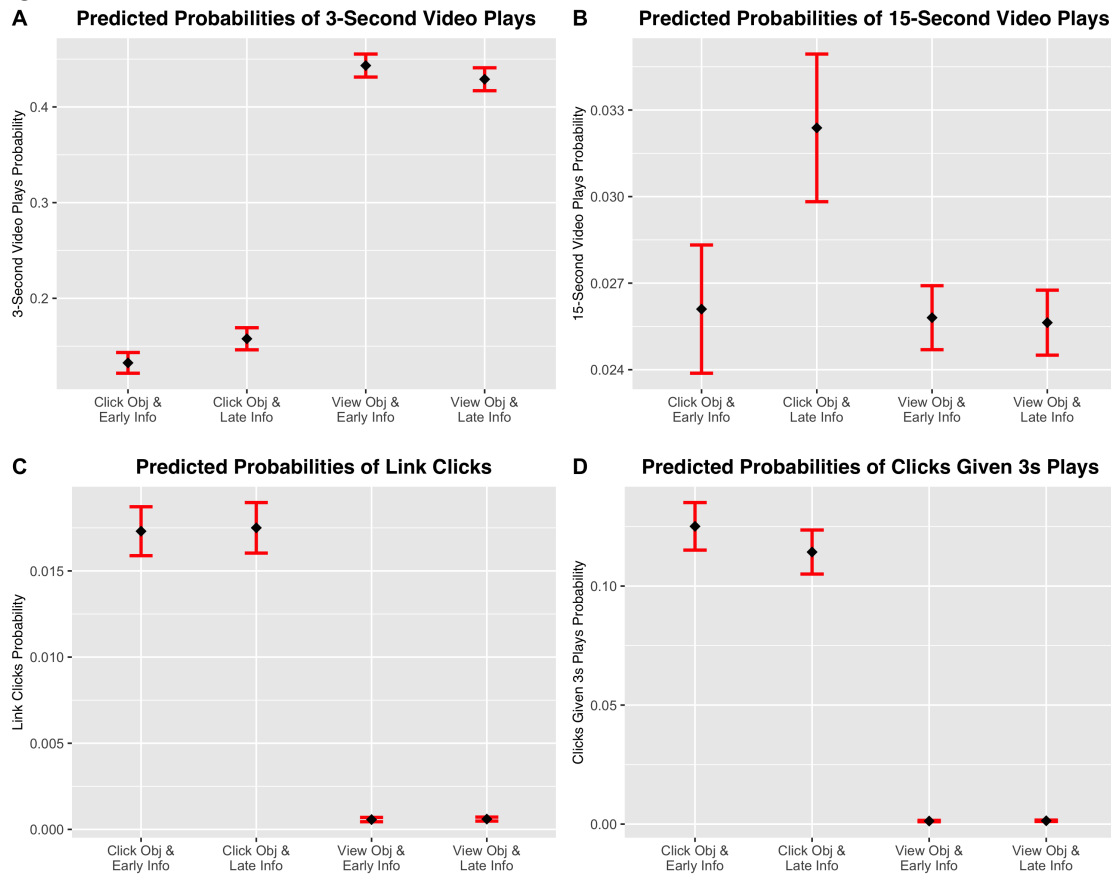




Figure 2-4: Age Distribution of Campaign Audience and Page Followers

**A** Audience Distribution Comparison by Age Group w/ View Obj.



**B** Audience Distribution Comparison by Age Group w/ Click Obj.



**C** Fb Page Audience Distribution Comparison by Age Group




Figure 2-5: Comparison in Audience Age Distribution Between Younger vs. Older Content Tailoring

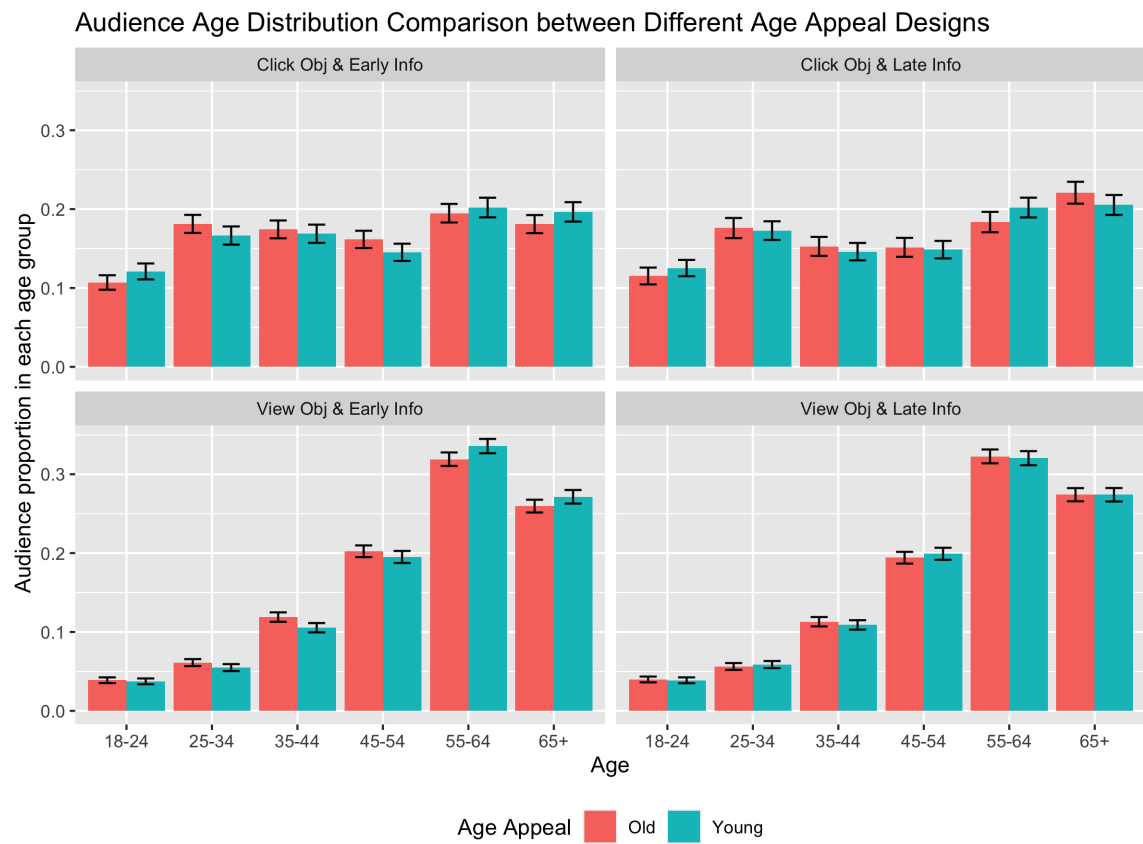


Figure 2-6: Predicted Probabilities of % Video Plays with 95% C.I. in Different Treatment Conditions

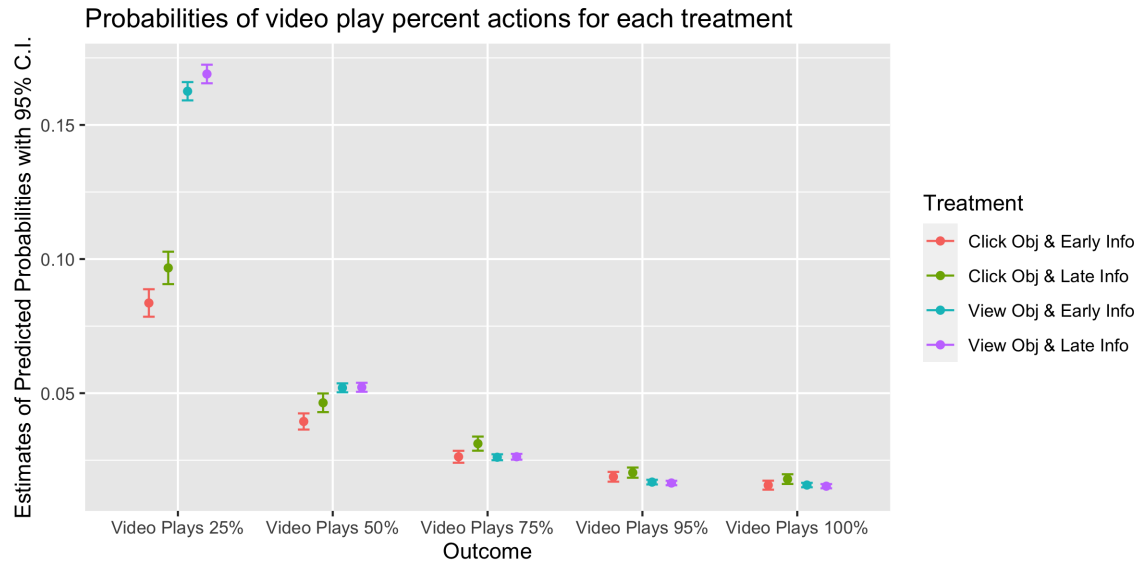


Figure 2-7: Predicted Probabilities with 95% C.I. in Different Treatment Conditions with View-through Optimization as an Additional Campaign Objective Treatment Condition

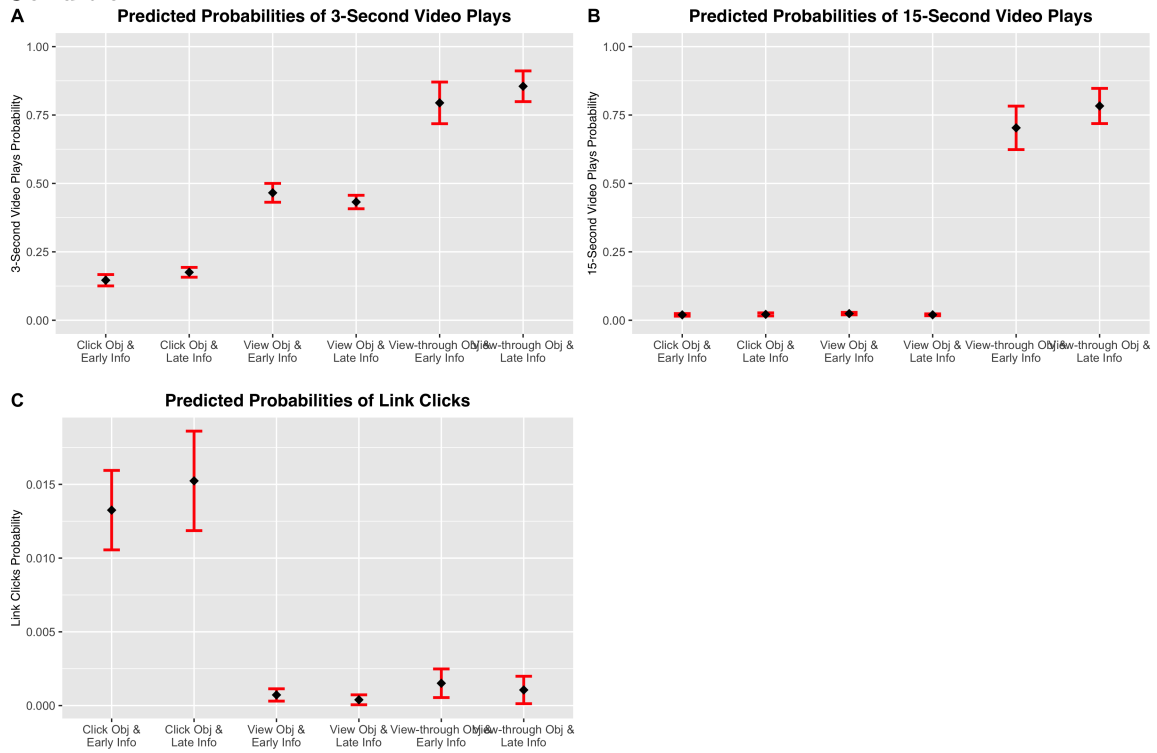


Figure 2-8: Predicted Probabilities with 95% C.I. in Different Treatment Conditions by Local vs. National Audience Appeal

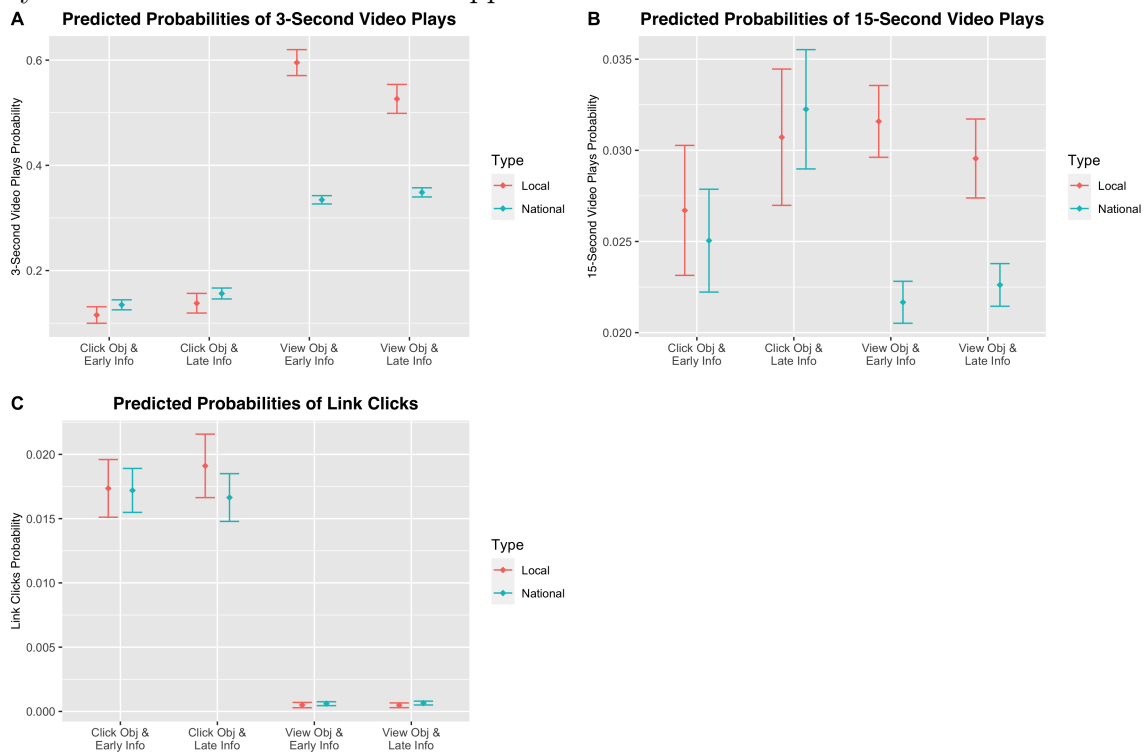




Figure 2-9: Predicted Probabilities with 95% C.I. in Different Treatment Conditions by Age

