



1-2018

Advertising Content and Consumer Engagement on Social Media: Evidence from Facebook

Dokyun Lee

Kartik Hosanagar
University of Pennsylvania

Harikesh Nair

Follow this and additional works at: https://repository.upenn.edu/marketing_papers

 Part of the [Advertising and Promotion Management Commons](#), [Business Administration, Management, and Operations Commons](#), [Business Analytics Commons](#), [Business and Corporate Communications Commons](#), [Communication Technology and New Media Commons](#), [Marketing Commons](#), [Mass Communication Commons](#), [Social Media Commons](#), and the [Technology and Innovation Commons](#)

Recommended Citation

Lee, D., Hosanagar, K., & Nair, H. (2018). Advertising Content and Consumer Engagement on Social Media: Evidence from Facebook. *Management Science*, <http://dx.doi.org/10.1287/mnsc.2017.2902>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/marketing_papers/339
For more information, please contact repository@pobox.upenn.edu.

Advertising Content and Consumer Engagement on Social Media: Evidence from Facebook

Abstract

We describe the effect of social media advertising content on customer engagement using data from Facebook. We content-code 106,316 Facebook messages across 782 companies, using a combination of Amazon Mechanical Turk and natural language processing algorithms. We use this data set to study the association of various kinds of social media marketing content with user engagement—defined as Likes, comments, shares, and click-throughs—with the messages. We find that inclusion of widely used content related to brand personality—like humor and emotion—is associated with higher levels of consumer engagement (Likes, comments, shares) with a message. We find that directly informative content—like mentions of price and deals—is associated with lower levels of engagement when included in messages in isolation, but higher engagement levels when provided in combination with brand personality-related attributes. Also, certain directly informative content, such as deals and promotions, drive consumers' path to conversion (click-throughs). These results persist after incorporating corrections for the nonrandom targeting of Facebook's EdgeRank (News Feed) algorithm and so reflect more closely user reaction to content than Facebook's behavioral targeting. Our results suggest that there are benefits to content engineering that combines informative characteristics that help in obtaining immediate leads (via improved click-throughs) with brand personality-related content that helps in maintaining future reach and branding on the social media site (via improved engagement). These results inform content design strategies. Separately, the methodology we apply to content-code text is useful for future studies utilizing unstructured data such as advertising content or product reviews.

Keywords

consumer engagement, social media, advertising content, content engineering, marketing communication, machine learning, natural language processing, selection, Facebook, EdgeRank, News Feed algorithm

Disciplines

Advertising and Promotion Management | Business | Business Administration, Management, and Operations | Business Analytics | Business and Corporate Communications | Communication Technology and New Media | Marketing | Mass Communication | Social Media | Technology and Innovation

The Effect of Advertising Content on Consumer Engagement: Evidence from Facebook*

Dokyun Lee
The Wharton School

Kartik Hosanagar
The Wharton School

Harikesh S. Nair
Stanford GSB

Abstract

We investigate the effect of social media content on customer engagement using a large-scale field study on Facebook. We content-code more than 100,000 unique messages across 800 companies engaging with users on Facebook using a combination of Amazon Mechanical Turk and state-of-the-art Natural Language Processing algorithms. We use this large-scale database of advertising attributes to test the effect of ad content on subsequent user engagement – defined as *Likes* and comments – with the messages. We develop methods to account for potential selection biases that arise from Facebook’s filtering algorithm, EdgeRank, that assigns posts non-randomly to users. We find that inclusion of persuasive content – like emotional and philanthropic content – *increases* engagement with a message. We find that informative content – like mentions of prices, availability and product features – *reduce* engagement when included in messages in isolation, but *increase* engagement when provided in combination with persuasive attributes. Persuasive content thus seems to be the key to effective engagement. Our results inform advertising design in social media, and the methodology we develop to content-code large-scale textual data provides a framework for future studies on unstructured natural language data such as advertising content or product reviews.

Keywords: advertising, social media, advertising content, large-scale data, natural language processing, selection, Facebook, EdgeRank.

*We thank seminar participants at the ISIS Conference (2013), Mack Institute Conference (Spring 2013), and SCECR Conference (Summer 2013) for comments, and a collaborating company that wishes to be anonymous for providing the data used in the analysis. The authors gratefully acknowledge the financial support from the Jay H. Baker Retailing Center and Mack Institute of the Wharton School and the Wharton Risk Center (Russell Ackoff Fellowship). All errors are our own.

1 Introduction

Social media is increasingly taking up a greater share of consumers' time spent online and, as a result, is becoming a larger component of firm's advertising budgets. Surveying 4,943 marketing decision makers at US companies, the 2013 Chief Marketing Officer survey (www.cmosurvey.org) reports that expected spending on social media marketing will grow from 8.4% of firms' total marketing budgets in 2013 to about 22% in the next 5 years. As firms increase their social media activity, the role of *content engineering* has become increasingly important. Content engineering seeks to develop ad content that better engage targeted users and drive the desired goals of the marketer from the campaigns they implement. Surprisingly however, despite the numerous insights from the applied psychology literature about the design of the ad-creative and its obvious relevance to practice, relatively little has been formally established about the empirical consequences of advertising content outside the laboratory, in real-world, field settings. Ad content also is under emphasized in economic theory. The canonical economic model of advertising as a signal (c.f. Nelson (1974); Kihlstrom and Riordan (1984); Milgrom and Roberts (1986)) does not postulate any direct role for ad content because advertising intensity conveys all relevant information about product quality in equilibrium to market participants. Models of informative advertising (c.f. Butters (1977); Grossman and Shapiro (1984)) allow for advertising to inform agents only about price and product existence – yet, casual observation and several studies in lab settings (c.f. Armstrong (2010)) suggest advertisements contain much more information and content beyond prices. In this paper, we investigate the role of content in driving consumer engagement in social media in a field setting and document that content matters significantly. We find that a variety of emotional, philanthropic and informative advertising content attributes affect engagement and that the role of content varies significantly across firms and industries. The richness of our engagement data and the ability to content code ads in a cost-efficient manner enables us to study the problem at a larger scale than much of the previous literature on the topic.

Our analysis is of direct relevance to industry in better understanding and improving firms' social media marketing strategies. Recent studies (e.g., Creamer 2012) report that only about 1% of an average firm's Facebook fans (users who have *Liked* the Facebook Page of the firm) actually engage with the brand by commenting on, *Liking* or sharing posts by the firm on the platform. As a result, designing better advertising content that achieves superior reach and engagement on social media is an important issue for marketing on this new medium. While many brands have established a social media presence, it is not clear what kind of content works better and for which firm, and in what way. For example, are posts seeking to inform consumers about product or price attributes more effective than persuasive messages? Are videos or photos more likely to engage users relative to simple status updates? Do messages explicitly soliciting user response (e.g., "*Like* this post if ...") draw more engagement or in fact turn users away? Does the same strategy apply across different industries? Our paper explores these kinds of questions and contributes to the formulation of better content engineering policies in practice.

Our empirical investigation is implemented on Facebook, which is the largest social media platform in the world. Many top brands now maintain a Facebook page from which they serve posts and messages to connected users. This is a form of free social media advertising that has increasingly become a popular and

important channel for marketing. Our data comprises information on about 100,000 such messages posted by a panel of about 800 firms over a 11-month period between September 2011 and July 2012. For each post, our data also contains time-series information on two kinds of engagement measures – *Likes* and comments – observed on Facebook. We supplement these engagement data with message attribute information that we collect using a large-scale survey we implement on Amazon Mechanical Turk (henceforth “AMT”), combined with a Natural Language Processing algorithm (henceforth “NLP”) we build to tag messages. We incorporate new methods and procedures to improve the accuracy of content tagging on AMT and our NLP algorithm. As a result, our algorithm achieves about 99% accuracy, recall and precision for almost all tagged content profiles. The methods we develop will be useful in future studies analyzing advertising content and product reviews.

Our data also has several advantages that facilitate a study of advertising content. First, Facebook posts have rich content attributes (unlike say, Twitter tweets, which are restricted in length) and rich data on user engagement. Second, Facebook requires real names and, therefore, data on user activity on Facebook is often more reliable compared to other social media sites. Third, engagement is measured on a daily basis (panel data) by actual post-level engagement such as *Likes* and comments that are precisely tracked within a closed system. These aspects make Facebook an almost ideal setting to study the effect of ad content.

Our strategy for coding content is motivated by the psychology, marketing and economic literatures on advertising (see Cialdini (2001); Chandy et al. (2001); Bagwell (2007); Vakratsas and Ambler (1999) for some representative overviews). In the economics literature, it is common to classify advertising as informative (shifting beliefs about product existence or prices) or persuasive (shifting preferences directly). The basis of information is limited to prices and/or existence, and persuasive content is usually treated as a “catch-all” without finer classification. Rather than this coarse distinction, our classification follows the seminal classification work of Resnik and Stern (1977), who operationalize informative advertising based on the number and characteristics of informational cues (see Abernethy and Franke, 1996 for an overview of studies in this stream). Some criteria for classifying content as informative include details about product deals, availability, price, and product related aspects that could be used in optimizing the purchase decision. Following this stream, any product oriented facts, and brand and product mentions are categorized as informative content. Following suggestions in the persuasion literature (Cialdini, 2001; Nan and Faber, 2004; Armstrong, 2010), we classify “persuasive” content as those that broadly seek to influence by appealing to *ethos*, *pathos* and *logos* strategies. For instance, the use of a celebrity to endorse a product or attempts to gain trust or good-will (e.g., via small talk, banter) can be construed as the use of *ethos* – appeals through credibility or character – and a form of persuasive advertising. Messages with philanthropic content that induce empathy can be thought of as an attempt at persuasion via *pathos* – an appeal to a person’s emotions. Lastly, messages with unusual or remarkable facts that influence consumers to adopt a product or capture their attention can be categorized as persuasion via *logos* – an appeal through logic. We categorize content that attempt to persuade and promote relationship building in this manner as persuasive content.

Estimation of the effect of content on subsequent engagement is complicated by the non-random allocation of messages to users implemented by Facebook via its *EdgeRank* algorithm. *EdgeRank* tends to serve to users posts that are newer and are expected to appeal better to his/her tastes. We develop corrections

to account for the filtering induced by *EdgeRank*. Our main finding from the empirical analysis is that persuasive content drives social media engagement significantly. Additionally, informative content tends to drive engagement positively only when combined with such content. Persuasive content thus seem to be the key to effective content engineering in this setting. The empirical results unpack the persuasive effect into component attribute effects and also estimate the heterogeneity in these effects across firms and industries.