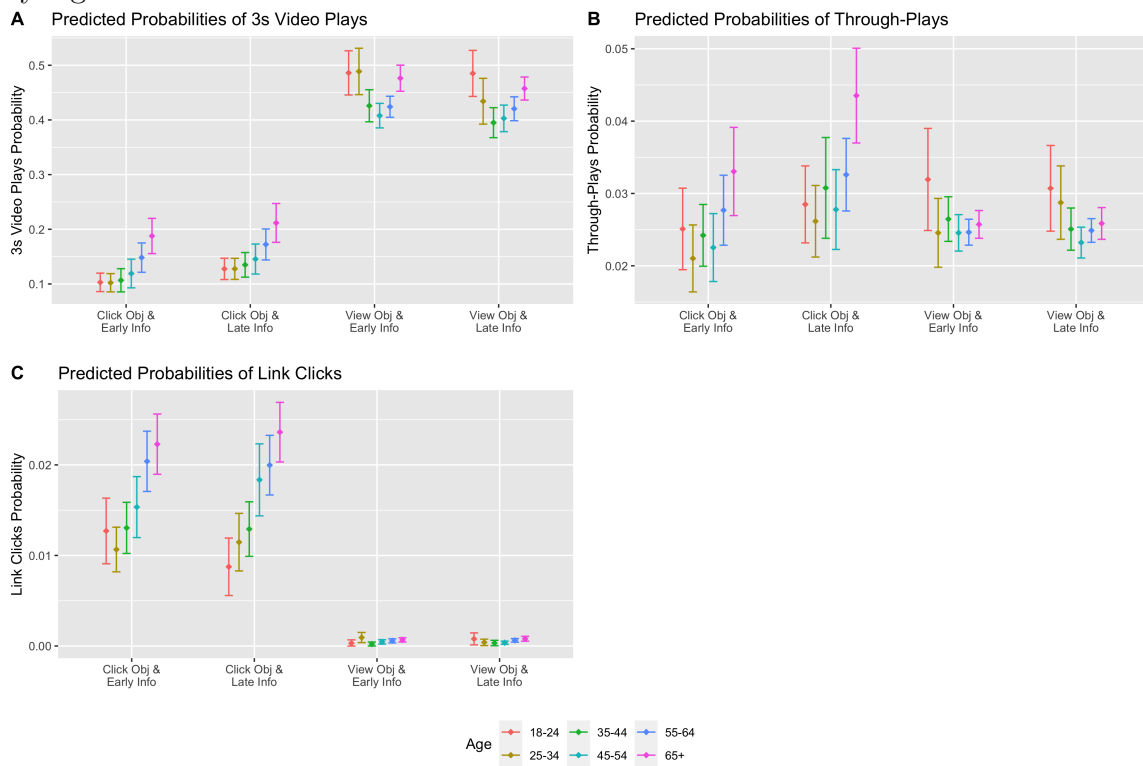


Figure 2-10: Hourly 3s View and Click Rates for Three Days (A & B) and Ten Days (C & D)

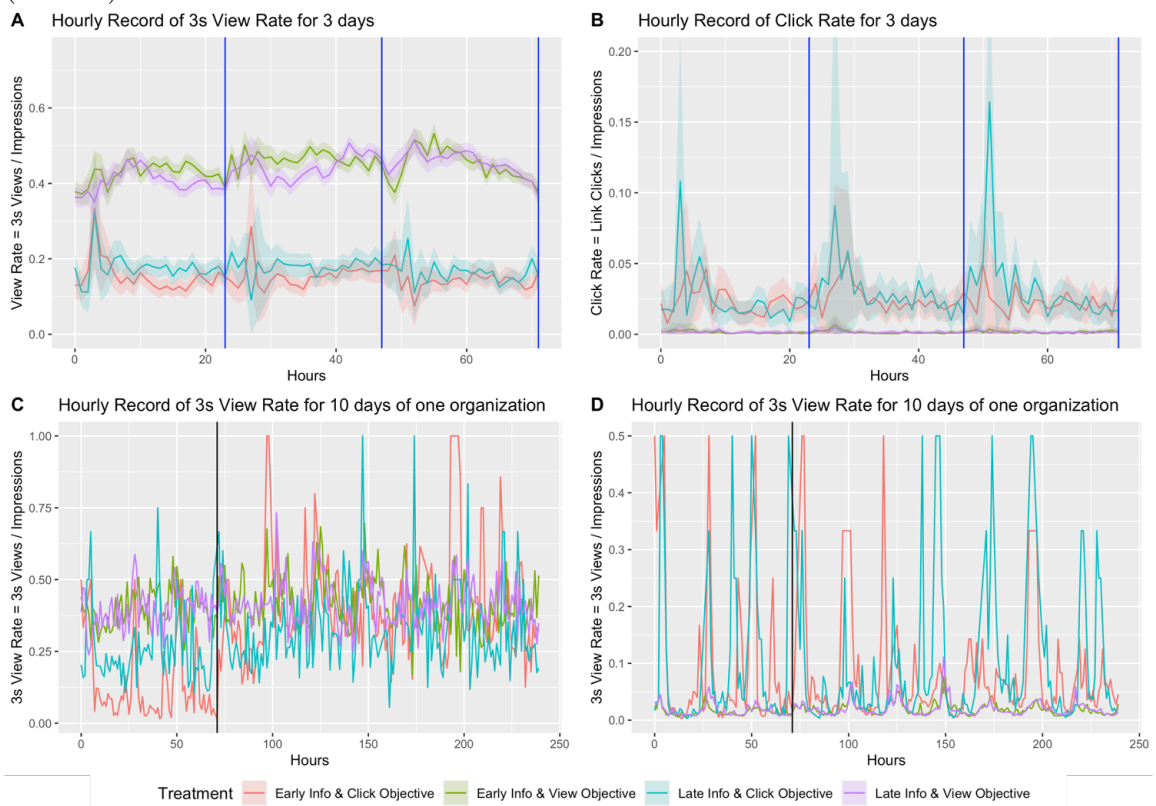
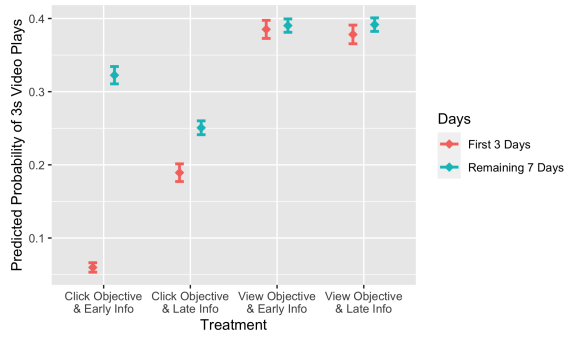



Figure 2-11: Predicted Probability Comparison Between First Three Days and Remaining Seven Days

**A** Three-second view rate comparison between the first 3 days and the remaining 7 days



**B** Link click rate comparison between the first 3 days and the remaining 7 days



## 2.9.2 Tables

Table 2.1: Summary Statistics of Outcome Metrics:

Metrics	Count	Percent/Rate
Impression	377,848	-
Impression (View Obj)	279,521	74.0%
Impression (Click Obj)	98,327	26.0%
3s Video Play	131,878	0.349
Through-Play	12,058	0.032
Link Click	2,014	0.005

*Note:* In the last column – Percent/Rate, the statistics for impressions under either view or click objective are in percentage, and the statistics for the other outcomes are in rate. Both percent and rate metrics are obtained by dividing the positive action count by the total impressions.

Table 2.2: Table of Main Binomial Regression Results

	<i>Dependent variable:</i>					
	3s Video Plays		Thru-Plays		Link Clicks	
	(1)	(2)	(3)	(4)	(5)	(6)
View Obj.	1.519*** (0.011)	1.651*** (0.016)	-0.131*** (0.024)	-0.012 (0.035)	-3.411*** (0.075)	-3.429*** (0.108)
Late Info		0.203*** (0.020)		0.222*** (0.041)		0.012 (0.049)
Late Info × View Obj.		-0.262*** (0.022)		-0.229*** (0.047)		0.037 (0.150)
(Intercept)	-1.776*** (0.177)	-1.879*** (0.178)	-3.504*** (0.104)	-3.619*** (0.106)	-4.034*** (0.131)	-4.039*** (0.133)
Observations	3,982	3,982	3,982	3,982	3,982	3,982
Log Likelihood	-20,429	-20,354	-5,971	-5,956	-2,386	-2,386

*Note:* In columns (1) and (2), dependent variable is the 3s video play. In columns (3) and (4), dependent variable is the through-play or 15s video play. In columns (5) and (6), dependent variable is the link clicks. The model incorporates organization-specific fixed effects. Standard errors are in parentheses. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 2.3: Marginal Effects of Changing One Treatment Variable While Keeping the Other Treatment Variable Constant

	Objectives from Click to View		Info Timing from Early to Late	
	Early Info (1)	Late Info (2)	Click Obj (3)	View Obj (4)
3-Second Plays	0.311*** (0.009)	0.271*** (0.009)	0.025** (0.008)	-0.014 (0.009)
Through-Plays	-0.0003 (0.0013)	-0.0068*** (0.0014)	0.0063*** (0.0018)	-0.0002 (0.0008)
Link Clicks	-0.0167*** (0.0007)	-0.0169*** (0.0007)	0.0002 (0.0010)	0.0000 (0.0001)

*Note:* The dependent variable of columns (1) and (2) is the difference in predicted probabilities of each outcome action when campaign objective changes from link click optimization to 2s view optimization, given either early or late informational content creative design. The dependent variable of columns (3) and (4) is the difference in predicted probabilities of each outcome action when ad creative changes from early informational content to late informational content, given a specified campaign objective. Standard errors are in parentheses. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 2.4: Chi-square Test Results of Demographic Age and Gender Distributions Between Older vs. Younger Appeal for Each Treatment Group

Campaign Objective	Timing of informational content	
	Early Info	Late Info
Link Click	18.3	16.9
Video View	32.2***	18.5

*Note:* \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$

*Note:* Statistics in the table are Chi-square values with 11 dof between the creatives with differing age appeals, given campaign objective and information timing. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 2.5: Table of Marginal Effects When Changing from Early Informational Content to Late by Organization’s Audience Appeal

	Local Appeal		National Appeal	
	Click Obj (1)	View Obj (2)	Click Obj (3)	View Obj (4)
3-Second Plays	0.0222 (0.0120)	-0.0670*** (0.0178)	0.0233** (0.0078)	0.0174** (0.0060)
Through-Plays	0.0054* (0.0025)	-0.0016 (0.0015)	0.0066** (0.0023)	0.0009 (0.0008)
Link Clicks	0.0013 (0.0015)	0.0000 (0.0001)	-0.0005 (0.0013)	0.0000 (0.0001)

*Note:* Dependent variable, marginal effects, is the difference in predicted probabilities of each advertising outcome action when we change the informational content timing from early to late. Marginal effects in columns (1) and (2) refer to organizations with local appeal, and columns (3) and (4) refer to organizations with national appeal. Standard errors are in parentheses. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .



Table 2.6: Binomial Regression Coefficients of 3-second Video Plays and Link Clicks

	<i>Dependent variable:</i>					
	3-Second Video Plays			Link Clicks		
	(1)	(2)	(3)	(4)	(5)	(6)
Late Info	0.206*** (0.020)	0.207*** (0.020)	0.209*** (0.020)	0.012 (0.049)	0.015 (0.049)	0.019 (0.050)
View Obj.	1.651*** (0.016)	1.656*** (0.016)	1.673*** (0.016)	-3.429*** (0.108)	-3.430*** (0.108)	-3.443*** (0.107)
Late Info × View Obj.	-0.264*** (0.022)	-0.263*** (0.022)	-0.265*** (0.022)	0.037 (0.150)	0.033 (0.150)	0.030 (0.150)
(Intercept)	-1.878*** (0.178)	-1.898*** (0.185)	-1.894*** (0.108)	-4.039*** (0.133)	-4.047*** (0.138)	-4.028*** (0.084)
Campaign FE	✓	✓		✓	✓	
Weekday FE		✓			✓	
Date FE			✓			✓
Observations	6,072	6,072	6,072	6,072	6,072	6,072
Log Likelihood	-24,766	-24,455	-23,606	-3,264	-3,263	-3,277

*Note:* In column (1)-(3), dependent variable is the 3s video plays. In column (4)-(6), dependent variable is the link clicks. Standard errors are in parentheses. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .



# Chapter 3

## How Does the Severity of the COVID-19 Pandemic Affect Digital Advertising?

### 3.1 Introduction

At the beginning of 2020, the world was hit by the COVID-19 outbreak, which subsequently resulted in a global pandemic. Due to its highly infectious nature and relatively high fatality rate, people had to practice shelter-in-place and other preventive health measures to reduce the risk of becoming infected. When more people got infected with the COVID-19 virus at the same time, the number of hospitalizations and deaths also went up correspondingly, which took up more medical resources. The lack of healthcare resources further increased the death toll due to a lack of medical care, and hence the pandemic situation became worse. People's self-isolation practice also varied as the severity of the pandemic changed. The more severe the pandemic seemed, the more cautious people were about going out. The constant variation in the COVID pandemic's severity certainly disrupted people's everyday lives and changed their behavior and emotions. Given the widespread nature of and difficulty in eradicating the COVID-19 virus, it is imperative to understand whether and how

many people have been affected from different perspectives, so that we can be better prepared when facing future waves of the mutating virus. Our research would like to study people’s behavioral change through the lens of consumer behavior, and particularly their responses to digital ads.