We do not address the separate but important question of how engagement affects product demand and firm’s profits so as to complete the link between ad-attributes and those outcome measures. First, the data required for the analysis of this question at a scale comparable to this study are still not widely available to researchers. Second, firms and advertisers care about engagement *per se* and seem to be willing to invest in advertising for generating engagement, even though numerous academic studies starting with the well-known “split-cable” experiments of Lodish et al. (1995) have found that the effect of advertising on short-term sales is limited. Our view is that advertising is a dynamic problem and a dominant role of advertising is to build long-term brand-capital for the firm. Even though the current period effects of advertising on demand is small, the long-run effect of advertising may be large, generated by intermediary activities like increased consumer engagement, increased awareness and inclusion in the consumer consideration set. Thus, studying the formation and evolution of these intermediary activities – like engagement – may be worthwhile in order to better understand the true mechanisms by which advertising affects outcomes in market settings, and to resolve the tension between the negative results in academia and the continued investments in advertising in industry. This is where we see this paper as making a contribution. The inability to connect this engagement to firms’ profits and demand is an acknowledged limitation of this study.

Our paper adds to an emerging literature on the effects of ad content. A recent theoretical literature has developed new models that allow ad content to matter in equilibrium by augmenting the canonical signaling model in a variety of ways (e.g. Anand and Shachar (2009) by allowing ads to be noisy and targeted; Anderson and Renault (2006) by allowing ad content to resolve consumers’ uncertainty about their match-value with a product; and Mayzlin and Shin (2011) and Gardete (2013) by allowing ad content to induce consumers to search for more information about a product). Our paper is most closely related to a small empirical literature that has investigated the effects of ad content in field settings. These include Bertrand et al. (2010) (effect of direct-mail ad content on loan demand); Anand and Shachar (2011); Liaukonyte et al. (2013) (effect of TV ad content on viewership and online sales); Tucker (2012a) (effect of ad persuasion on YouTube video sharing) and Tucker (2012b) (effect of “social” Facebook ads on philanthropic participation). Also related are recent studies exploring the effect of content more generally (and not specifically ad content) including Berger and Milkman (2012) (effect of emotional content in New York Times articles on article sharing) and Gentzkow and Shapiro (2010) (effect of newspaper’s political content on readership). Finally, our paper is related to empirical studies on social media (reviewed in Sundararajan et al. (2013); Aral et al. (2013)). Relative to this literature, our study makes two main contributions. First, from a managerial standpoint, we show that while persuasive ad content – especially emotional and philanthropic content – positively impacts consumer engagement in social media, informative content has a negative effect unless it is combined with persuasive content attributes. This is particularly important for marketing managers who wish to use their social media presence to promote their brand and products. We also show how the insights

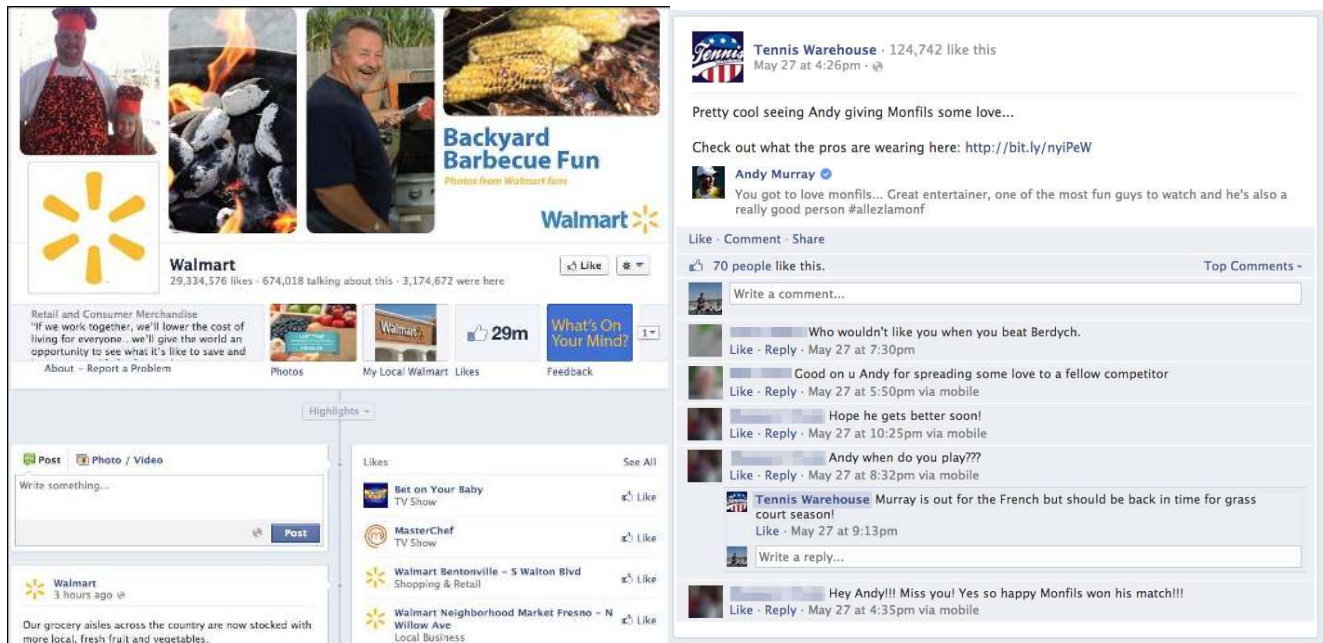


Figure 1: (Left) Example of a firm’s Facebook Page (Walmart). (Right) Example of a firm’s post and subsequent user engagement with that post (Tennis Warehouse). Example is not necessarily from our data.

differ by industry type. Second, none of the prior studies on ad content have been conducted at the scale of this study. The rigorous content-tagging methodology we develop, which combines surveys implemented on AMT with NLP-based algorithms, provides a framework to conduct large-scale studies analyzing content of advertising.

2 Data

Our dataset is derived from the “pages” feature offered by Facebook. The feature was introduced on Facebook in November 2007. Facebook Pages enable companies to create profile pages and to post status updates, advertise new promotions, ask questions and push content directly to consumers. The left panel of Figure 1 shows an example of Walmart’s Facebook Page, which is typical of the type of pages large companies host on the social network. In what follows, we use the terms pages, brands and firms interchangeably. Our data comprises posts served from firms’ pages onto the Facebook profiles of the users that are linked to the firm on the platform. To fix ideas, consider a typical post (see the right panel of Figure 1): “Pretty cool seeing Andy giving Monfils some love... Check out what the pros are wearing here: <http://bit.ly/nyiPeW>.”¹ In this status update, a tennis equipment retailer starts with small talk, shares details about a celebrity (Andy Murray and Gael Monfils) and ends with link to a product page. Each such post is a unit of analysis in our data.

¹Retailer picked randomly from an online search; not necessarily from our data.

2.1 Data Description

2.1.1 Raw Data and Selection Criteria

To collect the data, we partnered with an anonymous firm, henceforth referred to as Company X that provides analytical services to Facebook Page owners by leveraging data from Facebook’s *Insights*. *Insights* is an analytics tool provided by Facebook that allows companies to monitor the performance of their Facebook posts. Company X augments data from Facebook *Insights* across a large number of client firms with additional records of daily message characteristics, to produce a raw dataset comprising a post-day-level panel of messages posted by companies via their Facebook pages. The data also includes two consumer engagement metrics: the number of *Likes* and comments for each post each day. These metrics are commonly used in industry as measures of engagement. They are also more granular than other metrics used in extant research such as the number of fans who have *Liked* the page. Also available in the data are the number of impressions of each post per day (i.e., the total number of users the post is exposed to). In addition, page-day level information such as the aggregate demographics of users (fans) who *Liked* the page on Facebook or have ever seen posts by the page are collected by Company X on a daily level². This comprises the population of users a post from a firm can potentially be served to. We leverage this information in the methodology we develop later for accounting for non-random assignment of posts to users by Facebook. Once a firm serves a post, the post’s impressions, *Likes* and comments are recorded daily for an average of about 30 days (maximum: 126 days).³ The raw data contains about a million unique posts by about 2,600 unique companies. We clean the data to reflect the following criteria:

- Only pages located in the US.
- Only posts written in English.
- Only posts with complete demographics data.

After cleaning, the data span 106,316 unique messages posted by 782 companies (including many large brands) between September 2011 and July 2012. This results in about 1.3 million rows of post-level daily snapshots recording about 450 million page fans’ responses. Removing periods after which no significant activity is observed for a post reduces this to 665,916 rows of post-level snapshots (where activity is defined as either impressions, *Likes*, or comments). The companies in our dataset are categorized into 110 different industry categories as defined by Facebook. These finer categories are combined into 6 broader industry categories following Facebook’s page classification criteria. Table 1 shows these categories with examples.

2.1.2 Content-coded Data

We use a two-step method to label content. First, we contract with workers through AMT and tag 5,000 messages for a variety of content profiles. Subsequently, we build an NLP algorithm by combining several statistical classifiers and rule-based algorithms to extend the content-coding to the full set of 100,000 messages.

²In essence, our data is the most complete data outside of Facebook - the data includes more details and snapshots than what Facebook offers exclusively to page owners via the Application Programming Interface called Facebook Query Language.

³A vast majority of posts do not get any impression or engagement after 7 days. After 15 days, virtually all engagements and impressions (more than 99.9%) are accounted for.

CELEBRITY & PUBLIC FIGURE	ENTERTAINMENT	CONSUMER PRODUCTS & BRANDS
Actor & Director (Danny Boyle)	TV Shows (Star Trek)	Clothing (Ralph Lauren)
Athlete (Roger Federer)	Movies & Musics (Gattaca)	Book (Dune)
Musicians & Bands (Muse)	Recreation & Sports (Tennis)	Cars (Tesla Motors)
Government Official (Barack Obama)	Concert Tour (Coachella)	Food & Groceries (Trader Joe’s)
Author (Frank Herbert)	Entertainment (Monte Carlo)	Electronics (Nokia)
ORGANIZATIONS & COMPANY	WEBSITES	LOCAL PLACES & BUSINESSES
Health Agency (WHO)	Website (TED)	Local Business (The Halal Guys)
Non-profit Organization (Wikipedia)	Personal Website (Miller Photography)	Restaurants & Cafe (Olive Garden)
Government Organization (US Army)	App Pages (Google Search)	Museum & Art Gallery (MoMA)
University (University of Pennsylvania)		Hotel (Marriott)
Church & Religious (Catholic)		Legal & Law (American Bar Association)

Table 1: **Six Broader Categories of Pages and Some Examples of Finer Subcategories:** This table documents how base categories are merged into 6 broad categories. This follows the 6 broad page types listed on Facebook. Examples of actual pages (not necessarily from our data) are in parentheses.

This algorithm uses the 5,000 AMT-tagged messages as the training data-set. Best practices reported in the recent literature are used to ensure the quality of results from AMT and to improve the performance of the NLP algorithm (accuracy, recall, precision). The resulting NLP algorithm achieves around 99% accuracy, 99% recall and 99% precision for almost all the content profiles we consider with 10-fold cross validation. We describe these methods in more detail later in the paper.

The content in Facebook posts can be categorized as informative, persuasive, or both. Some messages inform consumers about deals and discounts about products, while other messages seek to connect with consumers on a personal level to promote brand personality, form relationships and are social in nature. We call the first type informative content, and the second persuasive content. Many messages do both at the same time by including casual banter and product information simultaneously (e.g., “Are you a tea person or a coffee person? Get your favorite beverage from our website”).

Table 2 outlines the finer classification of the attributes we code up, including precise definitions, summary statistics, and the source for coding the attribute. As mentioned, we content-code messages into various persuasive and informative attributes. In Table 2, the 8 variables: BRANDMENTION, DEAL, PRICECOMPARE, PRICE, TARGET, PRODAVAIL, PRODLOCATION, and PRODMENTION are informative. These variables enable us to assess the effect of search attributes, brand, price, and product availability information on engagement. The 8 variables: REMFACT, EMOTION, EMOTICON, HOLIDAYMENTION, HUMOR, PHILANTHROPIC, FRIENDLIKELY, and SMALLTALK are classified as persuasive. These definitions include emotional content, humor, banter and more complex content like the “FRIENDLIKELY” classification, which is a binary variable that reflect Mechanical Turk survey participant’s agreement that their “friends on social media are likely to post a message as the one shown.”

Besides these main variables of interest, controls and content-related patterns noted as important in industry reports are profiled. We include these content categories to investigate more formally considerations laid out in industry white papers, trade-press articles and blog reports about the efficacy of message attributes in social media engagement. It includes content that explicitly solicits readers to comment or includes blanks for users to fill out (thus providing an explicit option to facilitate engagement). Additionally,

characteristics like whether the message contained photos, website links, and the nature of the page-owner (business organization versus celebrity) are also coded. Other message-specific characteristics and controls include metrics such as message length in characters and SMOG (“Simple Measure of Gobbledygook”), an automatically computed reading complexity index that is used widely. Higher values of SMOG implies a message is harder to read. Table 3 shows sample messages taken from Walmart’s page in December 2012 and shows how we would have tagged them. The reader should note that some elements of content tagging and classification are necessarily subjective and based on human judgement. We discuss our methods (which involve obtaining agreement across 9 tagging individuals) in section 2.2. All things considered, we believe this is one of the most comprehensive attempts at tagging advertising content in the empirical literature.

2.1.3 Data Descriptive Graphics

This section presents descriptive statistics of the main stylized patterns in the data. Figure 2 shows box plots of the log of impressions, *Likes*, and comments versus the time (in days) since a post is released (τ). Both comments and *Likes* taper off to zero after two and six days respectively. The rate of decay of impressions is slower. Virtually all engagements and impressions (more than 99.9%) are accounted for within 15 days of release of a post.

Figure 3 shows the average number of *Likes* and comments by message type (photo, link, etc.) over the lifetime of a post. Messages with photos have the highest average *Likes* (94.7) and comments (7.0) over their lifetime. Status updates obtain more comments (5.5) on average than videos (4.6) but obtain less *Likes* than videos. Links obtain the lowest *Likes* on average (19.8) as well as the lowest comments (2.2). Figure 4 shows the same bar plots split across 6 industry categories. A consistent pattern is that messages with photos always obtain highest *Likes* across industries. The figure also documents interesting heterogeneity in engagement response across industries. The patterns in these plots echo those described in reports by many market research companies such as Wildfire and comScore.

Figure 5 presents the average number of *Likes* and comments by content attribute. Emotional messages obtain the most number of *Likes* followed by posts identified as “likely to be posted by friends” (variable: FRIENDLIKELY). Emotional content also obtain the highest number of comments on average followed by SMALLTALK and FRIENDLIKELY. The reader should note these graphs do not account for the market-size (i.e. the number of impressions a post reached). Later, we present an econometric model that incorporates market-size as well as selection by Facebook’s filtering algorithm to assess user engagement more formally.

Finally, Figure 6 shows the percentage of messages featuring a content attribute split by industry category. We represent the relative percentages in each cell by the size of the bubbles in the chart. The largest bubble is SMALLTALK for the celebrities category (60.4%) while the smallest is PRICECOMPARE for the celebrities category (0%). This means that 6 in 10 posts by celebrity pages in the data have some sort of small talk (banter) and/or content that does not relate to products or brands; and that there are no posts by celebrity owned pages that feature price comparisons. Interestingly, celebrity pages also do little targeting (i.e. via posts that explicitly call out to certain demographics or subpopulations with certain qualifications). “Remarkable facts” (our definition) are posted more by firms in the entertainment category and less by places and business-oriented pages. Consistent with intuition, consumer product pages and local places/businesses

Variable	Description	Source	Mean	SD	Min	Max
TAU (τ)	Time since the post release (Day)	Facebook	6.253	3.657	1	16
LIKES	Number of “Likes” post has obtained	Facebook	48.373	1017	0	324543
COMMENTS	Number of “Comments” post has obtained	Facebook	4.465	78.19	0	22522
IMPRESSIONS	Number of times message was shown to users	Facebook	9969.2	129874	1	4.5×10^7
SMOG	SMOG readability index (higher means harder to read)	Computed	7.362	2.991	3	25.5
MSGLEN	Message length in characters	Computed	157.41	134.54	1	6510
HTTP	Message contains a link	Computed	0.353	0.478	0	1
QUESTION	Message contains questions	Computed	0.358	0.479	0	1
BLANK	Message contains blanks (e.g. “My favorite artist is __”)	Computed	0.010	0.099	0	1
ASKLIKE	Explicit solicitation for “Likes” (e.g. “Like if ...”)	Computed	0.006	0.080	0	1
ASKCOMMENT	Explicit solicitation for “Comments”	Computed	0.001	0.029	0	1
PERSUASIVE						
REMACT	Remarkable fact mentioned	AMT	0.527	0.499	0	1
EMOTION	Any type of emotion present	AMT	0.524	0.499	0	1
EMOTICON	Contains emoticon or net slang (approximately 1000 scraped from web emoticon dictionary e.g. :D, LOL)	Computed	0.012	0.108	0	1
HOLIDAYMENTION	Mentions US Holidays	Computed	0.006	0.076	0	1
HUMOR	Humor used	AMT	0.375	0.484	0	1
PHILANTHROPIC	Philanthropic or activist message	AMT	0.498	0.500	0	1
FRIENDLY	Answer to question: “Are your friends on social media likely to post message such as the shown?”	AMT	0.533	0.499	0	1
SMALLTALK	Contains small talk or banter (defined to be content other than about a product or company business)	AMT	0.852	0.355	0	1
INFORMATIVE						
BRANDMENTION	Mentions a specific brand or organization name	AMT+Comp	0.264	0.441	0	1
DEAL	Contains deals: any type of discounts and freebies	AMT	0.620	0.485	0	1
PRICECOMPARE	Compares price or makes price match guarantee	AMT	0.442	0.497	0	1
PRICE	Contains product price	AMT+Comp	0.051	0.220	0	1
TARGET	Message is targeted towards an audience segment (e.g. demographics, certain qualifications such as “Moms”)	AMT	0.530	0.499	0	1
PRODAVAIL	Contains information on product availability (e.g. stock and release dates)	AMT	0.557	0.497	0	1
PRODLOCATION	Contains information on where to obtain product (e.g. link or physical location)	AMT	0.690	0.463	0	1
PRODMENTION	Specific product has been mentioned	AMT+Comp	0.146	0.353	0	1
MSGTYPE	Categorical message type assigned by the Facebook	Facebook				
- App	application related posts	Facebook	0.099	0.299	0	1
- Link	link	Facebook	0.389	0.487	0	1
- Photo	photo	Facebook	0.366	0.481	0	1
- Status Update	regular status update	Facebook	0.140	0.347	0	1
- Video	video	Facebook	0.005	0.070	0	1
PAGECATEGORY	Page category closely following Facebook’s categorization	Facebook				
- Celebrity	Singers, Actors, Athletes etc	Facebook	0.056	0.230	0	1
- ConsumerProduct	consumer electronics, packaged goods etc	Facebook	0.296	0.456	0	1
- Entertainment	Tv shows, movies etc	Facebook	0.278	0.447	0	1
- Organization	non-profit organization, government, school organization	Facebook	0.211	0.407	0	1
- PlaceBusiness	local places and businesses	Facebook	0.071	0.257	0	1
- Website	page about a website	Facebook	0.088	0.283	0	1

Table 2: **Variable Descriptions and Summary for Content-coded Data:** To interpret the “Source” column, note that “Facebook” means the values are obtained from Facebook, “AMT” means the values are obtained from Amazon Mechanical Turk and “Computed” means it has been either calculated or identified using online database resources and rule-based methods in which specific phrases or content (e.g. brands) are matched. Finally, “AMT+Computed” means primary data has been obtained from Amazon Mechanical Turk and it has been further augmented with online resources and rule-based methods.