The main research question that we study is whether and how the variation in perceived COVID-19 pandemic severity impacts people’s response to digital advertising. We gathered observation data from a collaborating digital ad creative and analytics firm. We managed to obtain the firm’s Pinterest ad data for three years, from 2019 to 2021, within the US. Section 3.3 describes the Pinterest ad performance data and creative element data, as well as the COVID data that reflect or associate with the severity of the pandemic, with time and location variations. The pandemic officially hit in March 2020, so we use the 2019 data as a baseline. We employ difference-in-differences and fixed-effects models, and we take 2s views, conversions defined by us, and conversions defined by the firm as the outcome variables.

In Section 3.5, we find that the increase in the stay-at-home percentage, which serves as a proxy for pandemic severity, caused an increase in conversions. The impact of increasing severity was especially prominent in the initial outbreak of the pandemic. The effect is expected, given that people had more restrictions on going out and thus tended to replace offline shopping with online. Meanwhile, we have identified two other variables with significant moderating effects – vaccination rate and the presence of people in the ad creative. We find that the vaccination rate has a negative moderating effect on views and conversions, and the presence of people in ads has a positive effect on views. The significance of the moderating effects implies that this is not just a simple replacement effect. In particular, the moderating effect of the presence of people implies that the ad audience prefers ads with people in the ad creative – very likely associated with loneliness due to self-isolation. Lastly, we explore the heterogeneous effects in Section 3.6. The ads of retail industries experienced an increase in conversions as the pandemic became more severe. Moreover, the travel industry took a hit in its advertising performance as the pandemic worsened.

Overall, we hope our research helps firms make better strategic marketing decisions

in a time of crisis when people are experiencing physical and psychological disruptions. Not only does the strategy apply to when and how much to advertise, but also how to make ads more appealing during such unsettling times by incorporating certain creative elements.

## 3.2 Literature Review

Our research contributes to three streams of academic literature. The first research stream concerns the impact of COVID-19 on people’s mental status and perception of the pandemic’s severity. Psychologically, people experienced a huge shock during the COVID-19 pandemic – a multitude of negative emotions has been found to be associated with the pandemic [10]. Research has identified the feeling of disgust towards others with disease symptoms [25]. Fear of getting infected also played an influential role in affecting consumer shopping behavior during the COVID-19 pandemic [22, 25]. Depression, anxiety, and stress, mainly arising from the lockdown and social distancing, are undoubtedly significant factors impacting people’s daily life [64, 54, 50]. Research has also shown that actual severity does not necessarily align well with perceived severity [33]. In fact, they are only moderately associated, and perceived severity tends to predict COVID-related distress and behavior better [58]. The literature provides a solid foundation for constructing our models. Specifically, we decided to use the perceived severity of the pandemic rather than the actual severity as the independent variable. It also motivated us to explore the psychological factors of the pandemic more deeply by inspecting the moderating effects of vaccination and ad creatives.

Second, our research also addresses the impact of the COVID pandemic on consumer behavior. Unusual purchases and self-isolation were two types of behavioral responses during the COVID-19 pandemic, with a strong link identified between the two. Both are driven by the perceived severity of the pandemic [38]. For example, negative information about COVID-related deaths increased perceived risk and hence stockpiling [26]. The temporary budget contraction due to the outbreak of

the pandemic would result in people re-evaluating their preferences and purchasing fewer products, even after the budget was restored [55]. Sometimes even the opposite behavior has been identified in the literature. Research shows that the pandemic made people relatively prefer more atypical products because they implicitly associate common products with many people [31]. Externally, the closure of stores during the pandemic disrupted people’s original purchasing habits, potentially leading to more option exploration and hence unusual purchases [68]. On the other hand, research also shows that people choose familiar products to reduce uncertainty [25]. It is also found that consumers might prefer natural products with lay theories that natural products are safer [59]. Our research contributes to this research stream by empirically exploring the impact on conversions and how that impact differs across different categories of products using empirical data.

Lastly, our research aligns well with the impact of the pandemic on digital marketing. We find a few research studies that focus on this topic. An exploratory study presented survey evidence showing consumers have increased their use of social media as a way to identify products, evaluate them and make purchases [42]. Under the threat of the pandemic, consumers tended to favor products with authentic advertising messages due to the intention to reduce uncertainty [48, 34]. The informational value of ads has a positive effect on online engagement among utilitarian and hedonic products [13]. Temporal framing in advertising messages could also sway people’s behavior [36]. Most current research focuses on advertising content and employs either lab experiments or survey methods. We would like to fill the gap by providing large-scale empirical evidence on the relative long-term impact of the pandemic on digital advertising. Moreover, we further explore the impact of ad creatives during the pandemic crisis.

### 3.3 Data Description

This section presents all the relevant data for our models and analysis. The first subsection introduces the Pinterest ad data, including the ad campaign data, the ad

creative element data, and ad performance data that we use as the dependent variables in our causal model. We explore some key ad campaign statistics, with industry and campaign objective breakdowns. The second subsection describes the mobility data, which we use as the primary treatment variable in our causal model. The last subsection describes the COVID-related data that correlate with the treatment variable.

### 3.3.1 Pinterest Ad Data

We use Pinterest ad performance data ranging from January 2019 to December 2021, collected from our collaborating advertising firm. During the three years, there were a total of 1,656 unique campaigns and 8,684 unique media assets within the US. The ads include both image and video, with around 70% of the ad creatives being images. Figure 3-1 shows the distribution of campaign industries and campaign objectives. In the upper figure, the top three industry categories are consumer packaged goods, retail, and automotive. In the bottom figure, the most popular campaign objectives include awareness, traffic, web conversion, and consideration. Note that the Pinterest ad performance data we collected is confined to the campaigns in the collaborating firm's database. It does not represent the distribution of all ad industries or campaign objectives on the platform.

In Table 3.1, we summarize some important campaign-related variables. The impressions refer to the number of times the ad has been shown to Pinterest users. The spend refers to the total expenditure for each ad campaign in US dollars. The mean is much larger than the median in both cases, indicating that both ad impressions and expenditure data are right-skewed. In terms of the outcome variables, we mainly look at three different types of ad performance metrics for our analysis. 2-second views are defined by the number of times the video ad or re-pins of the video ad played continuously for 2 seconds while at least 50% in view. The 2s view rate in Table 3.1 is obtained by dividing 2s views by impressions for each campaign. Again, we observe that the data is very right-skewed. Conversion refers to lower-funnel conversion actions, including add-to-carts, checkouts, and sign-ups. Similarly, the conversion rate

for each campaign in the fourth row is obtained by dividing conversions by impressions. The second conversion metric in the last row of the table, conversion (FD), refers to the firm-defined conversion metric. The ad's corresponding firm defines the conversion action while setting up the advertising campaign. Admittedly, one limitation of the metric is that the specific conversion action is unknown to us, and that might add some uncertainty and inconsistency to this metric. However, we included this firm-defined conversion metric because it is relatively more evenly represented in all types of ad campaigns, as we will explain in the next section.

Table 3.2 shows a summary of the ad performance metrics with industry breakdowns. The industries are in order of their proportions among all campaigns, from the highest to the lowest. Columns (1), (3), and (5) refer to the median values of the three metric rates, and columns (2), (4), and (6) refer to the maximum values. We skip the minimum values because they are primarily zeros. Looking at the median values, we observe that the firm-defined conversion rate in column (5) only has three industries with a value of 0 as the median rate. It is lower than the conversion rate, with five industries with a median value of 0, and the 2s view rate, with seven industries with a median value of 0. Moreover, comparing the maximum rate, the firm-defined conversion rate only has two industries with a 0 value, and these two industries hold the most negligible proportions among all campaigns. In contrast, the conversion rate has four industries with 0 as the maximum campaign conversion rate. Our observation indicates that the conversion rate is more confined to specific industries.

Similar observations can be found in Table 3.3, a similar summary of ad performance metric rates with campaign objective breakdowns. The campaign objectives are also in order of descending proportions. Again, we observe that the firm-defined conversion rate median and maximum for all the campaign objectives are non-zero, whereas there is at least one zero value for both the 2s view and conversion rate summaries, especially for the median rate values. The metric summary with different breakdowns in the two tables shows that the firm-defined conversion rate is less subject to changes in industry or campaign objective setup. Due to its relative stability,



we decided to include it in our assessment, in addition to the most direct 2s view and conversion metrics.

We also explore our data spatially and temporally. Figure 3-3 shows the impression distribution by US state. We observe, in general, that the impression distribution is consistent with the state population distribution. For example, the four most populous states in the US are California, Texas, Florida, and New York, based on US Census Bureau data. We observe that these four states also have the four most considerable shares of ad impressions. Figure 3-2 shows the time-series pattern of the total ad expenditure, total ad impressions, and cost per impression for the three years from 2019 to 2021. There are certainly ups and downs for all three. A typical pattern is that the total ad expenditure increases at the end of each year, probably due to the winter holiday promotion season. Another general pattern is that the higher the ad expenditure, the higher the number of impressions. The exception is that after mid-2020, the cost per impression started to rise dramatically, possibly resulting from increasing competition from more ads bidding with each other.

We collected the data from ad creative summaries. Specifically, we mainly look at whether there is any speech, person, or text detected in the ad creative. Table 3.4 summarizes the proportion of the ad creatives that have each creative element detected for each industry, and industries are again ranked from the highest share to the lowest. We observe that, in general, the distribution of the presence of person and text is reasonably consistent across all industries. Except for the tech and entertainment industries, which have smaller sample sizes, the presence of people in ad creative ranges between 0.1 and 0.35, and the presence of text goes above 0.84. However, multiple industries are below 0.1 for speech, and a few are above 0.2.

### 3.3.2 Spectus Mobility Data

The aggregated stay-at-home data was provided by Spectus<sup>1</sup> through its Spectus Social Impact program Shelter-in-Place Analysis. The Shelter-in-Place Analysis represents the percentage of users staying at home in any given county and is further aggregated to the state level for this analysis. It is calculated daily by measuring how many users moved less than 330 feet from home. Figure 3-11 gives a comprehensive overview of all the stay-at-home percentage time-series patterns for all states. However, except for Washington DC, which is visibly higher than the rest of the country during the pandemic, all the other states' patterns overlay, so it is difficult to tell them apart. Figure 3-4 gives a closer look at three example states: New York (NY), California (CA), and Texas (TX). These are three very large states in the US. They are similar in that they all have very populous metro cities. We could observe that the stay-at-home trend for the three states before COVID was fairly consistent. In particular, the patterns of CA and TX almost overlap with each other before COVID occurred. Right after COVID started, we observe a much larger stay-at-home percentage difference. For example, between June 2020 and August 2020, the stay-at-home percentage kept decreasing in NY, experienced an inverse U-turn in TX, and remained relatively flat in CA. The percentages for both NY and CA are much higher than TX. Another example is that there was a spike in stay-at-home behavior in NY at the beginning of February 2021. We observe another much larger spike in TX, but two weeks later than in NY, and no spike in CA. Overall, we want to highlight that there have been much more variations in stay-at-home patterns across different states since COVID began.

---

<sup>1</sup>Spectus data was collected from de-identified mobile phone users who had opted-in to provide access to their mobility data anonymously, through a CCPA-compliant framework. Prior to sharing data with researchers, Spectus aggregated data to the county level. In order to further preserve privacy, Spectus discarded data from counties with low user counts. Researchers accessed the data under a strict agreement, which precludes attempts to disaggregate the data.

### 3.3.3 COVID Data

We collected the COVID-related data within the US over time. The main COVID variables that we have explored include the weekly new case rate, the weekly COVID-related death rate, and the cumulative number of fully vaccinated people per hundred members of the population. We found all three variables are correlated with the stay-at-home data and with each other.

We observe the correlation between the new case rate and the COVID-related death rate in the top and middle plots of Figure 3-5. There was a spike in both the new case rate and death rate in NY in April 2020. Meanwhile, we also observe that the stay-at-home spike at the start of the pandemic was the highest for NY among the three example states. TX has three peaks for both the new case rate and death rate around July 2020, January 2021, and August 2021, and the death rate peaks consistently lag behind the new case rate by a few weeks. All three states experienced a peak around January 2021 for new case and death rates, and correspondingly the stay-at-home percentage during that period also increased. Finally, in the bottom plot of Figure 3-5, we observe that the vaccination rate started to rise in March 2021 and achieved a relatively high level in June 2021. Correspondingly, we observe the stay-at-home percentage values dropping from March 2021.

## 3.4 Methodology

We first estimated the causal effects of the change in perceived pandemic severity on advertising effectiveness, measured by ad views and conversions. We use people's mobility data, and specifically the stay-at-home percentage, as a proxy for the treatment variable – pandemic severity. The immediate concern is how much we can trust the stay-at-home variable to serve as a proxy for COVID severity. Pandemic severity is undoubtedly related to the fatality rate, the number of new cases, and vaccination. The fatality rate indicates the lives lost; the number of new cases indicates the potential future lives lost and how wide and fast the spread of the virus is. The cumulative vaccination rate indicates the level of enhanced immunization against the

virus or severe symptoms, and that could reduce both new cases and fatalities. All three variables are critical in defining people’s perception of the pandemic’s severity.

The previous data section finds that the three variables all correlate with the stay-at-home percentage to a certain extent. We run a simple fixed-effects linear regression with these three COVID-related statistics as the independent variables and the state and week fixed effects. We use the stay-at-home percentage as the dependent variable and limit the data timeframe from mid-March 2020 to the end of 2021. The model assesses how correlated the COVID data are with the chosen proxy (stay-at-home data). The results are shown in Table 3.5. As expected, all three variables have highly significant coefficients. The  $R^2$  value goes as high as 0.97. This result confirms that the COVID-related information, including the death rate, the new case rate, and the vaccination rate alone, plus the location and time fixed effects, already have an extremely high predictive power of people’s stay-at-home behavioral tendency. The high goodness-of-fit reassures us about using the stay-at-home percentage as a proxy for perceived pandemic severity.