Sample Messages	Content Tags
<i>Cheers! Let Welch's help ring in the New Year.</i>	BRANDMENTION, SMALLTALK, HOLIDAYMENTION, EMOTION
<i>Maria's mission is helping veterans and their families find employment. Like this and watch Maria's story. http://walmarturl.com/VzWFlh</i>	PHILANTHROPIC, SMALLTALK, ASKLIKE, HTTP
<i>On a scale from 1-10 how great was your Christmas?</i>	SMALLTALK, QUESTION, HOLIDAYMENTION
<i>Score an iPad 3 for an iPad2 price! Now at your local store, \$50 off the iPad 3. Plus, get a \$30 iTunes Gift Card. Offer good through 12/31 or while supplies last.</i>	PRODMENTION, DEAL, PRODLOCATION, PRODAVAIL, PRICE
<i>They're baaaaaack! Now get to snacking again. Find Pringles Stix in your local Walmart.</i>	EMOTION, PRODMENTION, BRANDMENTION, PRODLOCATION

Table 3: **Examples of Messages and Their Content Tags:** The messages are taken from 2012 December posts on Walmart's Facebook page.

post the most about products (PRODMENTION), product availability (PRODAVAIL), product location (PRODLOC), and deals (DEAL). Emotional (EMOTION) and philanthropic (PHILAN) content have high representation in pages classified as celebrity, organization and websites. Similarly, the AMT classifiers identify a larger portion of messages posted by celebrity, organization and website-based pages to be similar to posts by friends.

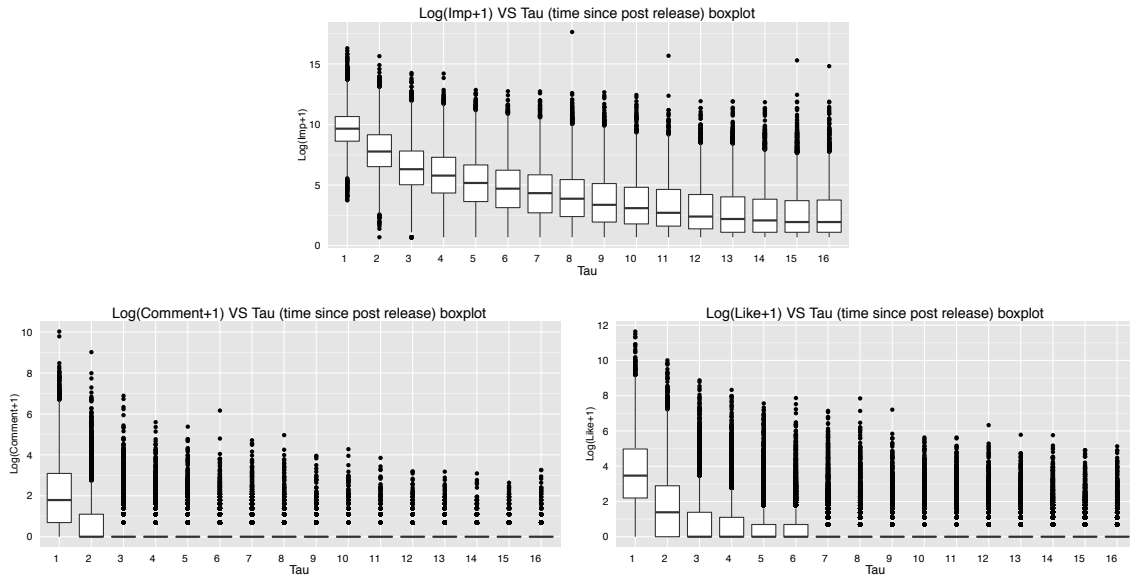


Figure 2: **Box Plots of Log(engagement+1) vs Time since Post Release:** Three graphs show the box plots of (log) impressions, comments and *Like* vs. τ respectively. Both comments and *Likes* taper to zero after two and six days respectively. On the other hand, impressions die out slower. After 15 days, virtually all engagements and impressions (more than 99.9%) are accounted for. There are many outliers.




Figure 3: **Average Likes and Comments by Message Type:** This figure shows the average number of *Likes* and comments obtained by posts over their lifetime on Facebook, split by message type.

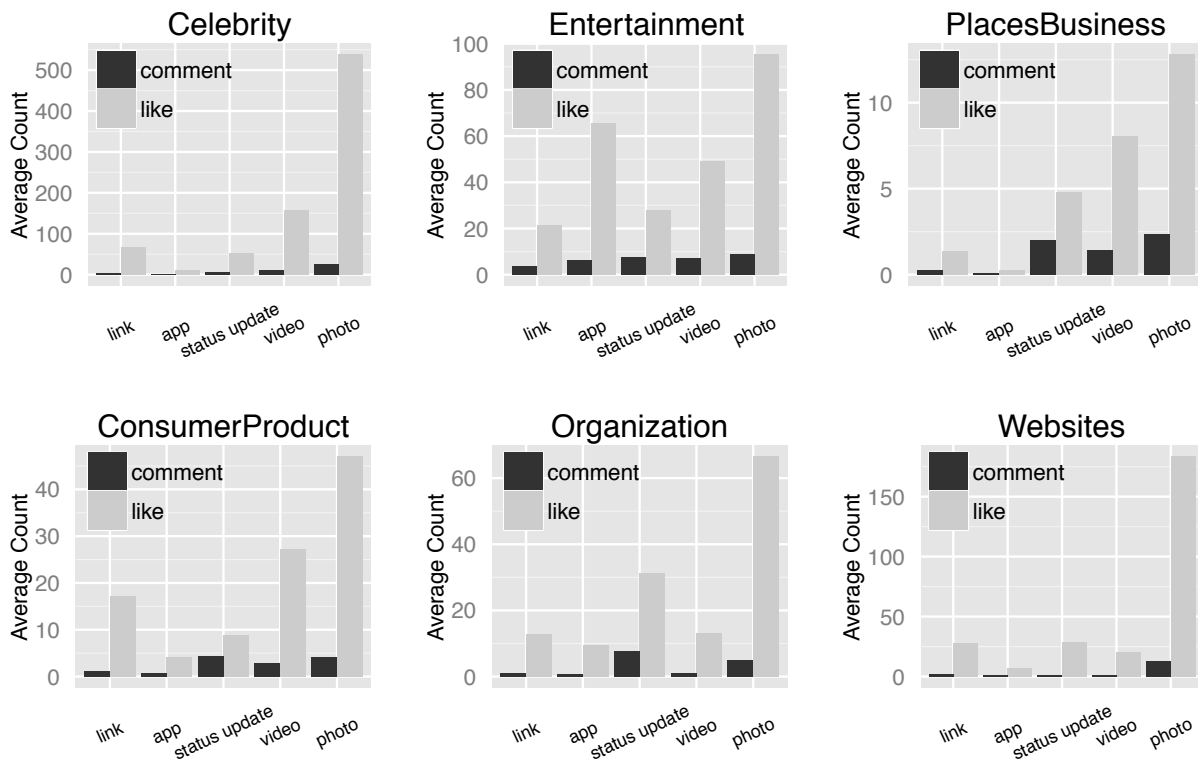


Figure 4: **Average Likes and Comments by Message Type by Industry:** This figure shows the average number of *Likes* and comments obtained by posts over their lifetime split by message type for each industry.

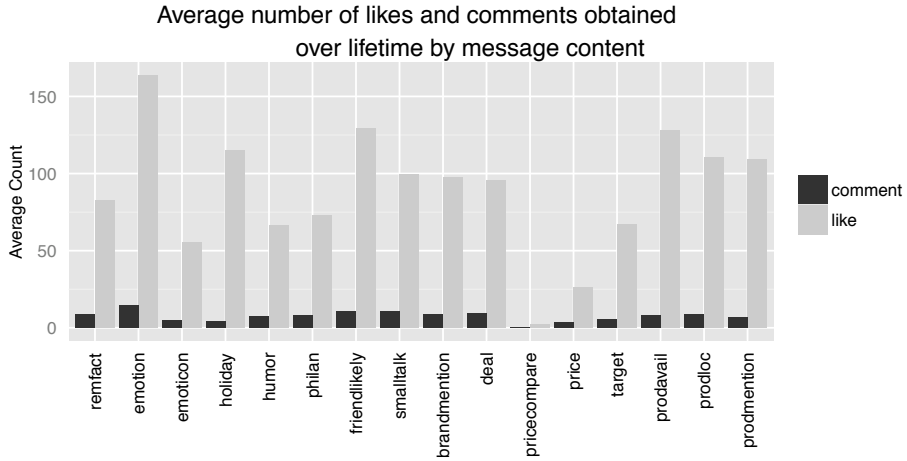


Figure 5: **Average Likes and Comments by Message Content:** This figure shows the average number of *Likes* and comments obtained by posts over their lifetime split by message content.

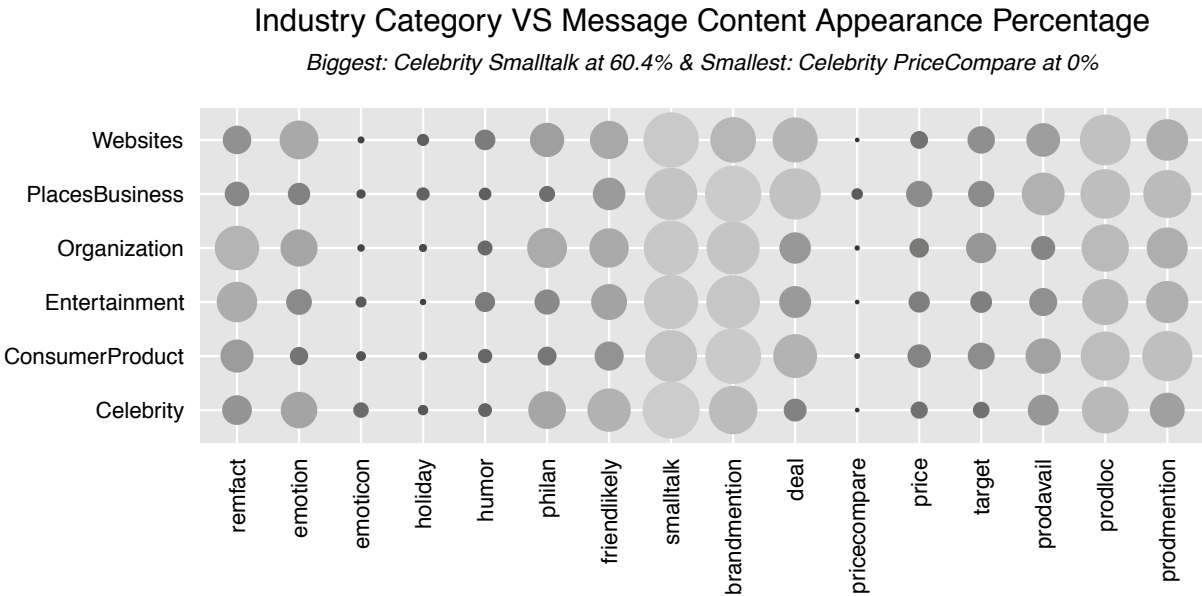


Figure 6: **Bubble Chart of Broader Industry Category vs Message Content:** This chart shows the relative percentage of message contents appearing within industry categories for 5,000 messages. Larger and lighter bubbles imply a higher percentage of messages in that cell. The largest bubble (60.4%) corresponds to SMALLTALK for the celebrity page category and the smallest bubble (0%) corresponds to PRICECOMPARE for the celebrity category.

2.2 Amazon Mechanical Turk

We now describe our methodology for content-coding messages using AMT. AMT is a crowdsourcing marketplace for simple tasks such as data collection, surveys and text analysis. It has now been successfully leveraged in several academic papers for online data collection and classification. To content-code our messages, we create a survey instrument comprising of a set of binary yes/no questions which we pose to workers

(or “Turkers”) on AMT. Please see Appendix 1 for the final survey instrument.

Following best-practices in the literature, we employ the following strategies to improve the quality of classification by the Turkers in our study.

1. For each message, at least 9 different Turkers’ inputs are recorded. We obtain the final classification by a majority-voting rule.
2. We restrict the quality of Turkers included in our study to comprise only those with at least 100 reported completed tasks and 97% or better reported task-approval rates.
3. We use only Turkers from the US so as to filter out those potentially not proficient in English, and to closely match the user-base from our data (recall, our data has been filtered to only include pages located in the US).
4. We refined our survey instrument through an iterative series of about 10 pilot studies, in which we asked Turkers to identify confusing or unclear questions. In each iteration, we asked 10-30 Turkers to identify confusing questions and the reasons they found those questions confusing. We refined the survey in this manner till almost all queried Turkers stated no questions were confusing.
5. To filter out participants who were not paying attention, we included an easily verifiable test question “does the message have a dollar sign (\$)?”. Responses from Turkers that failed the verification test are dropped from the data.
6. In order to incentivize workers, we awarded additional bonuses of \$2-\$5 to the top 20 workers with exceptional accuracy and throughput.
7. On average, we found that message tagging took a little over 3 minutes and it typically took at least 20 seconds or more to completely read the tagging questions. We defined less than 30 seconds to be too short, and discarded any message tags with completion times shorter than that duration to filter out inattentive Turkers and automated programs (“bots”).
8. Once a Turker tags more than 20 messages, a couple of tagged samples are randomly picked and manually examined for quality and performance. This process identified about 20 high-volume Turkers who completed all surveys in less than 10 seconds and tagged several thousands of messages (there were also Turkers who took time to complete the surveys but chose seemingly random answers). We concluded these were automated programs. These results were dropped, and the Turkers “hard blocked” from the survey, via the blocking option provided in AMT.

We believe our methodology for content-classification has good external validity. The binary classification task that we serve to the AMT Turkers in our study is relatively simpler than the more complex tasks for which AMT-based data have been employed successfully in the literature. The existing AMT literature has documented evidence that several of the strategies implemented above improves the quality of the data generated (Mason and Suri (2012); Ipeirotis et al. (2010); Paolacci et al. (2010)). Snow et al. (2008) show that combining results from a few Turkers can produce data equivalent in quality to that of expert labelers




Figure 7: **Cronbach's Alphas for 5,000 Messages**: This bar graph shows the inter-rater reliability measure of Cronbach's Alpha among at least 9 distinct Turkers' inputs for each 5,000 messages. The mean is 0.82 and the median is 0.84. We replicated the study with only those above 0.7 and found the result to be robust.

for a variety of text tagging tasks. Similarly, Sheng et al. (2007) document that repeated labeling of the type we implement wherein each message is tagged by multiple Turkers, is preferable to single labeling in which one person tags one sentence. Finally, evaluating AMT based studies, Buhrmester et al. (2011) concludes that (1) Turkers are demographically more diverse than regular psychometric studies samples, and (2) the data obtained are at least as reliable as those obtained via traditional methods as measured by psychometric standards such as Cronbach's Alpha, a commonly used inter-rater reliability measure. Figure 7 presents the histogram of Cronbach's Alphas obtained for the 5,000 messages. The average Cronbach's Alpha for our 5,000 tagged messages is 0.82 (median 0.84), well above typically acceptable thresholds of 0.7. About 87.5% of the messages obtained an alpha higher than 0.7, and 95.4% higher than 0.6. For robustness, we replicated the study with only those messages with alphas above 0.7 (4,378 messages) and found that our results are qualitatively similar.

At the end of the AMT step, approximately 2,500 distinct Turkers contributed to content-coding 5,000 messages. This constitutes the training dataset for the NLP algorithm used in the next step.

2.3 Natural Language Processing (NLP) for Attribute Tagging

Natural Language Processing is an interdisciplinary field composed of techniques and ideas from computer science, statistics and linguistics for enabling computers to parse, understand, store, and convey information in human language. Some notable applications of NLP are in search engines such as Google, machine translation, and IBM's Watson. As such, there are many techniques and tasks in NLP (c.f., Liu, 2011; Jurafsky and Martin, 2008). For our purposes, we use NLP techniques to label message content from

Facebook posts using the AMT labeled messages as the training data. Typical steps for such labeling tasks include: 1) breaking the sentence into understandable building blocks (e.g., words or lemmas) and identifying different sentence-attributes similar to what humans do when reading; 2) obtaining a set of training sentences with labels tagged from a trusted source identifying whether the sentences do or do not have a given content profile (in our case, this source comprise the 5000 AMT-tagged messages); 3) using statistical tools to infer which sentence-attributes are correlated with content outcomes, thereby learning to identify content in sentences. When presented with a new set of sentences, the algorithm breaks these down to building blocks, identifies sentence-level attributes and assigns labels using the statistical models that were fine-tuned in the training process.

Recent research in the social sciences has leveraged a variety of NLP methods to mine textual data and these techniques have gained traction in business research (see for e.g., Netzer et al. (2012); Archak et al. (2011); Ghose et al. (2012)). Our NLP methods closely mirror cutting edge multi-step methods used in the financial services industry to automatically extract financial information from textual sources (e.g., Hassan et al. (2011)) and are similar in flavor to winning algorithms from the recent Netflix Prize competition.⁴ The method we use combines five statistical classifiers with rule-based methods via heterogeneous “ensemble learning” methods. The statistical classifiers are binary classification machine learning models that take attributes as input and output predicted classification probabilities. The rule-based methods usually use large data sources (a.k.a dictionaries) or use specific *if-then* rules inputted by human experts, to scan through particular words or occurrences of linguistic entities in the messages to generate a classification. Rule-based methods work well for classifying attributes when an exhaustive set of rules and/or dictionaries are available, or if the text length is short as is our case. For example, in identifying brand and product mentions, we augment our AMT-tagged answers with several large lists of brands and products from online sources and a company list database from Thomson Reuters. We then utilize rule-based methods to identify brand and product mentions by looking up these lists. Further, to increase the range of our brand name and product database, we also ran a separate AMT study with 20,000 messages in which we asked AMT Turkers to identify any brand or product name included in the message. We added all the brand and product names we harvested this way to our look-up database. Similarly, in identifying emoticons in the messages, we use large dictionaries of text-based emoticons freely available on the internet.

Finally, we utilize ensemble learning methods that combine classifications from the many classifiers and rule-based algorithms we use. Combining classifiers is very powerful in the NLP domain since a single statistical classifier cannot successfully overcome the classic precision-recall tradeoff inherent in the classification problem.⁵ The final combined classifier has higher precision and recall than any of the constituent classifiers. To the best of our knowledge, the cutting edge multi-step NLP method used in this paper has not been used in business research journals.⁶

⁴See <http://www.netflixprize.com>.

⁵The performance of NLP algorithms are typically assessed on the basis of accuracy (the total % correctly classified), precision (out of predicted positives, how many are actually positive), and recall (out of actual positives, how many are predicted as positives). An important tradeoff in such algorithms is that an increase in precision often causes decrease in recall or vice versa. This tradeoff is similar to the standard bias-variance tradeoff in estimation.

⁶Although there exist business research papers combining statistical classifiers and rule-based algorithms, to our knowledge, none utilize ensemble learning methods which are critical in increasing accuracy, precision, and recall. For example, these methods were a key part of the well-known Netflix-Prize winning algorithms. One of the contributions of this paper is the

For interested readers, the NLP algorithm’s training and classification procedures are described in the following steps. Figure 8 shows the process visually.

Training The Algorithm

1. The raw textual data of 5,000 messages in the training sample are broken down into basic building blocks of sentences using *stop-words removal* (removing punctuation and words with low information such as the definite article “the”), *tokenization* (the process of breaking a sentence into words, phrases, and symbols or “tokens”), *stemming* (the process of reducing inflected words to their root form, e.g., “playing” to “play”), and *part-of-speech tagging* (determining part-of-speech such as nouns). For reference see Jurafsky and Martin (2008). In this process, the input to the algorithm is a regular sentence and the output is an ordered set of fundamental linguistic entities with semantic values. We use a highly regarded python NLP framework named NLTK (Bird et al., 2009) to implement this step.
2. Once the messages are broken down as above, an algorithm extracts sentence-level attributes and sentence-structure rules that help identify the included content. Some examples of sentence-level attributes and rules include: frequent noun words (bag-of-words approach), bigrams, the ratio of part-of-speech used, tf-idf (term-frequency and inverse document frequency) weighted informative word weights, and whether “a specific key-word is present” rule. For completeness, we describe each of these in Table 4. The key to designing a successful NLP algorithm is to figure out what we (humans) do when identifying certain information. For example, what do we notice about the sentences we have identified as having emotional content? We may notice the use of certain types of words, use of exclamation marks, the use of capital letters, etc. At the end of this step, the dataset consists of sentence-level attributes generated as above (the x -variables), corresponding to a series of binary (content present/not-present) content labels generated from AMT (the y -variables).
3. For each binary content label, we then train a classification model by combining statistical and rule-based classifiers. In this step, the NLP algorithm fits the binary content label (the y -variable) using the sentence-level attributes as the x -variables. For example, the algorithm would fit whether or not a message has emotional content as tagged by AMT using the sentence attributes extracted from the message via step 2. We use a variety of different classifiers in this step including logistic regression with L1 regularization (which penalizes the number of attributes and is commonly used for attribute selection for problems with many attributes; see (Hastie et al., 2009)), Naive Bayes (a probabilistic classifier that applies Bayes theorem based on presence or absence of features), and *support vector machines* (a gold-standard algorithm in machine learning that works well for high dimensional problems) with different flavors of regularization and kernels ⁷.
4. To train the ultimate predictive classifier, we use ensemble methods to combine results from the multiple statistical classifiers we fit in step 3. The motivation for ensemble learning is that different classifiers

application of ensemble learning methods, which we believe hold much promise in future social science research based on text data.