Our primary method employs difference-in-differences with a fixed-effects regression. Figure 3-14 illustrates an example of the aggregated 2s view outcomes over time. Before COVID-19 started, the metrics of the two example states – CA and NY – overlapped. After COVID, especially after the initial outbreak, we observe there are more variations between the two states. The regression includes time fixed effects by month of the year  $\theta_t$ , location fixed effects by state  $\eta_s$ , and the interaction of a dummy for the start of the COVID pandemic  $P_t$ , indicating the start of the treatment period, with a variable that measures the treatment intensity  $D_{st}$ : COVID severity. Specifically, COVID severity is represented by the percentage of people who stay at home the entire day in a given region. This variable is one of the possible proxies we have chosen. The follow-up studies also explore using the COVID death rate and other mobility data as proxies for actual COVID severity. The outcome  $Y_{ist}$  is a binary variable indicating either the ad view action or the conversion action following an impression. For example, if we consider ad conversion as the outcome variable, we use 1 to indicate conversion success for one case and 0 for no conversion for this case.

Below is the model specification:

$$Y_{ist} = \alpha_i + \theta_t + \eta_s + \delta(P_t \times D_{st}) + X_{ist}\beta + \epsilon_{ist} \quad (3.1)$$

On top of the basic model, we also include several more covariates  $X$  and fixed effects  $\alpha_i$  to control for campaign- and ad-level differences. We include the weekly cost per impression (CPM) and the weekly ad spend (log spend) to control for shifts in ad competitiveness, which might affect targeting effectiveness. We include the duration an ad has been running on the platform (campaign duration) to control for changes in ad effectiveness, because the longer the ad has been running, the more likely it is that people will see the same ad multiple times. People who are repeatedly exposed to the same ad might have different reactions compared to those who see the ad for the first time. We include indicators for the basic creative elements: whether the ad has the presence of people, text, speech, and logo detected. Apart from time and location fixed effects, we also included campaign ID, campaign objective, and industry category fixed effects. Finally, we used campaign and time two-way cluster-robust standard errors.

## 3.5 Results & Discussion

This section presents the main model results from using the causal method explained in the previous section. Besides the main effect, we further incorporate some interaction terms into the model to investigate the underlying mechanisms.

Tables 3.6, 3.7 and 3.8 show the results of the fixed-effects models for 2s ad views, conversions and firm-defined conversions, respectively. In Table 3.7, perceived COVID severity, as represented by the stay-at-home measure, exerts a significantly positive effect on conversions. The coefficient 0.000289 of COVID start  $\times$  stay-at-home percentage in column (1) can be interpreted as meaning that a ten percentage-point increase in the stay-at-home population will lead to a 0.000029 increase in

the conversion probability. Similarly, the same term in Table 3.8 has a significantly positive value of 0.0143, meaning that a ten percentage-point increase in the stay-at-home population of the given state and time will result in a 0.00143 increase in the firm-defined conversion probability. Meanwhile, the same term in Table 3.6 column (1) has an insignificant coefficient, with high cluster-robust standard errors.

So far, we find that the severity of the COVID pandemic has a significantly positive effect on conversions and firm-defined conversions but no significant effect on ad views. A natural follow-up question is what caused these effects. We propose there are two possible underlying factors: the substitution effect and psychological factors. The substitution effect refers to the case that people simply replace their offline shopping with online shopping, due to the shelter-in-place practice. Psychological factors refer to negative feelings that we observe in the prior literature, such as loneliness or anxiety, which lead to unusual purchases. To investigate the underlying mechanism, we interact the treatment term with the cumulative vaccination rate and the presence of people in the ad creative in columns (2) and (3), respectively. The vaccination rate is represented as the number of fully vaccinated people per 100 people. As shown in Figure 3-13, the vaccine became available to the public around April 2021, and the vaccination rate in each state ramped up at a different speed. Research has also suggested that the availability of an effective COVID vaccine helped to boost people's confidence in the economy [18].

In column (2) of Tables 3.6 - 3.8, the difference-in-differences treatment term interacts with the vaccination rate, and we find the latter significantly and negatively moderates the effect of increasing perceived pandemic severity on ad conversions and views. However, we do not observe any moderating effect on firm-defined conversions. Specifically, when the fully vaccinated population increases by 10 per 100, the effect of 10 percentage points more people practicing stay-at-home on two-second view probability decreases by 0.0037. The negative moderating effect means that when the COVID situation becomes more serious, ad views and conversions increase, but to a significantly lesser extent when there is a higher vaccination rate than when there is lower vaccination coverage. This finding is consistent with the existing literature that

when people’s confidence increases with the increased vaccination rate, people feel the pandemic situation is more under control and therefore reduce unusual purchases due to less fear and panic. Moreover, the possible substitution effect occurs when people reduce their current online shopping in anticipation of being able to go out for offline shopping in the near future, as the vaccination rate goes up.

In Table 3.6 column (3), we interact the treatment term with the indicator of whether there are any people present in the ad creative. We find the presence of people in the ad creative significantly moderates the relationship between pandemic stay-at-home behavior and ad views in a positive direction. When there are people in the ad creative, a ten percentage-point increase in people staying at home during the COVID pandemic increases the two-second view probability by 0.070, compared to ad creative without people. This again proves the existence of psychological factors. When people are practicing more self-isolation due to COVID concerns, they tend to feel lonely and thus are more willing to see other people, even if it is in video ads. However, we find no significant moderating effect of people in ad creative on conversions or firm-defined conversions, as shown in Tables 3.7 and 3.8 column (3). This is consistent with our findings in the previous two chapters that visual manipulation tends to have a larger impact on upper-funnel view behavior compared to lower-funnel conversion behavior.

Overall, the interaction terms in column (2) of Tables 3.6 & 3.7 and column (3) of Table 3.6 prove there are psychological effects. Specifically, people’s social isolation tends to enforce a sense of loneliness, and hence they are more attracted by the appearance of people in ads. Correspondingly, we find the presence of people in ad creative does have a significantly positive moderating effect on views. Moreover, vaccination boosts confidence that the economy and social life will recover. Its moderating effect is also proof that the effect of varying COVID severity on advertising conversion is not just a substitution effect between offline and digital purchase options, but also an actual increase in conversion that contributes to an increase in profitability.

### 3.5.1 Further Discussion

In this subsection, we take a closer look at the main effects. The pandemic lasted for a relatively long period, and we want to gain a better understanding of the effect time limit. Tables 3.9 - 3.11 address the question of when the effect was stronger or weaker. In our new model, instead of aggregating the effects of the entire pandemic duration from March 2020 to December 2021, we divide the pandemic period from March 2020 to December 2021 into four segments: 1) from the date that the US federal government declared the COVID-19 pandemic as a national emergency on March 13, 2020, to June 30, 2020, 2) from July 1, 2020, to December 31, 2020, 3) from January 1, 2021, to June 30, 2021, and 4) from July 1, 2021, to December 31, 2021. Due to space constraints, we cannot present all the model variables in the three tables, but we show the key treatment variables that reflect the effect sizes.

In Table 3.9 column (1), we observe that there was a significantly positive effect of increasing COVID severity on 2s views immediately after the COVID-19 pandemic started. We missed this effect when looking at the aggregated effects in Table 3.6. However, the effect quickly diminishes after the first period and becomes more and more negative going forward. Similarly, in Table 3.10 column (1), we observe that the effect of COVID on ad conversions is only significant in the first period, i.e., the first three months of the pandemic. As time goes by, although the effect remains positive, it is no longer significant. In contrast to these two ad outcomes, in Table 3.11 column (1), we observe the effects to be significantly positive in the first and fourth periods of the pandemic. Moreover, it is also weakly significant in the second period. The result indicates that fluctuations in pandemic severity had a longer-lasting impact on firm-defined conversions.

In Table 3.9 column (2), we observe the positive moderating effect of the presence of people in ad creative persisted in periods 1 and 3. The magnitude became consistently smaller in periods 2 and 4, likely due to seasonality. In Table 3.10 column (2), we observe a significantly negative moderating effect of the presence of people in ad creative on the relationship between COVID severity and conversions in the



second pandemic period. In Table 3.11 column (2), we observe no significant effects, consistent with findings from the main result in Table 3.8.

### 3.5.2 Robustness Checks

This subsection further discusses and validates the choice of treatment variable (i.e., perceived pandemic severity) and using the stay-at-home percentage as a proxy to quantify pandemic severity. The main concern about the choice of independent variable is whether the stay-at-home mobility data serves as an appropriate proxy for pandemic severity. In this subsection, we compare two alternative treatment variables to address the concern.

People might doubt whether the observed variations in stay-at-home percentage values across different regions are predominantly driven by the work-from-home and return-to-office policies that the firms in that region adopted. Such policies might vary a lot by region and time and are not necessarily entirely based on actual pandemic severity. For example, with a high concentration of technology firms, the San Francisco Bay Area can adapt to working from home much better than an area that thrives on tourism. Therefore, we used the mobility index for malls as an alternative proxy for the estimation because people’s visits to malls are generally less subject to whatever work policies firms adopt. Rather, it is more related to how comfortable people feel about being in indoor spaces with strangers and how serious people consider the pandemic situation to be. The Cuebiq Visit Index (CVI) is a metric provided by Spectus and is aggregated at the Designated Market Area region level and also by vertical, e.g., retail. It is calculated as the aggregated number of visits to locations in that vertical weighted by population, divided by Spectus’s active user base times 1000. The mall CVI only includes measured visits to mall locations. Columns (1) and (3) in Table 3.12 show the results of using Spectus’s mobility index for malls as the treatment variable. The model results show that the decrease in mall mobility index (i.e., a proxy for an increase in COVID severity) exerts a significantly positive impact on ad conversions and a weakly significantly positive impact on firm-defined ad conversions. The interpreted findings are consistent with the main results.

The consistency between the malls' mobility patterns and the stay-at-home percentage patterns over time further supports the consistent findings we have when using them as the independent variables. Figures 3-6 and 3-7 show the time-series patterns of the three independent variables in Massachusetts (MA) and CA, respectively. The top and bottom plots correspond to the stay-at-home percentage and the mall visit index. We reverse the mall visit index pattern by multiplying by -1 when making the plot so that the patterns are more easily comparable. The higher the stay-at-home percentage, the fewer visits to malls, and hence the higher the negative mall visit index. The stay-at-home and mall visit index plots in both example states show that the patterns are fairly consistent.

Another concern about the dependent variable lies in the actual versus perceived severity of the COVID pandemic situation. Figure 3-7 illustrates that how people in CA felt about the COVID situation could differ from the actual severity of COVID. Both the stay-at-home percentage (top) and mall visit index (middle) plots could serve as proxies for how comfortable people were with going out and thus how they thought about the risk/severity of the pandemic after the onset of COVID. For example, we observe a massive surge around March 2020, when the COVID pandemic was just beginning. People felt very anxious about COVID-19 because the entire world still knew very little about the virus. There was another peak at the start of 2021 due to the Delta variant, but it was much smaller than the first wave. Starting in the spring of 2021, the stay-at-home percentage and mall visit index become much smaller as the vaccine was taken up widely. Meanwhile, the COVID fatality rate could be seen as a more accurate representation of pandemic severity, but there was only a minor surge in the death rate in March 2020, whereas the biggest spike occurred during the Delta wave. Overall, the COVID-related death plot (bottom) does not align very well with the two plots above it. Moreover, the death rate tends to be a lagging indicator of the COVID case situation, based on what we observe in the Data section of Figure 3-5 and the related literature [65].

Due to the disparity between the actual severity and perceived severity patterns, people might also doubt which exerts a greater effect on people's behavior. Columns

(2) and (4) in Table 3.12 show the results of using the COVID death number per million population as the treatment variable. It has no significant effect on ad conversions or firm-defined conversions. In contrast, in column (1) of Tables 3.7 and 3.8, we already see that the perceived severity (i.e., the stay-at-home percentage) exerts a significantly positive effect on the two ad conversions. The model results show that people’s perceived severity rather than the actual severity plays a dominant role in affecting advertising effectiveness. The results not only back up our choice of the independent variable in the main results but are also consistent with our sense that people’s behavior might be more influenced by what they believe in, not the truth.

In addition, we also perform robustness checks by using more granular location breakdowns and adding more image-/video-related covariates in the models. Table 3.15 presents the model results using DMAs (Designated Marketing Areas) as location breakdowns instead of states. DMAs are media markets within the United States, typically defined based on metropolitan areas with surrounding suburbs. There are 210 DMAs across the entire United States. We observe that the effect on 2s views is insignificant, and the effects on two types of conversions are significantly positive. Not only is the significance consistent with our state-level models but the magnitude is also very close, i.e., 0.0148 vs. 0.0143 for firm-defined conversions and 0.00025 vs. 0.00029 for conversions. Table 3.16 presents the model results when adding additional creative element covariates. These covariates are the features that we use for feature importance ranking in Chapter 1. The results are again very consistent with the original model results in terms of significance and magnitude.

### **3.5.3 Generalizability – Alternative Social Media Platforms**

This subsection presents the fixed-effects model results using Facebook ad data. Again, the data comes from the collaborating ad creative and analysis firm, which sampled the Facebook advertising performance data and corresponding tag data from its database. Overall, the ad conversion results are consistent with the model findings from the Pinterest data, indicating that the observations are potentially generalizable.

We recognize certain differences and limitations regarding the Facebook data com-

pared with the Pinterest data. First, the Facebook ad performance data has much coarser location breakdowns. Specifically, the collaborating firm can only provide the ad metrics with country breakdowns. There are possibly many variations within a country regarding COVID severity, people’s responses to the pandemic, and the use of social media. Second, the stay-at-home data source changes. The Spectus mobility data is only available for data within the US, so we switched to Facebook mobility data to gain the stay-at-home data for all countries around the globe. This change directly leads to the third limitation: the data used here only ranges from March 2020 to December 2021 because Facebook mobility data starts from March 2020. Since the data no longer starts in January 2019, we do not have 2019 data as a baseline in our model. As a result, the data we use ranges from March 2020 to December 2021. The ad performance data includes 12.5k unique campaigns and 84.6k unique media assets and covers 91 countries. The ad impressions are disproportionately distributed among countries, with the US taking 41.7% and Brazil taking 32.4%. Mexico has the 3rd largest share, at just 0.03%. These limitations could make both the results and interpretation slightly different from the main model, and thus we decided not to include the Facebook data findings in our primary findings. Instead, we discuss the results from the Facebook data in this generalizability section.

Table 3.13 presents the model results for the ad conversion outcome. We define either a purchase or an app install as a conversion action. Column (1) shows the results of the basic model with time, location, and ad-setup fixed effects, with campaign- and time-level cluster-robust standard errors. We observe that the stay-at-home coefficient is 0.134 and weakly significant. This effect means that when the stay-at-home ratio increases by 0.10, the conversion probability increases by 0.0134. As shown in column (2), the full vaccination rate has a weakly negative moderating effect. Lastly, the presence of people in ads does not moderate the relationship between the stay-at-home ratio and ad conversions. All three observations are consistent with the main results of the conversion outcome.

However, the results for ad views are not generalizable. In Table 3.17, we observe that an increase in stay-at-home behavior has a significantly negative effect on the

three-second view probability and no moderating effects of either the full vaccination rate or the presence of people in the ad creative on the 3s-view metric. The possible reasons include differences in how Facebook and Pinterest define the metrics, data granularity, the data timeframe used for modeling, the lack of a 2019 baseline for Facebook data, the interface, the types of ad creative on the two platforms, etc. Due to the considerable differences in our data use and modeling, we cannot tell what exactly causes the differences between the Facebook and Pinterest results. Meanwhile, the fact that the effect direction on the conversion metric remains similar, even with so many differences, confirms the generalizability of conversion results across different platforms and countries.

## 3.6 Heterogeneous Results

In this section, we describe some relevant heterogeneous results. We first identify and explore the ads of some industries that have prominent effects in either direction. The second part discusses how the heterogeneous effects vary by campaign objective.

### 3.6.1 Effects by Industry

In this subsection, we explore how the ad performance of different industries is affected to varying extents. In Figure 3-8, we list the effects of industries that are significantly different from zero. The top plot shows the significant effects on two-second views, the middle plot shows conversions, and the bottom plot shows firm-defined conversions.

In the top plot, we observe that increasing perceived COVID severity has significant adverse effects on views of health & wellness ads and home improvement ads. The negative view tendency for health & wellness ads might be related to information avoidance. At the time of the pandemic crisis, social media exposure to healthcare and medical information tended to bring about information overload, which led to information anxiety and eventually information avoidance [62, 24, 43]. In the middle plot, we observe that the effects of increasing pandemic severity on conversion are positive for the retail industry but negative for the travel industry. The results are

relatively reasonable, given that people tended to stockpile during the pandemic, especially when it became more severe. Concurrently, people also avoided any travel because it involved more risk of being exposed to the virus. In the last plot, the effect of pandemic severity on firm-defined ad conversions is again positive for the retail industry. Moreover, the effect is also positive for the health and wellness industry.

In Figure 3-9, we present how the moderating effect of the presence of people in ad creative varies by industry. The top plot shows the effects on 2s views. We observe positive effects for ads in the automotive and retail categories. Meanwhile, we also observe negative effects for ads in the alcohol, government or nonprofit, health and wellness, and telecommunications industries. The bottom plot shows the person moderating effects on the relationship between the stay-at-home percentage and ad conversions. We observe a positive effect for travel ads but negative effects for retail and health ads. We observe the opposite effect direction for retail ads in terms of the moderating effects on 2s views and conversions. Meanwhile, the negative moderating effects of the presence of people in health and wellness ads are consistent in both ad performance measures. We do not present the effects on firm-defined conversions because we do not observe significant moderating effects.

### **3.6.2 Effects by Campaign Objective**

In this subsection, we would like to learn whether the varying campaign objectives could cause a difference in the effects of pandemic severity. We do not find much difference among the campaign objectives for the 2s view and conversion outcomes. None of these effects are significantly different from zero. However, we do observe a significant difference in the firm-defined conversion outcome, as shown in Figure 3-10. The change in pandemic severity had significantly positive effects on ads with the web conversion campaign objective. The effect was also significantly higher than all other campaign objective effects.

## 3.7 Conclusion

Overall, our research has explored whether and how changes in perceived pandemic severity affect advertising effectiveness on social media platforms. We find positive effects on ad conversions when the pandemic is perceived to be more severe, as represented by people’s stay-at-home practices. Moreover, at the beginning of the pandemic, there were significant positive effects on both the view and conversion metrics when the pandemic became more severe. As time went on, the effects generally tended to diminish, except for firm-defined conversions. We then interacted COVID severity with other variables, including the vaccination rate and the presence of people in ad creative. We find that the vaccination rate negatively moderates the effect of stay-at-home behavior on ad views and conversions, and the presence of people positively moderates the effect of stay-at-home behavior on ad views. We further performed some robustness checks to confirm the choice of our treatment variable and the generalizability of the conversion outcome findings on another social media platform. Lastly, we looked at the heterogeneous effects of different campaign objectives and industry categories.

Our research has important managerial implications. As the federal government has renewed the national emergency declaration for the COVID-19 pandemic, the disease and its variants are very likely to be around for a long time. The virus has been through ups and downs in severity, in terms of spread and associated deaths. Based on our research findings, firms can still proactively boost their advertising effectiveness in the face of the COVID-19 outbreak and future crises. For example, when more people stay at home due to virus concerns, it might be wise for firms to increase online advertising expenditure, as people tend to make more online purchases. Moreover, firms could also vary their ad creative design to better cater to people’s psychological needs based on their advertising goals. For example, ads in retail categories might consider having people present in the ad to boost views.

Admittedly, our research is not without limitations. First of all, it would be beneficial to better understand the ad metrics, especially the firm-defined conversion

metric. For now, we do not have any access to how each firm defines its conversion action for each ad. Moreover, our data is only limited to what our collaborating firm has from the Pinterest platform. Further research could incorporate more data from different platforms and industries to gain a more generalizable view of the effects.