⁷We tried support vector machines with L1 and L2 regularization and various kernels including linear, radial basis function, and polynomial kernels. For more details, refer to Hastie et al. (2009).

perform differently based on underlying characteristics of data or have varying precision or recall in different locations of the feature vector space. Thus, combining them will achieve better classification output either by reducing variance (e.g. Bagging (Brieman, 1996)) or reducing bias (e.g. Boosting (Freund and Schapire, 1995)). Please see Xu and Krzyzak (1992); Bennett (2006) for further reading on ensemble methods. This step involves combining the prediction from individual classifiers by weighted-majority voting, unweighted-majority voting, or a more elaborate method called isotonic regression (Zadrozny and Elkan, 2002) and choosing the best performing method in terms of accuracy, precision and recall for each content profiles. In our case, we found that support vector machine based classifiers delivered high precision and low recall, while Naive Bayes based classifiers delivered high recall but low precision. By combining these, we were able to develop an improved classifier that delivers higher precision and recall and in effect, higher accuracy. Table 5 shows the improvement of the final ensemble learning method relative to using only one support vector machine. As shown, the gains from combining classifiers are substantial.

5. Finally, we assess the performance of the overall NLP algorithm on three measures, viz., accuracy, precision, and recall (as defined in Footnote 4) using the “10-fold cross validation” method. Under this strategy, we split the data randomly into 10 equal subsets. One of the subsets is used as the validation sample, and the algorithm trained on the remaining 9 sets. This is repeated 10 times, each time using a different subset as the validation sample, and the performance measures averaged across the 10 runs. The use of 10-fold cross-validation reduces the risk of overfitting and increases the external validity of the NLP algorithm we develop. Note, 10-fold cross-validation of this sort is computationally intensive and impacts performance measures negatively and is not implemented in some existing papers in business research. While the use of 10-fold cross-validation may negatively impact the performance measures, it is necessary to increase external validity. Table 5 shows these metrics for different content profiles. The performance is extremely good and comparable to performance achieved by the leading financial information text mining systems (Hassan et al., 2011).
6. We repeat steps 2-5 until desired performance measures are achieved.

Tagging New Messages

1. For each new messages repeat steps 1-2 described above.
2. Use the ultimate classifier developed above to predict whether a particular type of content is present or not.

One can think of this NLP algorithm as emulating the Turkers’ collective opinion in content-coding.

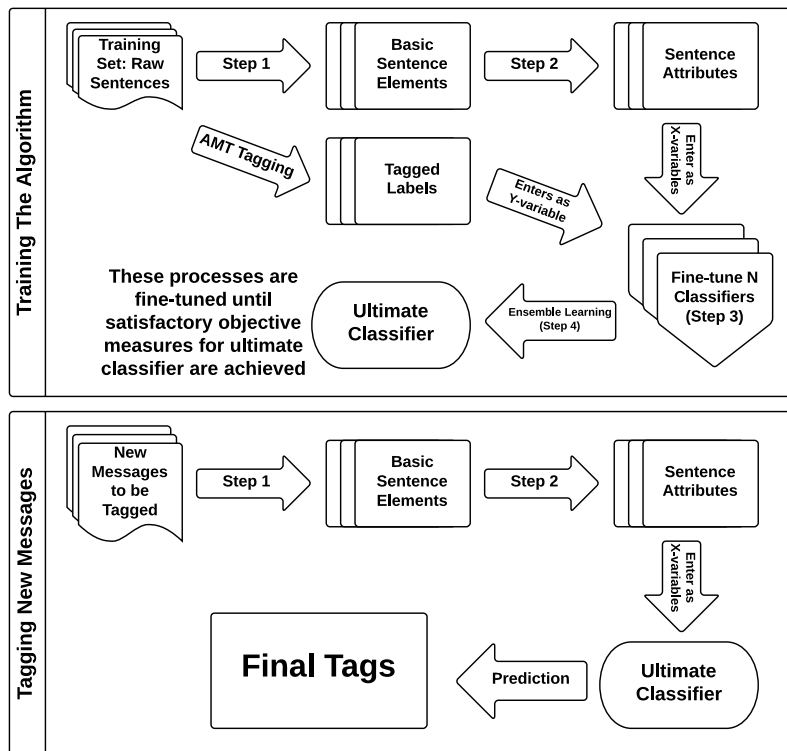


Figure 8: **Diagram of NLP Training and Tagging Procedure:** This diagram shows the steps of training the NLP algorithm and using the algorithm to tag the remaining messages. These steps are described in Section 2.3.

Rules and Attributes	Description
Bag of Words	Collects all the words and frequency for a message. Different variations include collecting top N most occurring words.
Bigram	A bigram is formed by two adjacent words (e.g. “Bigram is”, “is formed” are bigrams).
Ratio of part-of-speech	Part-of-speech (noun, verb, etc) ratio in each message.
TF-IDF weighted informative word	Term-Frequency and Inverse Document Frequency weighs each word based on their occurrence in the entire data and in a single message.
Specific Keywords	Specific keywords for different content can be collected and searched. e.g., Philanthropic messages have high change of containing the words “donate” and “help”. For brand and product identification, large online lists were scraped and converted into dictionaries for checking.
Frequency of different punctuation marks	Counts the number of different punctuations such as exclamation mark and question mark. This helps to identify emotion, questions, appearance of deals etc.
Count of non-alphanumerics	Counts the number of characters that are not A-Z and 0-9.

Table 4: **A Few Examples of Message Attributes Used in Natural Language Processing Algorithm**

	With Ensemble Learning (The Best Performing Algorithm)			Without Ensemble Learning (Support Vector Machine version 1 + Rule-based)		
	Accuracy	Precision	Recall	Accuracy	Precision	Recall
REMACT	0.998	0.998	0.998	0.939	1	0.556
EMOTION	0.996	0.992	0.999	0.951	0.987	0.390
HUMOR	0.999	0.999	1	0.977	1	0.142
PHILANTHROPIC	0.999	0.999	1	0.983	1	0.803
FRIENDLY	0.997	0.996	0.998	0.942	1	0.514
SMALLTALK	0.858	0.884	0.803	0.821	0.960	0.670
DEAL	0.996	0.999	0.994	0.97	1	0.805
PRICECOMPARE	0.999	0.999	1	0.999	1	0.857
TARGETING	0.999	0.998	1	0.966	1	0.540
PRODAVAILABILITY	0.999	0.998	1	0.917	1	0.104
PRODLOCATION	0.970	0.999	0.901	0.939	0.990	0.887

Table 5: **Performance of Text Mining Algorithm on 5000 Messages Using 10-fold Cross Validation:** This table presents metrics for performance of the classification algorithms used. The left 3 columns show the metrics for the final algorithm which combines classifiers via ensemble learning method while the right 3 columns show the metric for a support vector machine algorithm. Notice that the support vector machine classifier tends to have low recall and high precision. Naive Bayes tends to have high recall but low precision. Classifiers on their own cannot successfully overcome precision-recall tradeoff (if one is higher, one is lower). But combining many different classifiers with ensemble learning can increase both precision and recall.

3 Empirical Strategy

Our empirical goal is to investigate the effect of message ad content on subsequent customer engagement. Engagement – the y -variable – is observed in the data; and content – the x -variables – has been tagged as above and is also observed. If messages are randomly allocated to users, the issue of assessing the effect of message-content on engagement is straightforward; one simply projects x on y . Unfortunately, a complication arises because Facebook’s policy of delivery of messages to users is non-random: users more likely to find a post appealing are more likely to see the post in their newsfeed, a filtering implemented via Facebook’s “EdgeRank” algorithm. The filtering implies a selection problem in estimation of the effect of post-characteristics on engagement – if we see that posts with photos are more likely to be commented on by users, we do not know if this is effect of including a photo in a post, or whether Facebook is more likely to show posts with photos to users who are more likely to comment on them. The issue has been ignored in the literature on social media analysis so far. We address the selection issue via a two-step procedure, first by building a semiparametric model of “EdgeRank” that delivers an estimate of the expected number of impressions a post is likely to receive, and then, by incorporating this model to run a selectivity-corrected projection of *Likes* and comments on post characteristics in the second-stage (Blundell and Powell, 2003). For the first-stage, we exploit the fact that we observe the aggregated decisions of Facebook to serve impressions to users, and that “EdgeRank” is based on three variables as revealed by Facebook.

Addressing the problem is complicated by the secretiveness of EdgeRank and by data availability. We know from publicly available documentation that EdgeRank’s assignment of a post to a user is based on the so-called “3 Ts”: Type, Tie, and Time.⁸

⁸As disclosed first at the 2010 “f8” conference. See <http://whatIsEdgeRank.com> for a brief description of EdgeRank.