## 3.8 Appendix

### 3.8.1 Figures

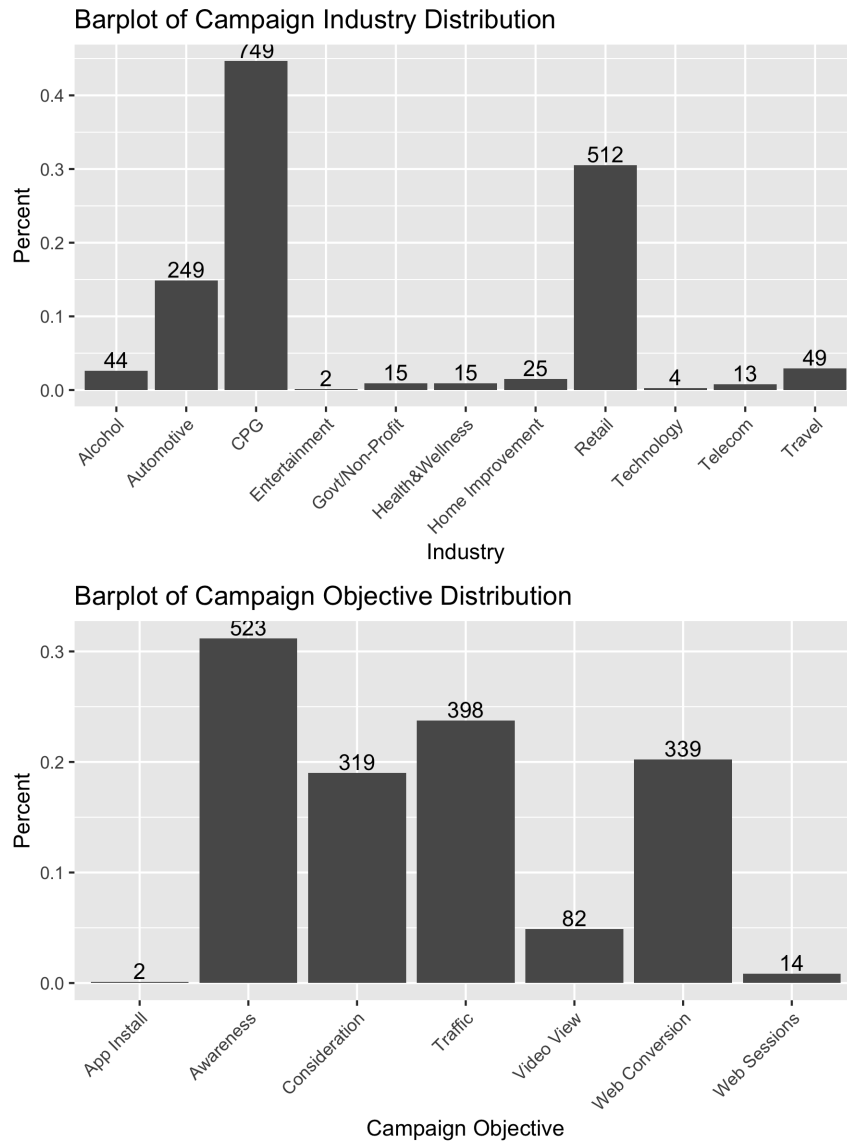


Figure 3-1: Summary of the campaign industry and objective percent

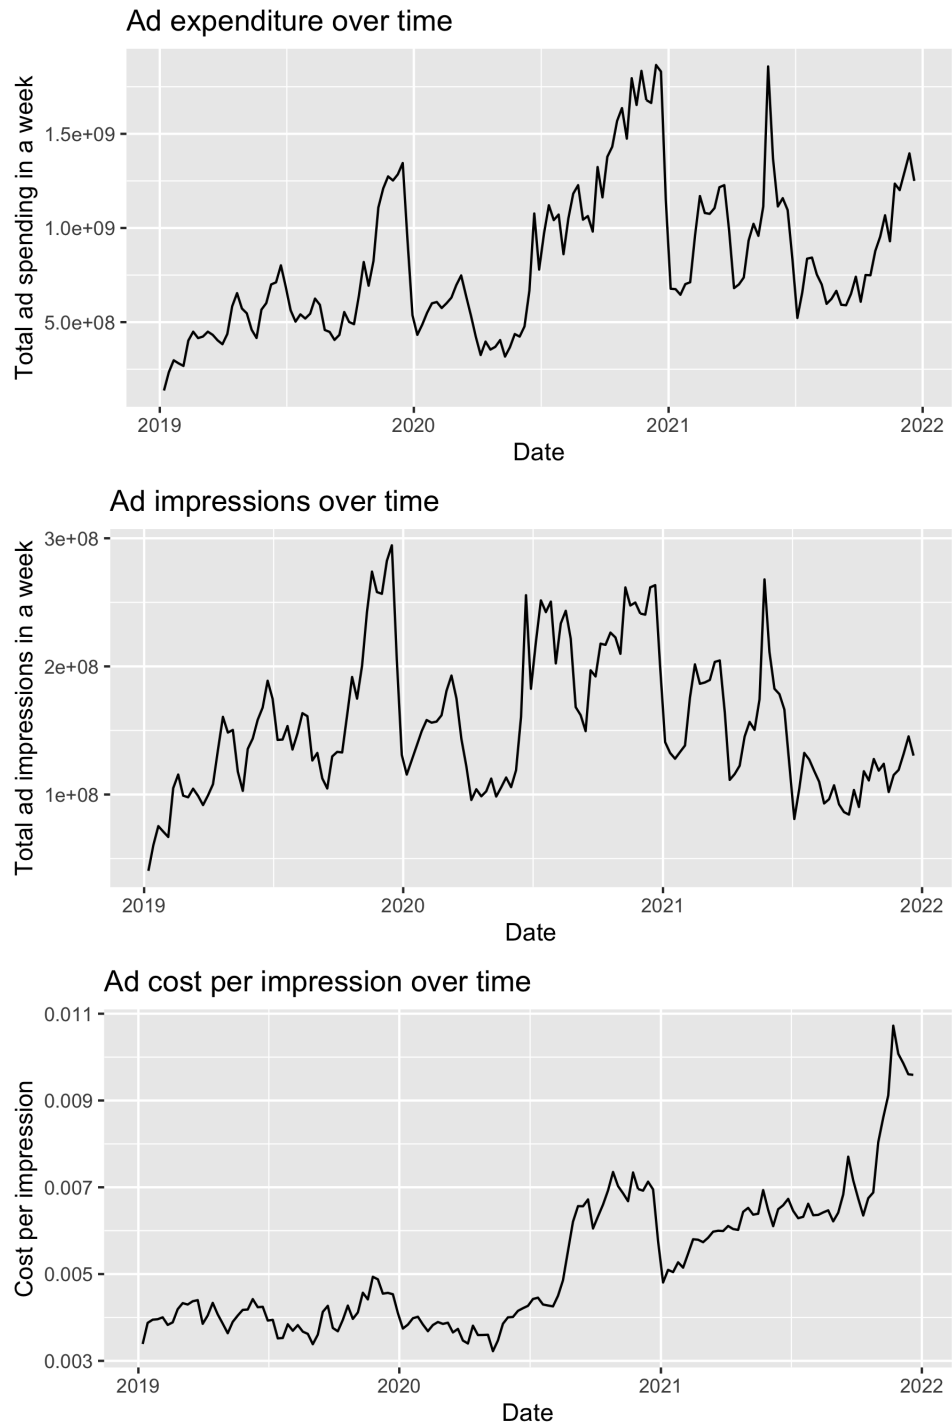


Figure 3-2: Time series plots of ad expenditure and impressions




Figure 3-3: Region distribution of Pinterest ad impressions




Figure 3-4: Stay-at-home time series data of example states




Figure 3-5: COVID-related data over time of NY, CA and TX




Figure 3-6: Correlation between stay-at-home and death rate in MA




Figure 3-7: Correlation between stay-at-home and death rate in CA

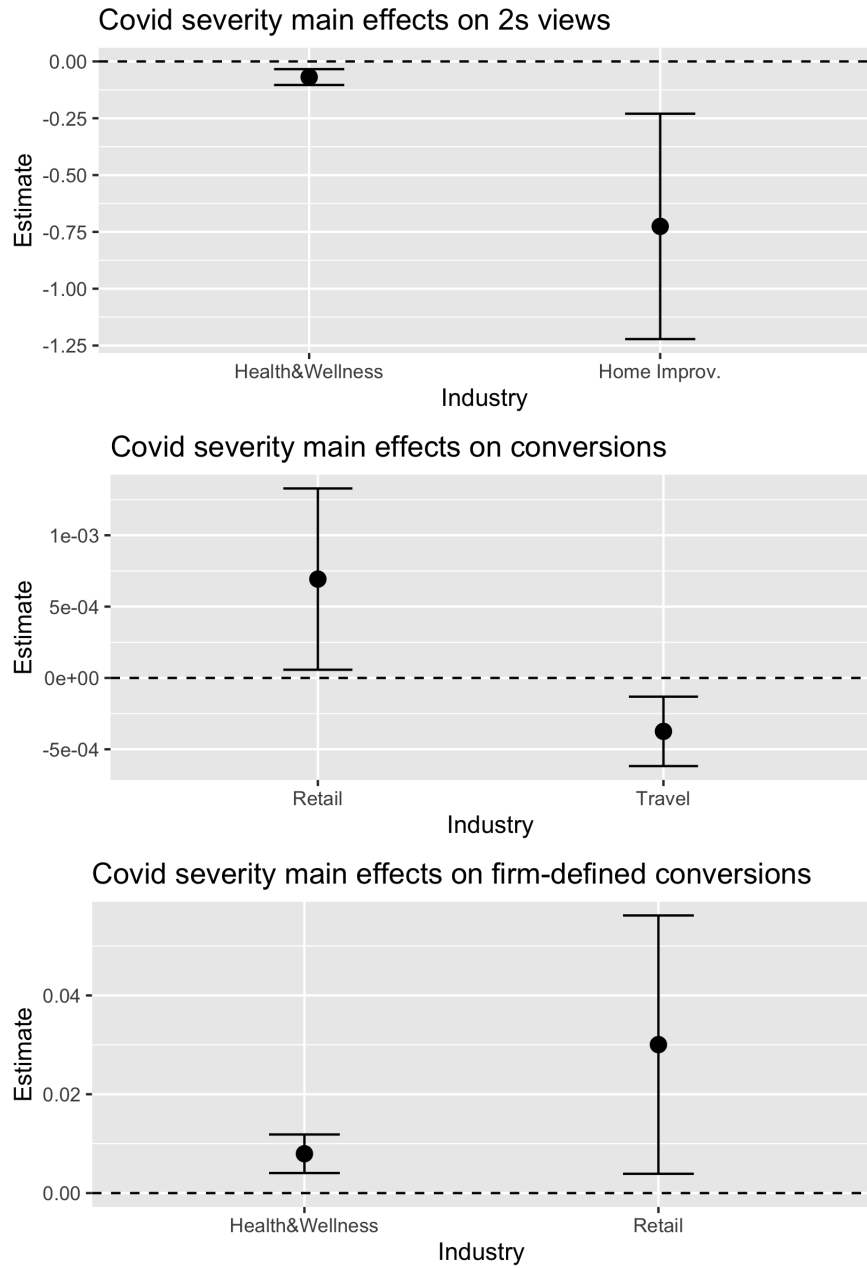


Figure 3-8: Effects of stay-at-home on ad effectiveness by industry categories



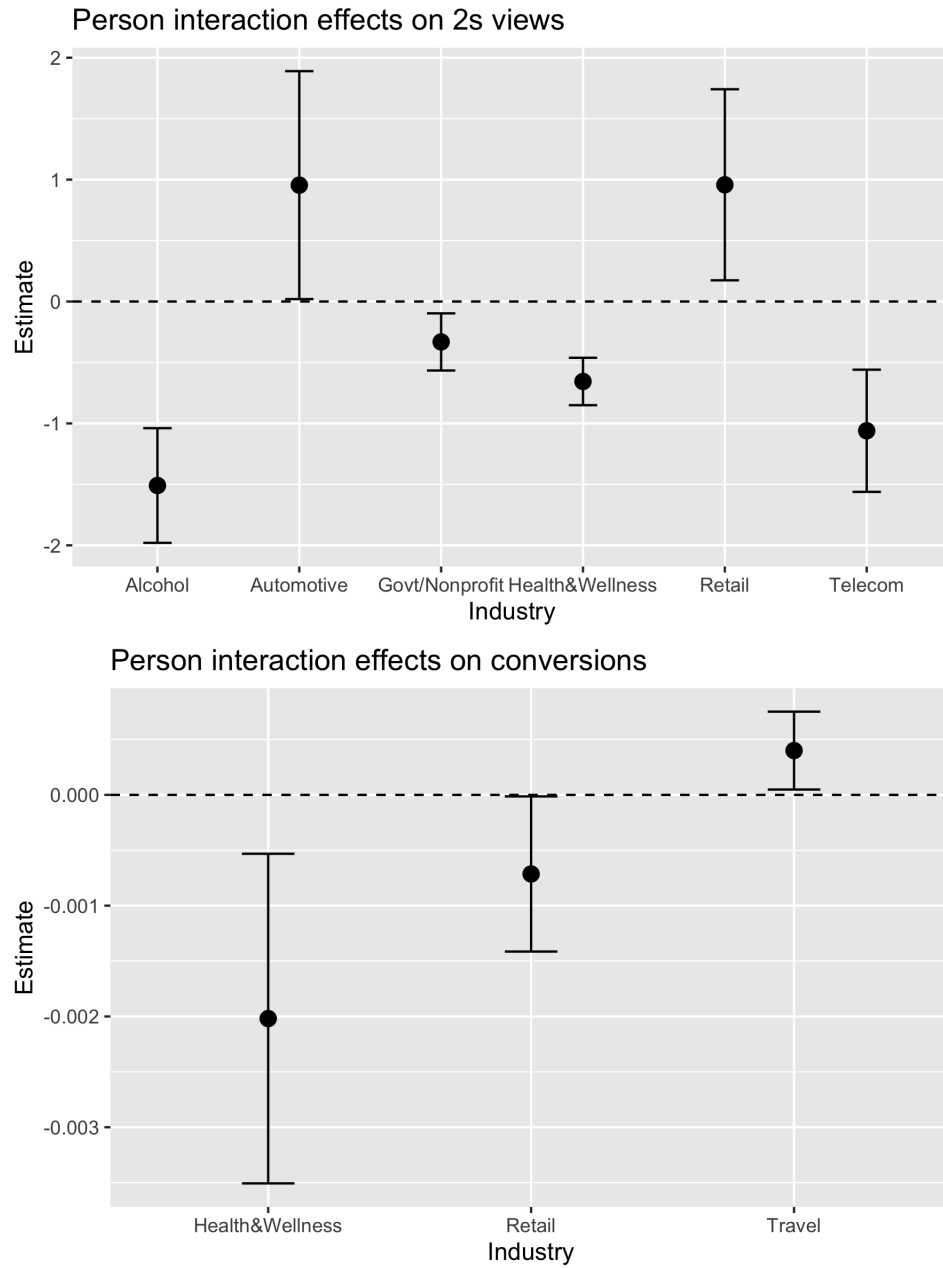


Figure 3-9: Moderating effects of person creative by industry categories

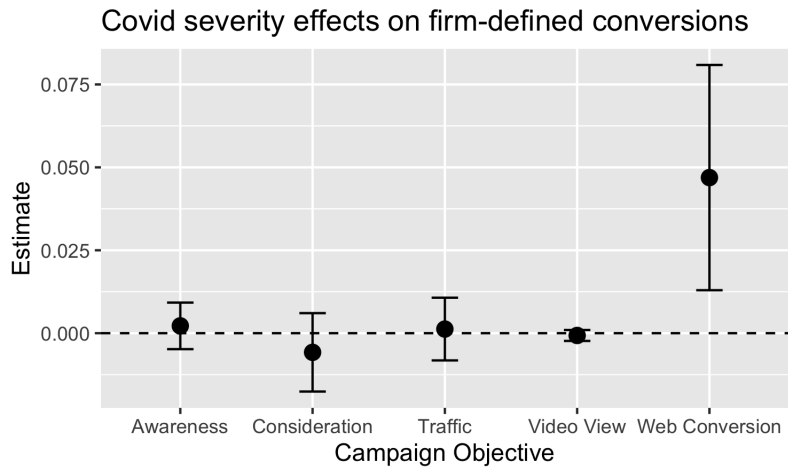


Figure 3-10: Effects of stay-at-home on ad effectiveness by industry categories




Figure 3-11: Stay-at-home time series data of all states

Covid-related deaths over time for each state




Figure 3-12: COVID-19 fatality time series data of all states




Figure 3-13: Vaccination rate time series data of all states




Figure 3-14: Aggregated 2s view rate time series data of example states – parallel trend

### 3.8.2 Tables

Table 3.1: Summary of campaign metrics

	Mean	Min.	1st Qu.	Median	3rd Qu.	Max.
Impressions	14.4MM	1	1.2MM	4.8MM	15.4MM	366.6MM
Spend	75,897	0	7,582	29,117	88,080	2,015,398
2s View Rate	0.1787	0.0000	0.0000	0.0000	0.3735	0.7593
Conversion Rate	0.00007	0.00000	0.00000	0.00000	0.00004	0.00191
Conv.(FD) Rate	0.0053	0.0000	0.0001	0.0007	0.0036	0.2698

Table 3.2: Table of ad campaign performance metric summary by industry

Industry	2s view		Conversion		Conversion (FD)	
	Median	Max	Median	Max	Median	Max
	(1)	(2)	(3)	(4)	(5)	(6)
CPG	0.00	0.72	0.000000	0.00191	0.00013	0.1389
Retail	0.00	0.76	0.000011	0.00131	0.00249	0.2698
Automotive	0.00	0.73	0.000039	0.00170	0.00064	0.0064
Travel	0.00	0.60	0.000060	0.00056	0.00657	0.0526
Alcohol	0.38	0.69	0.000000	0.00000	0.00015	0.0044
Home Improvement	0.00	0.62	0.000009	0.00028	0.00121	0.0223
Govt./Non Profit	0.00	0.62	0.000000	0.00000	0.00000	0.0049
Health & Wellness	0.00	0.65	0.000073	0.00050	0.00012	0.0086
Telecom	0.15	0.69	0.000015	0.00005	0.00011	0.0126
Technology	0.48	0.51	0.000000	0.00000	0.00000	0.0000
Entertainment	0.54	0.54	0.000000	0.00000	0.00000	0.0000



Table 3.3: Table of ad campaign performance metric summary by campaign objective

Objective	2s view		Conversion		Conversion (FD)	
	Median	Max	Median	Max	Median	Max
	(1)	(2)	(3)	(4)	(5)	(6)
Awareness	0.25	0.62	0.000000	0.00025	0.00015	0.0369
Traffic	0.00	0.69	0.000001	0.00095	0.00022	0.1273
Web Conversion	0.00	0.76	0.000094	0.00170	0.00199	0.2698
Consideration	0.00	0.72	0.000008	0.00131	0.00179	0.1168
Video View	0.43	0.73	0.000001	0.00015	0.00046	0.0237
Web Sessions	0.07	0.72	0.000061	0.00191	0.00462	0.1389
App Install	0.00	0.00	0.000000	0.00000	0.00004	0.0001

Table 3.4: Table of main ad creative presence summary by industry

Industry	Person	Text	Speech
CPG	0.24	0.96	0.10
Retail	0.13	0.84	0.06
Automotive	0.34	0.95	0.29
Travel	0.11	0.98	0.06
Alcohol	0.33	0.99	0.24
Home Improvement	0.20	0.94	0.02
Government/Non Profit	0.31	0.92	0.08
Health & Wellness	0.22	1.00	0.08
Telecom	0.18	0.95	0.09
Technology	1.00	1.00	0.58
Entertainment	1.00	1.00	1.00

Table 3.5: Fixed effect model of the COVID indicators on stay-at-home

	<i>Dependent variable:</i>
	stay-at-home
Weekly death rate	243.66*** (12.15)
Weekly new case rate	1.65*** (0.22)
Accumulative vaccination	-0.18*** (0.01)
Observations	4,743
R <sup>2</sup>	0.97
Adjusted R <sup>2</sup>	0.97
Residual Std. Error	0.01 (df = 4597)
<i>Note:</i>	*p<0.05; **p<0.01; ***p<0.001

Table 3.6: Main results for 2s view outcome

	<i>Dependent variable:</i>		
	Two-Second Views		
	(1)	(2)	(3)
stay-at-home	-0.0543 (0.0499)	-0.0061 (0.0428)	0.0736 (0.0623)
COVID start	0.0198 (0.0153)	0.0113 (0.0135)	0.0452* (0.0186)
full vacc. rate		0.0023** (0.0007)	
log spend	0.0043** (0.0015)	0.0043** (0.0015)	0.0041** (0.0015)
campaign duration	-0.0010** (0.0004)	-0.0012** (0.0004)	-0.0010** (0.0004)
CPM	7.1703** (2.5548)	2.3919 (2.4067)	7.4102** (2.5734)
person	0.3147*** (0.0219)	0.3141*** (0.0218)	0.3979*** (0.0723)
logo	0.0066 (0.0192)	0.0107 (0.0190)	0.0057 (0.0192)
text	-0.0143 (0.0276)	-0.0146 (0.0273)	-0.0121 (0.0274)
speech	0.0650*** (0.0172)	0.0649*** (0.0171)	0.0629*** (0.0174)
stay-at-home×COVID start	0.0238 (0.0428)	0.0442 (0.0391)	-0.1174 (0.0625)
stay-at-home×full vacc. rate		-0.0048** (0.0016)	
stay-at-home×person			-0.6372* (0.2947)
COVID start×person			-0.0804 (0.0819)
stay-at-home×COVID start×person			0.7016* (0.3373)
Observations	8,695,668	8,695,668	8,695,668
R <sup>2</sup>	0.3531	0.3535	0.3538
Adjusted R <sup>2</sup>	0.3530	0.3533	0.3537

*Note:* The fixed effects include month, state, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 3.7: Main results for conversion outcome

	<i>Dependent variable:</i>		
	Conversions		
	(1)	(2)	(3)
stay-at-home	0.000067 (0.000107)	0.000055 (0.000106)	0.000045 (0.000108)
COVID start	-0.000092 (0.000066)	-0.000105 (0.000070)	-0.000103 (0.000073)
full vacc. rate		0.000001* (0.000001)	
log spend	0.000006*** (0.000001)	0.000006*** (0.000001)	0.000006*** (0.000001)
campaign duration	-0.000001 (0.0000005)	-0.000001 (0.0000005)	-0.000001 (0.0000005)
CPM	-0.011235 (0.009191)	-0.010319 (0.008876)	-0.011128 (0.009105)
person	-0.000001 (0.000006)	-0.000002 (0.000006)	-0.000020 (0.000022)
logo	0.000002 (0.000007)	0.000003 (0.000006)	0.000002 (0.000007)
text	-0.000004 (0.000006)	-0.000003 (0.000006)	-0.000003 (0.000006)
speech	-0.000014** (0.000005)	-0.000014** (0.000005)	-0.000014** (0.000005)
stay-at-home×COVID start	0.000289* (0.000146)	0.000329* (0.000157)	0.000334 (0.000177)
stay-at-home×full vacc. rate		-0.000007** (0.000002)	
stay-at-home×person			0.000095 (0.000088)
COVID start×person			0.000046 (0.000041)
stay-at-home×COVID start×person			-0.000212 (0.000190)
Observations	8,695,668	8,695,668	8,695,668
R <sup>2</sup>	0.000424	0.000424	0.000424
Adjusted R <sup>2</sup>	0.000231	0.000231	0.000231

*Note:* The fixed effects include month, state, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 3.8: Main results for firm-defined conversion outcome

	<i>Dependent variable:</i>		
	Firm-Defined Conversions		
	(1)	(2)	(3)
stay-at-home	-0.00735 (0.00671)	-0.00931 (0.00737)	-0.00843 (0.00729)
COVID start	-0.00126 (0.00131)	-0.00081 (0.00163)	-0.00157 (0.00136)
full vacc. rate		-0.00008 (0.00007)	
log spend	0.00017** (0.00006)	0.00017** (0.00005)	0.00017** (0.00006)
campaign duration	-0.000004 (0.00001)		-0.000004 (0.00001)
CPM	-0.10860 (0.21322)		-0.10848 (0.21360)
person	0.00009 (0.00022)	0.00011 (0.00022)	-0.00063 (0.00084)
logo	0.00069 (0.00047)	0.00054 (0.00038)	0.00070 (0.00047)
text	-0.00039 (0.00033)	-0.00037 (0.00036)	-0.00040 (0.00033)
speech	-0.00012 (0.00018)	-0.00012 (0.00020)	-0.00011 (0.00019)
stay-at-home×COVID start	0.01433* (0.00685)	0.01340 (0.00710)	0.01585* (0.00758)
stay-at-home×full vacc. rate		0.00017 (0.00021)	
stay-at-home×person			0.00515 (0.00389)
COVID start×person			0.00110 (0.00100)
stay-at-home×COVID start×person			-0.00734 (0.00487)
Observations	8,695,668	8,695,668	8,695,668
R <sup>2</sup>	0.03235	0.03237	0.03235
Adjusted R <sup>2</sup>	0.03217	0.03218	0.03217

*Note:* The fixed effects include month, state, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 3.9: Table of 2s views – segmenting the pandemic duration

	<i>Dependent variable:</i>	
	Two-Second Views	
	(1)	(2)
stay-at-home×pandemic duration 1	0.1067** (0.0407)	−0.0619 (0.0558)
stay-at-home×pandemic duration 2	0.0209 (0.0370)	−0.0616 (0.0621)
stay-at-home×pandemic duration 3	−0.0781 (0.0489)	−0.3035*** (0.0790)
stay-at-home×pandemic duration 4	−0.0394 (0.0472)	−0.0790 (0.0601)
stay-at-home×pandemic duration 1×person		0.9538** (0.3165)
stay-at-home×pandemic duration 2×person		0.4486 (0.3186)
stay-at-home×pandemic duration 3×person		1.0007** (0.3361)
stay-at-home×pandemic duration 4×person		0.4499 (0.3308)
Observations	8,695,668	8,695,668
R <sup>2</sup>	0.3536	0.3547
Adjusted R <sup>2</sup>	0.3534	0.3545

*Note:* The fixed effects include month, state, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 3.10: Table of conversions – segmenting the pandemic duration

	<i>Dependent variable:</i>	
	Conversions	
	(1)	(2)
stay-at-home×pandemic duration 1	0.000467*	0.000492
	(0.000228)	(0.000252)
stay-at-home×pandemic duration 2	0.000154	0.000205
	(0.000099)	(0.000113)
stay-at-home×pandemic duration 3	0.000143	0.000164
	(0.000125)	(0.000141)
stay-at-home×pandemic duration 4	0.000098	0.000140
	(0.000103)	(0.000124)
stay-at-home×pandemic duration 1×person		−0.000140
		(0.000242)
stay-at-home×pandemic duration 2×person		−0.000219*
		(0.000106)
stay-at-home×pandemic duration 3×person		−0.000084
		(0.000109)
stay-at-home×pandemic duration 4×person		−0.000120
		(0.000109)
Observations	8,695,668	8,695,668
R <sup>2</sup>	0.000428	0.000428
Adjusted R <sup>2</sup>	0.000235	0.000234

*Note:* The fixed effects include month, state, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .



Table 3.11: Table of firm-defined conversions – segmenting the pandemic duration

	<i>Dependent variable:</i>	
	Firm-Defined Conversions	
	(1)	(2)
stay-at-home×pandemic duration 1	0.0158*	0.0165*
	(0.0077)	(0.0082)
stay-at-home×pandemic duration 2	0.0113	0.0119
	(0.0062)	(0.0067)
stay-at-home×pandemic duration 3	0.0113	0.0123
	(0.0071)	(0.0077)
stay-at-home×pandemic duration 4	0.0199*	0.0275*
	(0.0089)	(0.0129)
stay-at-home×pandemic duration 1×person		−0.0026
		(0.0058)
stay-at-home×pandemic duration 2×person		−0.0029
		(0.0045)
stay-at-home×pandemic duration 3×person		−0.0051
		(0.0047)
stay-at-home×pandemic duration 4×person		−0.0199
		(0.0114)
Observations	8,695,668	8,695,668
R <sup>2</sup>	0.0324	0.0324
Adjusted R <sup>2</sup>	0.0322	0.0322

*Note:* The fixed effects include month, state, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 3.12: Table of alternative independent variables on ad conversions

	<i>Dependent variable:</i>			
	Conversions		Firm-defined Conversions	
	(1)	(2)	(3)	(4)
mall cvi	-0.000573 (0.000292)		-0.011060 (0.006605)	
COVID start	0.000014 (0.000030)		0.002711 (0.001860)	
weekly.death.per.million		0.00000003 (0.0000001)		0.000005 (0.000004)
log spend	0.000006*** (0.000001)	0.000006*** (0.000001)	0.000178** (0.000057)	0.000179** (0.000057)
CPM	-0.012075 (0.009653)	-0.018072 (0.013586)	-0.125246 (0.217248)	-0.169080 (0.265620)
campaign duration	-0.000001 (0.0000005)	-0.000001 (0.000001)	-0.000003 (0.000009)	-0.0000005 (0.000007)
person	-0.000001 (0.000006)	-0.000001 (0.000006)	0.000092 (0.000218)	0.000101 (0.000216)
logo	0.000002 (0.000007)	0.000003 (0.000007)	0.000707 (0.000479)	0.000893 (0.000575)
text	-0.000003 (0.000006)	-0.000005 (0.000006)	-0.000392 (0.000333)	-0.000480 (0.000340)
speech	-0.000014** (0.000005)	-0.000014** (0.000005)	-0.000116 (0.000182)	-0.000120 (0.000189)
mall cvi×COVID start	-0.000237* (0.000118)		-0.004759 (0.002412)	
Observations	8,181,696	8,695,668	8,181,696	8,695,668
R <sup>2</sup>	0.000424	0.000422	0.032390	0.032313
Adjusted R <sup>2</sup>	0.000219	0.000229	0.032191	0.032126

*Note:* The fixed effects include month, state, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level.  $p < 0.10$ ; \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 3.13: Table of conversions – Facebook platform

	<i>Dependent variable:</i>		
	Conversion		
	(1)	(2)	(3)
stay-at-home	0.1341 (0.0754)	0.0775 (0.0639)	0.1034 (0.0936)
full vacc. rate		−0.0003 (0.0006)	
log spend	−0.0004 (0.0007)	0.0001 (0.0007)	−0.0004 (0.0007)
campaign duration	−0.0010 (0.0005)	−0.0008 (0.0005)	−0.0010 (0.0005)
CPM	1.2391** (0.4647)	0.9450* (0.3857)	1.2446** (0.4617)
person	−0.0028 (0.0029)	−0.0024 (0.0026)	−0.0117 (0.0131)
logo	0.0002 (0.0048)	−0.0019 (0.0048)	0.0002 (0.0048)
text	0.0047 (0.0028)	0.0042 (0.0027)	0.0047 (0.0028)
speech	−0.0155* (0.0063)	−0.0154* (0.0063)	−0.0155* (0.0064)
stay-at-home×full vacc. rate		−0.0037 (0.0019)	
stay-at-home×person			0.0431 (0.0605)
Observations	16,144,988	15,945,138	16,144,988
R <sup>2</sup>	0.6089	0.6101	0.6089
Adjusted R <sup>2</sup>	0.6086	0.6098	0.6086

*Note:* The fixed effects include month, country, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level.  $p < 0.10$ ;  $*p < 0.05$ ;  $**p < 0.01$ ;  $***p < 0.001$ .

Table 3.14: Table of alternative independent variables on 2s views

	<i>Dependent variable:</i>	
	Two-Second Views	
	(1)	(2)
mall cvi	0.0888 (0.0556)	
COVID start	0.0306** (0.0108)	
weekly.death.per.million		0.00004 (0.00003)
log spend	0.0044** (0.0015)	0.0044** (0.0015)
CPM	7.3826** (2.5675)	8.4819** (2.7836)
campaign duration	-0.0010** (0.0004)	-0.0010** (0.0004)
person	0.3147*** (0.0219)	0.3148*** (0.0219)
logo	0.0066 (0.0191)	0.0081 (0.0194)
text	-0.0143 (0.0276)	-0.0148 (0.0277)
speech	0.0649*** (0.0172)	0.0651*** (0.0173)
mall cvi×COVID start	-0.0379 (0.0214)	
Observations	8,181,696	8,695,668
R <sup>2</sup>	0.3530	0.3531
Adjusted R <sup>2</sup>	0.3529	0.3529

*Note:* The fixed effects include month, state, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 3.15: Robustness check – using DMA as location unit instead of state

	<i>Dependent variable:</i>		
	2s Views	Conversions(FD)	Conversions
	(1)	(2)	(3)
stay-at-home	−0.08115* (0.03917)	−0.00472 (0.00622)	0.000043 (0.000083)
COVID start	0.01435 (0.01513)	−0.00240 (0.00156)	−0.000094 (0.000058)
CPM	4.76552 (2.64711)	−0.09000 (0.22029)	−0.00905 (0.00934)
campaign duration	−0.00100*** (0.00029)	−0.00002 (0.00002)	−0.00000014 (0.00000041)
log spend	0.00151 (0.00088)	0.00015** (0.00005)	0.0000050*** (0.0000010)
person	0.31837*** (0.03005)	0.00006 (0.00029)	−0.00000 (0.00001)
logo	−0.00993 (0.02335)	0.00140 (0.00117)	0.0000037 (0.0000063)
text	−0.01265 (0.02635)	−0.00027 (0.00039)	−0.0000011 (0.000006)
speech	0.05467** (0.01848)	0.00038 (0.00043)	−0.000011 (0.000007)
stay-at-home×COVID start	0.0574 (0.03809)	0.0148* (0.00676)	0.00025* (0.00012)
Observations	47,990,560	47,990,560	47,990,560
R <sup>2</sup>	0.36828	0.03804	0.00038
Adjusted R <sup>2</sup>	0.36826	0.03800	0.00034

*Note:* The fixed effects include month, DMA, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 3.16: Robustness Check – additional video feature covariates

	<i>Dependent variable:</i>		
	2s Views (1)	Conversions(FD) (2)	Conversions (3)
stay-at-home	0.00304 (0.03681)	−0.00697 (0.00647)	0.000079 (0.000109)
COVID start	0.01302 (0.01127)	−0.00124 (0.00129)	−0.000093 (0.000066)
log spend	0.00512*** (0.00124)	0.00017** (0.00005)	0.000006*** (0.000001)
campaign duration	−0.00032 (0.00022)	−0.000005 (0.000009)	−0.000001 (0.0000005)
CPM	2.80209 (2.43649)	−0.11175 (0.21288)	−0.011484 (0.009219)
person	0.22498*** (0.02689)	0.00054 (0.00048)	0.000019 (0.000017)
logo	0.00851 (0.01283)	0.00067 (0.00045)	−0.000001 (0.000007)
text	−0.18483*** (0.03825)	0.00010 (0.00039)	0.000010 (0.000009)
speech	0.06079*** (0.01401)	−0.00034 (0.00021)	−0.000020** (0.000006)
stay-at-home:COVID start	0.00478 (0.029806)	0.0139* (0.006624)	0.00028* (0.000141)
Observations	8,695,668	8,695,668	8,695,668
R <sup>2</sup>	0.379894	0.032375	0.000425
Adjusted R <sup>2</sup>	0.379772	0.032184	0.000228

*Note:* The fixed effects include month, state, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level. The additional 33 video feature covariates include the duration and average size of each creative element, number of people present, average confidence of emotion sub-elements, the duration of each gaze direction, gender and predicted age range sub-elements.

\* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .

Table 3.17: Table of 3s views – Facebook platform

	<i>Dependent variable:</i>		
	Three-second Views		
	(1)	(2)	(3)
stay-at-home	−0.0362*	−0.0413**	−0.0168
	(0.0141)	(0.0133)	(0.0190)
full vacc. rate		−0.0002	
		(0.0001)	
log spend	0.0010	0.0010	0.0011
	(0.0007)	(0.0007)	(0.0007)
campaign duration	−0.0001*	−0.0001*	−0.0001*
	(0.0001)	(0.0001)	(0.0001)
CPM	−0.0493	−0.0589	−0.0528
	(0.0516)	(0.0532)	(0.0514)
person	0.0164***	0.0164***	0.0220***
	(0.0024)	(0.0024)	(0.0046)
logo	−0.0068*	−0.0068*	−0.0068*
	(0.0034)	(0.0034)	(0.0034)
text	0.0040	0.0040	0.0040
	(0.0022)	(0.0023)	(0.0022)
speech	0.0267***	0.0267***	0.0267***
	(0.0037)	(0.0037)	(0.0037)
stay-at-home×full vacc. rate		0.0008	
		(0.0006)	
stay-at-home×person			−0.0272
			(0.0210)
Observations	16,144,988	15,945,138	16,144,988
R <sup>2</sup>	0.1478	0.1479	0.1478
Adjusted R <sup>2</sup>	0.1472	0.1472	0.1472

*Note:* The fixed effects include month, country, campaign ID, industry, and objective. Robust standard errors clustered at campaign and time level. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ .