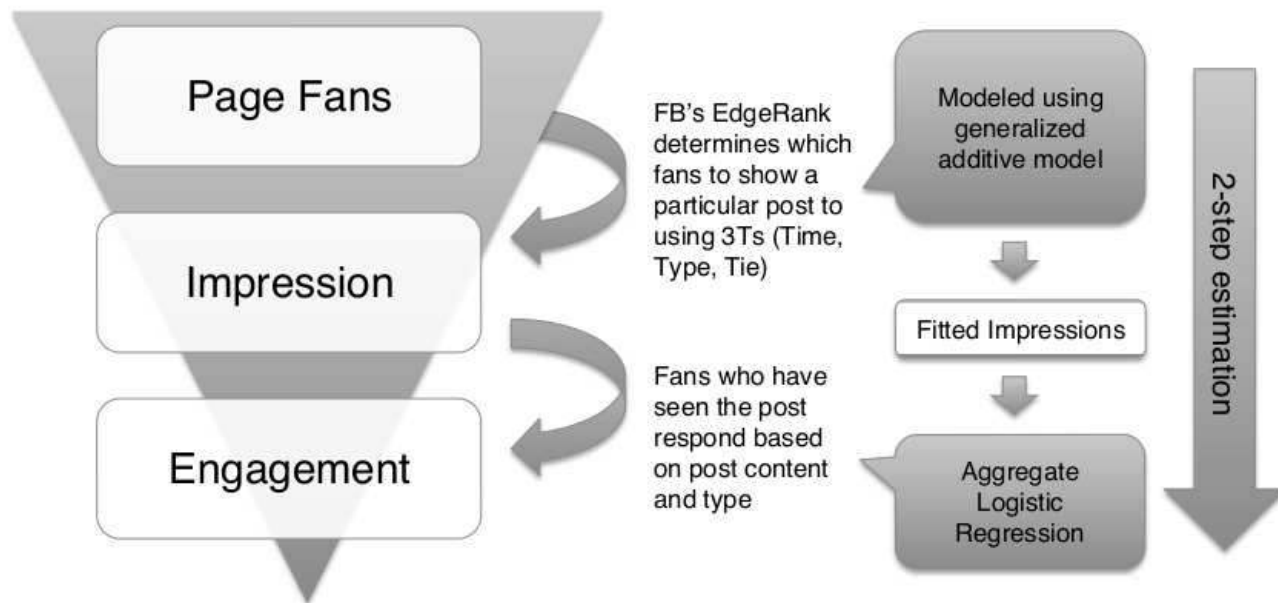


Figure 9: **Impression-Engagement Funnel:** Facebook’s EdgeRank chooses subset of Page fans to show posts released by the page and fans who’ve seen the post engage with the post based on content and type. EdgeRank is modeled with generalized additive model and the final engagement is estimated through aggregate logistic regression. Details of estimation are in Sections 3.1 and 3.2.

- *Type* (z) refers to the type of post. Facebook categorizes post-type into 5 classes: status update, photo, video, app, or link.
- *Tie* (h_{ijt}) refers to the affinity score between page j (company) and the Facebook user i (viewer of the post) at time t which is based on the strength and frequency of the interaction history between the user and the page.
- *Time* (τ) refers to the time since the post.

Our dataset contains direct observations on the variables *Type* and *Time*. We do not have individual-level data on a user’s history with pages to model tie strengths. However, we exploit the fact that we have access to demographics data on the set of users who *could potentially have been shown* a post released by a firm, versus *who were actually shown* the post. The difference reflects the selection by EdgeRank, which we utilize as a proxy measure of Tie-strength based targeting. Since we do not know the exact functional form of EdgeRank’s targeting rule, we work with a semiparametric specification, utilizing flexible splines to capture the effect of EdgeRank. At the end of this step, we thus develop a flexible approximation to EdgeRank’s targeting. In the second step, we can then measure the effect of ad content on *Likes* and comments, by controlling for the non-random targeting using our first-stage model. Figure (9) shows the empirical strategy visually. The econometrics below sets up estimation using the aggregate post-level panel data split by demographics that we observe, while acknowledging the fact that non-random targeting is occurring at the individual-level.

3.1 First-stage: Approximating EdgeRank’s Assignment

We represent post k ’s type in a vector z_k , the time since post k was released in τ_k , and the history of user i ’s past engagement with company j on Facebook in a vector h_{ijt} . Table 6 summarizes the notation.

To understand our procedure, let $n_{kjt}^{(d)}$ denote the number of users of demographic type $d = 1, \dots, D$ who were shown post k by firm j at time t . We refer to $n_{kjt}^{(d)}$ as impressions. $n_{kjt}^{(d)}$ is indirectly reported in the data and can be reverse-engineered from Company X’s reports. A description of this procedure is provided in Appendix 2. Let $N_{jt}^{(d)}$ denote the total number of users of demographic type d for firm j on day t to whom the post can potentially be delivered. $N_{jt}^{(d)}$ is directly observed in the data, and comprises all users of demographics d who have *Liked* the firm on Facebook. To be clear, note that *Liking* a post is different from *Liking* a page – *Liking* a page provides the firm that maintains that page an opportunity to serve its posts to that user via Facebook’s Newsfeed. $N_{jt}^{(d)}$ is a count of all such users.

Now, note that by EdgeRank’s assignment rule, the aggregated impressions for demographic type d , $n_{kjt}^{(d)}$, is an (unknown) function of liked-fans $N_{jt}^{(d)}$, the tie strength between users within demographic bucket d and the posting firm, $h_{ijt}^{(d)}$, the type of post z_k , and time since post release τ_k ,

$$\mathbb{E}(n_{kjt}^{(d)}) = g(N_{jt}^{(d)}, h_{ijt}^{(d)}, z_k, \tau_k) \tag{1}$$

We do not observe individual-level data on each users i ’s interaction with every post which could be the basis of estimating Equation (1). Instead, we can construct the aggregated number of impressions and liked-fans within a set of demographic buckets in the data. To use this variation as a source of approximating EdgeRank, we approximate the RHS of Equation (1) as,

$$\mathbb{E}(n_{kjt}^{(d)}) \approx g_d(N_{jt}^{(d)}, \theta_{1j}^{(d)}, z_k, \tau_k) \tag{2}$$

where, we use a firm-demographic bin specific fixed effect, $\theta_{1j}^{(d)}$, to capture the effect of user history. This approximation would literally be true if all individuals within demographic bucket d had the same history with firm j . In practice, this is not the case, and this may induce approximation errors into the procedure, because additional history-heterogeneity within demographic buckets is not modeled (or is assumed into the

<i>Notation</i>	<i>Description</i>
i	User
j	Firm
k	Post
t	Time (day)
z_k	post k ’s media type (5 options: photo, video, status update, app, link)
τ_k	Time since post k was released
h_{ijt}	History of user i ’s past engagement with firm j
$g(\cdot)$	EdgeRank score approximating function
$n_{kjt}^{(d)}$	Impressions of post k by page j at time t by users in demographics bin d
$N_{jt}^{(d)}$	Number of users of demographics bin d who <i>Liked</i> page j as of time t
$\theta_0^{(d)}$	Intercept term for each demographics d
$\theta^{(d)}$	Parameters in EdgeRank approximation for demographics bin d

Table 6: User-level Setup Notation

error term). This is a caveat to our analysis. Access to individual-level data could be the basis of improving this procedure and relaxing this assumption. We view Equation (2) as a flexible approximation that allows us to leverage the observed variation in firm-level impressions across demographics, while enabling us to include firm and demographic-level fixed effects into a procedure that best approximates EdgeRank based on what we as researchers (and firms) know about Facebook’s filtering algorithm. We will also estimate the right-hand function $g_d(\cdot)$ separately for each demographic bucket, in effect allowing for slope heterogeneity in demographics in addition to intercept heterogeneity across demographics.

The next step relates to approximating the function $g_d(\cdot)$. Since we do not know the exact functional form of the above selection equation, we approximate the function semiparametrically via a Generalized Additive Model (GAM) (c.f., Hastie and Tibshirani (1990)). The GAM is a generalized linear model with additive predictors consisting of smoothed (e.g. interpolation and curve fitting) covariates. The GAM fits the following flexible relationship between a set of covariates X and dependent variable Y ,

$$\mu(\mathbb{E}(Y|X_1, X_2, \dots, X_p)) = \alpha + s_1(X_1) + s_2(X_2) + \dots + s_p(X_p)$$

where μ is a link function (e.g. gaussian, poisson, gamma), and s_1, s_2, \dots, s_p are nonparametric smoothing functions such as cubic splines or kernel smoothers. We model the EdgeRank selection equation for each demographic d as the following,

$$\begin{aligned} h_d \left[\log(n_{kjt}^{(d)} + 1) \right] &= \theta_0^{(d)} + \theta_{1j}^{(d)} + \theta_2^{(d)} \mathbb{N}_{jt}^{(d)} + s_1(\mathbb{N}_{jt}^{(d)}; \theta_3^{(d)}) + \sum_{r=2}^5 \theta_{4r}^{(d)} \mathbb{I}(z_k = r) \\ &+ \sum_{r=2}^{16} \theta_{5r}^{(d)} I(\tau_k = r) + \epsilon_{kjt}^{(d)} \end{aligned} \quad (3)$$

where, $h_d \equiv g_d^{-1}(\cdot)$ is the identity (Gaussian) link function, $\theta_0^{(d)}$ is an intercept term unique to each demographic, d , and $\theta_{1j}^{(d)}$ is a firm-demographic fixed effect that captures the tie strength between the firm j and demographics d .⁹ $\mathbb{N}_{jt}^{(d)}$ is the number of fans of demographic d for firm j at time t and denotes the potential audience for a post. s_1 is a cubic spline smoothing function, essentially a piecewise-defined function consisting of many cubic polynomials joined together at regular intervals of the domain such that the fitted curve, the first and second derivatives are continuous. We represent the interpolating function $s_1(\cdot)$ as a linear combination of a set of basis functions $b(\cdot)$ and write: $s_1(\mathbb{N}_{jt}^{(d)}; \theta_3^{(d)}) = \sum_{r=3}^q b_r(\mathbb{N}_{jt}^{(d)}) \theta_{3r}^{(d)}$, where the $b_r(\cdot)$ are a set of basis functions of dimension q to be chosen and $\theta_{3r}^{(d)}$ are a set of parameters to be estimated. We follow a standard method of generating basis functions, $b_r(\cdot)$, for the cubic spline interpolation as defined in Wood (2006). Fitting the spline also requires choosing a smoothing parameter, which we tune via generalized cross-validation. We fit all models via the R package `mgcv` described in Wood (2006).

Finally, we include dummy variables for post-type (z_k) and for each day since release of the post (τ_k ; up to 16 days), to capture the effect of post-type and time-since-release semiparametrically. These are allowed to be d -specific. We collect the set of parameters to be estimated for each demographic bucket in a vector,

⁹We also tried Poisson and Negative Binomial link functions (since $n_{kjt}^{(d)}$ is a count variable), as well as the identity link function without logging the y -variable. Across these specifications, we found the identity link function with $\log(y)$ resulted in the best fit, possibly due to many outliers. We also considered specifications with numerous interaction of the covariates included, but found they were either not significant or provided trivial gains in the R^2 .