# Bibliography

- [1] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. Discrimination through optimization: How facebook’s ad delivery can lead to biased outcomes. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–30, 2019.
- [2] Ali Alkhatib and Michael Bernstein. Street-level algorithms: A theory at the gaps between policy and decisions. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2019.
- [3] Judy I Alpert and Mark I Alpert. Background music as an influence in consumer mood and advertising responses. *ACR North American Advances*, 1989.
- [4] Kyle Bagwell. The economic analysis of advertising. *Handbook of industrial organization*, 3:1701–1844, 2007.
- [5] Kyle Bagwell and Garey Ramey. Coordination economies, advertising, and search behavior in retail markets. *The American Economic Review*, pages 498–517, 1994.
- [6] Saeideh Bakhshi, David A Shamma, and Eric Gilbert. Faces engage us: Photos with faces attract more likes and comments on instagram. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 965–974, 2014.
- [7] Daniel Belanche, Carlos Flavián, and Alfredo Pérez-Rueda. Understanding interactive online advertising: Congruence and product involvement in highly and lowly arousing, skippable video ads. *Journal of Interactive Marketing*, 37:75–88, 2017.
- [8] Marianne Bertrand, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman. What’s advertising content worth? evidence from a consumer credit marketing field experiment. *The Quarterly Journal of Economics*, 125(1):263–306, 2010.
- [9] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [10] Margaret C Campbell, J Jeffrey Inman, Amna Kirmani, and Linda L Price. In times of trouble: A framework for understanding consumers’ responses to threats, 2020.

- [11] Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- [12] M Chattington, N Reed, D Basacik, A Flint, and A Parkes. Investigating driver distraction: the effects of video and static advertising. *TRL Published Project Report*, 2010.
- [13] Chiao-Chieh Chen and Yu-Ping Chiu. Advertising content and online engagement on social media during the covid-19 epidemic in taiwan. *Journal of Marketing Communications*, pages 1–19, 2021.
- [14] Khai Chiong, Matthew Shum, Ryan Webb, and Richard Chen. Split-second decision-making in the field: Response times in mobile advertising. *SSRN*, 3289386:19, 2018.
- [15] Khai Chiong, Matthew Shum, Ryan Webb, and Richard Chen. Combining choices and response times in the field: a drift-diffusion model of mobile advertisements. *Available at SSRN 3289386*, 2020.
- [16] J. Clement. Most used social media 2020, 11 2020.
- [17] Bo Cowgill and Catherine Tucker. Economics, fairness and algorithmic bias. *preparation for: Journal of Economic Perspectives*, 2019.
- [18] Pragyan Deb, Davide Furceri, Daniel Jimenez, Siddharth Kothari, Jonathan D Ostry, and Nour Tawk. The effects of covid-19 vaccines on economic activity. *Swiss Journal of Economics and Statistics*, 158(1):1–25, 2022.
- [19] Mathias Dewatripont, Ian Jewitt, and Jean Tirole. Multitask agency problems: Focus and task clustering. *European economic review*, 44(4-6):869–877, 2000.
- [20] Anthony J Dukes, Qihong Liu, and Jie Shuai. Interactive advertising: The case of skippable ads. *Available at SSRN 3169629*, 2019.
- [21] Dean Eckles, Brett R Gordon, and Garrett A Johnson. Field studies of psychologically targeted ads face threats to internal validity. *Proceedings of the National Academy of Sciences*, 115(23):E5254–E5255, 2018.
- [22] Ludvík Eger, Lenka Komárková, Dana Egerová, and Michal Mičík. The effect of covid-19 on consumer shopping behaviour: Generational cohort perspective. *Journal of Retailing and Consumer Services*, 61:102542, 2021.
- [23] Ernst Fehr and Klaus M Schmidt. Fairness and incentives in a multi-task principal–agent model. *scandinavian Journal of Economics*, 106(3):453–474, 2004.
- [24] Shaoxiong Fu, Hongxiu Li, Yong Liu, Henri Pirkkalainen, and Markus Salo. Social media overload, exhaustion, and use discontinuance: Examining the effects of information overload, system feature overload, and social overload. *Information Processing & Management*, 57(6):102307, 2020.

- [25] Chelsea Galoni, Gregory S Carpenter, and Hayagreeva Rao. Disgusted and afraid: Consumer choices under the threat of contagious disease. *Journal of Consumer Research*, 47(3):373–392, 2020.
- [26] Marilyn Giroux, Jooyoung Park, Jae-Eun Kim, Yung Kyun Choi, Jacob C Lee, Seongseop Kim, Seongsoo Jang, Hector Gonzalez-Jimenez, and Jungkeun Kim. The impact of communication information on the perceived threat of covid-19 and stockpiling intention. *Australasian Marketing Journal*, page 18393349211028670, 2021.
- [27] Baptiste Gregorutti, Bertrand Michel, and Philippe Saint-Pierre. Grouped variable importance with random forests and application to multiple functional data analysis. *Computational Statistics & Data Analysis*, 90:15–35, 2015.
- [28] Orit Hershler and Shaul Hochstein. At first sight: A high-level pop out effect for faces. *Vision research*, 45(13):1707–1724, 2005.
- [29] Bengt Holmstrom and Paul Milgrom. Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *JL Econ. & Org.*, 7:24, 1991.
- [30] Yu-Chen Hsieh and Kuo-Hsiang Chen. How different information types affect viewer’s attention on internet advertising. *Computers in human Behavior*, 27(2):935–945, 2011.
- [31] Yunhui Huang and Jaideep Sengupta. The influence of disease cues on preference for typical versus atypical products. *Journal of Consumer Research*, 47(3):393–411, 2020.
- [32] June Hwang, Kosuke Imai, and Alex Tarr. Automated coding of political campaign advertisement videos: An empirical validation study. *PolMeth*, 36(1):1, 2019.
- [33] Tina Iachini, Francesca Frassinetti, Francesco Ruotolo, Filomena Leonela Sbordone, Antonella Ferrara, Maria Arioli, Francesca Pazzaglia, Andrea Bosco, Michela Candini, Antonella Lopez, et al. Social distance during the covid-19 pandemic reflects perceived rather than actual risk. *International Journal of Environmental Research and Public Health*, 18(11):5504, 2021.
- [34] Mengtian Jiang, Jing Yang, Eunsin Joo, and Taeyoung Kim. The effect of ad authenticity on advertising value and consumer engagement: A case study of covid-19 video ads. *Journal of Interactive Advertising*, pages 1–9, 2022.
- [35] Garrett A Johnson, Randall A Lewis, and Elmar I Nubbemeyer. Ghost ads: Improving the economics of measuring online ad effectiveness. *Journal of Marketing Research*, 54(6):867–884, 2017.
- [36] Ha Eun Kim, Yoon-Na Cho, and Nara Youn. Covid-19 uncertainty and temporal framing in advertising for online experiential consumption. *Journal of Advertising*, 50(3):280–289, 2021.

- [37] Sally M Kuehn and Pierre Jolicoeur. Impact of quality of the image, orientation, and similarity of the stimuli on visual search for faces. *Perception*, 23(1):95–122, 1994.
- [38] Samuli Laato, AKM Najmul Islam, Ali Farooq, and Amandeep Dhir. Unusual purchasing behavior during the early stages of the covid-19 pandemic: The stimulus-organism-response approach. *Journal of Retailing and Consumer Services*, 57:102224, 2020.
- [39] Anja Lambrecht and Catherine Tucker. Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads. *Management Science*, 65(7):2966–2981, 2019.
- [40] Stephen RH Langton, Anna S Law, A Mike Burton, and Stefan R Schweinberger. Attention capture by faces. *Cognition*, 107(1):330–342, 2008.
- [41] Xuan Liu, Savannah Shi, Thales Teixeira, and Michel Wedel. Video content marketing: The making of clips. *Journal of Marketing*, 82(4):86–101, 2018.
- [42] Andrew N Mason, John Narcum, and Kevin Mason. Social media marketing gains importance after covid-19. *Cogent Business & Management*, 8(1):1870797, 2021.
- [43] Jörg Matthes, Kathrin Karsay, Desirée Schmuck, and Anja Stevic. “too much to handle”: Impact of mobile social networking sites on information overload, depressive symptoms, and well-being. *Computers in Human Behavior*, 105:106217, 2020.
- [44] Dina Mayzlin and Jiwoong Shin. Uninformative advertising as an invitation to search. *Marketing Science*, 30(4):666–685, 2011.
- [45] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635*, 2019.
- [46] Lukas Meier, Sara Van De Geer, and Peter Bühlmann. The group lasso for logistic regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(1):53–71, 2008.
- [47] Nicolai Meinshausen and Peter Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473, 2010.
- [48] Jooyoung Park, Jungkeun Kim, Daniel C Lee, Seongseop S Kim, Benjamin G Voyer, Changju Kim, Billy Sung, Hector Gonzalez-Jimenez, Fernando Fastoso, Yung K Choi, et al. The impact of covid-19 on consumer evaluation of authentic advertising messages. *Psychology & Marketing*, 39(1):76–89, 2022.
- [49] Sahil Patel. 85 percent of facebook video is watched without sound, 5 2016.

- [50] Ana Luisa Pedrosa, Letícia Bitencourt, Ana Cláudia Fontoura Fróes, Maria Luíza Barreto Cazumbá, Ramon Gustavo Bernardino Campos, Stephanie Bruna Camilo Soares de Brito, and Ana Cristina Simões e Silva. Emotional, behavioral, and psychological impact of the covid-19 pandemic. *Frontiers in psychology*, page 2635, 2020.
- [51] Erica Perry. 2020 video marketing and statistics: What brands need to know, 10 2019.
- [52] Rik Pieters, Luk Warlop, and Michel Wedel. Breaking through the clutter: Benefits of advertisement originality and familiarity for brand attention and memory. *Management science*, 48(6):765–781, 2002.
- [53] Rik Pieters and Michel Wedel. Attention capture and transfer in advertising: Brand, pictorial, and text-size effects. *Journal of Marketing*, 68(2):36–50, 2004.
- [54] Gabriele Prati and Anthony D Mancini. The psychological impact of covid-19 pandemic lockdowns: a review and meta-analysis of longitudinal studies and natural experiments. *Psychological medicine*, 51(2):201–211, 2021.
- [55] Gretchen R Ross, Margaret G Meloy, and Kurt A Carlson. Preference refinement after a budget contraction. *Journal of Consumer Research*, 47(3):412–430, 2020.
- [56] John R Rossiter and Larry Percy. Attitude change through visual imagery in advertising. *Journal of Advertising*, 9(2):10–16, 1980.
- [57] Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.
- [58] Norman B Schmidt, Alex D Martin, Nicholas P Allan, Brittany M Mathes, Kevin G Saulnier, and David S March. Actual versus perceived infection rates of covid-19: Impact on distress, behavior and disability. *Journal of Psychiatric Research*, 137:525–533, 2021.
- [59] Sydney E Scott, Paul Rozin, and Deborah A Small. Consumers prefer “natural” more for preventatives than for curatives. *Journal of Consumer Research*, 47(3):454–471, 2020.
- [60] Sivaramakrishnan Siddarth and Amitava Chattopadhyay. To zap or not to zap: A study of the determinants of channel switching during commercials. *Marketing Science*, 17(2):124–138, 1998.
- [61] Linden Ryan Skeens. Council post: Video advertising trends going into 2021, 10 2020.
- [62] Saira Hanif Soroya, Ali Farooq, Khalid Mahmood, Jouni Isoaho, and Shan-e Zara. From information seeking to information avoidance: Understanding the health information behavior during a global health crisis. *Information processing & management*, 58(2):102440, 2021.

- [63] Danang Tandyonomanu. Ads on youtube: Skip or watch? In *1st International Conference on Social Sciences (ICSS 2018)*, pages 325–328. Atlantis Press, 2018.
- [64] Michael L Tee, Cherica A Tee, Joseph P Anlacan, Katrina Joy G Aligam, Patrick Wincy C Reyes, Vipat Kuruchittham, and Roger C Ho. Psychological impact of covid-19 pandemic in the philippines. *Journal of affective disorders*, 277:379–391, 2020.
- [65] Christian Testa, Nancy Krieger, Jarvis Chen, and William Hanage. Visualizing the lagged connection between covid-19 cases and deaths in the united states: An animation using per capita state-level data (january 22, 2020–july 8, 2020). *Working Paper*, 2020.
- [66] Robert L Thorndike. Who belongs in the family? *Psychometrika*, 18(4):267–276, 1953.
- [67] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [68] Amir Tohidi, Dean Eckles, and Ali Jadbabaie. Habits in consumer purchases: Evidence from store closures. *Available at SSRN 4077391*, 2022.
- [69] Michel Wedel and Rik Pieters. Eye fixations on advertisements and memory for brands: A model and findings. *Marketing science*, 19(4):297–312, 2000.
- [70] Michel Wedel and Rik Pieters. A review of eye-tracking research in marketing. *Review of marketing research*, 4(2008):123–147, 2008.
- [71] Jeremy Yang, Juanjuan Zhang, and Yuhan Zhang. First law of motion: Influencer video advertising on tiktok. *Available at SSRN 3815124*, 2021.
- [72] Joonhyuk Yang, Yingkang Xie, Lakshman Krishnamurthi, and Purushottam Patla. High-energy audio in advertising: A large-scale investigation of tv commercials. *Available at SSRN*, 2020.
- [73] Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- [74] Zhe Zhao, Lichan Hong, Li Wei, Jilin Chen, Aniruddh Nath, Shawn Andrews, Aditee Kumthekar, Maheswaran Sathiamoorthy, Xinyang Yi, and Ed Chi. Recommending what video to watch next: a multitask ranking system. In *Proceedings of the 13th ACM Conference on Recommender Systems*, pages 43–51, 2019.