$\theta^{(d)}$, which we estimate by GAM estimation. The estimated parameter vector, denoted $\hat{\theta}^{(d)}$, $d = 1, \dots, D$, serves as an input to the second stage of the estimation procedure.

3.2 Second-stage: Modeling Engagement given Post-Assignment

We operationalize engagement via two actions, *Likes* and comments on the post. The selection problem was that users can choose to *Like* or comment on a post only if they were served impressions, which generates non-random censoring because impression assignment was endogenous to the action. We address the censoring by including a correction for the fact that a user was shown a post non-randomly, estimated semiparametrically as above. Suppose $\hat{\Psi}_{kjt}^{(d)}$ denotes the fitted estimate from the first-stage of the expected number of impressions of post k for firm j amongst users of type d at time t ,

$$\hat{\Psi}_{kjt}^{(d)} = g_d \left(N_{jt}^{(d)}, z_k, \tau_k; \hat{\theta}^{(d)} \right)$$

For future reference, note the expected number of impressions of post k for firm j at time t across all demographic buckets is simply the sum,

$$\hat{\Psi}_{kjt} = \sum_{d=1}^D g_d \left(N_{jt}^{(d)}, z_k, \tau_k; \theta^{(d)} \right)$$

Now, we let the probability that users will *Like* a post given the full set of post characteristics and auxiliary controls, M_{kt} , be logistic with parameters ψ ,

$$\pi(M_{kt}; \psi) = \frac{1}{1 + e^{-M_{kt}\psi}} \quad (4)$$

The parameter vector, ψ , is the object of inference in the second stage.¹⁰ We observe Q_{kjt} , the number of *Likes* of the post in each period in the data. To see the intuition for our correction, note that we can aggregate Equation (4) across users, so that the expected number of *Likes* is,

$$\mathbb{E}(Q_{kjt}) \approx \sum_{d=1}^D \hat{\Psi}_{kjt}^{(d)} \times \left[\frac{1}{1 + e^{-M_{kt}\psi}} \right] \quad (5)$$

with $\hat{\Psi}_{kjt}^{(d)}$ are treated as known. The right-hand side is a weighted sum of logit probabilities of *Liking* a post. Intuitively, the decision to *Like* a post is observed by the researcher only for a subset of users who were endogenously assigned an impression by FB. The selection functions $\hat{\Psi}_{kjt}^{(d)}$ serve as weights that reweigh the probability of *Liking* to account for the fact that those users were endogenously sampled, thereby correcting for the non-random nature of post assignment when estimating the outcome equation.

We could use the expectation in Equation (5) as the basis of an estimation equation. Instead, for efficiency, we estimate the parameter vector ψ by maximum likelihood. We specify the probability that Q_{kjt} out of the $\hat{\Psi}_{kjt}$ assigned impressions are observed to *Like* the post, and that $\hat{\Psi}_{kjt} - Q_{kjt}$ of the remaining impressions are observed not to, is binomial with probability, $\pi(M_{kt}; \psi)$,

$$Q_{kjt} \sim \text{Binomial}(\hat{\Psi}_{kjt}, \pi(M_{kt}; \psi)) \quad (6)$$

¹⁰Allowing ψ to be d -specific in Equation (4) is conceptually straightforward. Unfortunately, we do not have *Likes* or comments split by demographics in order to implement this.

Maximizing the implied binomial likelihood across all the data, treating $\hat{\Psi}_{kjt}$ as given, then delivers estimates of ψ . The intuition for the selection correction here is the same as that encapsulated in Equation (5). We can repeat the same procedure using the number of comments on the post as the dependent variable so as to recover the effect of post-characteristics on commenting as well. This two-step procedure thus delivers estimates of the causal effects of post-characteristics on the two outcomes of interest.

Discussion of Identification Identification in the model derives from two sources. First, we exploit the observed discrepancy in demographic distributions between the set of individuals to whom a post could have been served, versus those who were actually served. The discrepancy reflects the filtering by EdgeRank. Our first stage essentially projects this discrepancy onto post-type, time-since-release, page and demographic characteristics in a flexible way. This essentially serves as a “quasi” control function that corrects for the selectivity in the second stage (Blundell and Powell, 2003), where we measure the effect of post characteristics on outcomes. The second source of identification arises from exploiting the implied exclusion restriction that the rich set of AMT-content-coded attributes affect actual engagement, but are not directly used by EdgeRank to assign posts to users. The only post-characteristics used by EdgeRank for assignment is z_k , which is controlled for. Thus, any systematic correlation in outcomes with AMT-content-coded characteristics, holding z_k fixed, do not reflect selection-related considerations.

4 Results

4.1 First-Stage

The first-stage model, as specified in Equation 3, approximates EdgeRank’s post assignment algorithm. We run the model separately for each of the 14 age-gender bins used by Facebook. These correspond to two gender and seven age bins. For a given bin, the model relates the number of users of demographic type d who were shown post k by firm j at time t to the post type (z_k), days since post (τ) and tie between the firm and the user. Table 7 presents the results. The intercepts ($\theta_0^{(d)}$) indicate that posts by companies in our dataset are shown most often to Females ages 35-44, Females 45-54 and Males 25-34. The lowest number of impressions are for the 65+ age group. In our model, tie between a user and a firm is proxied by a fixed-effect for each firm-demographic pair. This implies 800×14 fixed effects corresponding to 800 firms and 14 demographic bins. Due to space constraints, we do not present all the estimated coefficients. Table 7 presents the coefficients for two randomly chosen firms. The first is a new-born clothing brand and the second is a protein bar brand. For ease of visualization, these fixed effects are shown graphically in Figure 10 (only the statistically significant coefficients are plotted). For posts by the the new-born clothing brand, the most impressions are among from females in the age-groups of 25-34, 18-24 and 35-44. Among males, ages 25-34 receive the most number of impressions. For posts by the protein bar brand, impressions are more evenly distributed across the different demographic bins, with the Male 18-24 group receiving the most impressions. These estimated coefficients are consistent with our expectations for the two brands.

Female							
	F 13-17	F 18-24	F 25-34	F 35-44	F 45-54	F 55-64	F 65+
Intercept	5.528***	6.071***	6.446***	7.165***	7.209***	6.133***	4.887***
Page 1 fixed effect - new born clothing brand	-0.210	2.458***	2.685***	1.544**	0.888	0.813	0.489
Page 2 fixed effect - protein bar brand	-0.573***	1.285***	1.466***	0.928***	0.016	1.671***	1.518***
Message Type - App is the base							
Link	0.010	0.045***	0.063***	0.042***	0.051***	0.051***	0.048***
Photo	0.253***	0.318***	0.340***	0.309***	0.297***	0.267***	0.249***
Status Update	0.100***	0.161***	0.175***	0.152***	0.152***	0.129***	0.114***
Video	0.033	0.041	0.061**	0.041	0.021	0.024	0.030
$N_{jt}^{(d)}$ (Fan Number)	2.0×10^{-6} ***	1.8×10^{-6} ***	7.2×10^{-6} ***	1.9×10^{-5} ***	1.9×10^{-5} ***	3.8×10^{-5} ***	8.5×10^{-5} ***
$s(N_{jt}^{(d)})$ significance	***	***	***	***	***	***	***
R-Squared	0.78	0.78	0.77	0.78	0.78	0.78	0.77
Male							
	M 13-17	M 18-24	M 25-34	M 35-44	M 45-54	M 55-64	M 65+
Intercept	5.486***	6.118***	7.075***	6.635***	6.125***	5.151***	4.011***
Page 1 fixed effect - new born clothing brand	0.156	0.932	1.673**	1.082	0.722	0.209	0.111
Page 2 fixed effect - protein bar brand	1.867***	2.423***	0.907***	0.670***	1.158***	1.575***	1.502***
Message Type - App is the base							
Link	-0.005	0.025***	0.033***	0.034***	0.038***	0.049***	0.030***
Photo	0.226***	0.284***	0.295***	0.277***	0.254***	0.230***	0.212***
Status Update	0.077***	0.124***	0.126***	0.120***	0.106***	0.103***	0.084***
Video	0.014	0.039	0.044*	0.031	0.016	0.007	0.023
$N_{jt}^{(d)}$ (Fan Number)	3.6×10^{-6} ***	1.0×10^{-6} ***	6.7×10^{-6} ***	2.5×10^{-5} ***	3.8×10^{-5} ***	5.2×10^{-5} ***	2.3×10^{-4} ***
$s(N_{jt}^{(d)})$ significance	***	***	***	***	***	***	***
R-Squared	0.79	0.80	0.79	0.78	0.78	0.77	0.76
*App is the base for message type. Significance Level: '****' < 0.001 '***' < 0.01 '**' 0.05							

Table 7: **EdgeRank Model Estimates:** This table presents the coefficients obtained from 14 generalized additive models for EdgeRank, calculated for each demographic bin. There are 14 demographic (gender-age) bins provided by Facebook. F13-17 means all females in the age between 13 and 17. Time since post (τ), and page-level fixed effects are not included in the table and presented graphically separately.

The estimates for message type are roughly the same in all demographic bins. For all demographics, the photo type has the highest coefficient (around 0.25) suggesting that photos are preferred to all other media types by EdgeRank. This is likely because users have historically engaged better with photos causing Facebook to show photos more often. The next most preferred post type is the status update with coefficients averaging around 0.12 followed by videos and links. The baseline post type, apps, is the message type that is least preferred by EdgeRank. The rank ordering of coefficients for message type do not strictly follow the rank ordering of number of posts released by firms, which is shown in Table 2. Whereas links are posted more often, photos get more impressions relative to posts of other types, clearly highlighting the role of EdgeRank. Days since post (τ) are not presented in Table 7 due to space constraints. However, Figure 11 presents a box plot of the coefficients for τ across all 14 demographic bins. All coefficients are negative and significant

and also more negative for higher values of τ , implying that EdgeRank prefers to show more recent posts. Finally, the coefficients for number of fans, $N_{jt}^{(d)}$, are positive and significant but they have relatively low magnitude. This is because our model includes a smoothed term of the number of fans, $s(N_{jt}^{(d)})$, which soaks up both the magnitude and nonlinearity. The smoothed fan-numbers are all significant.

The generalized additive model of EdgeRank recovers coefficients that make intuitive sense and are consistent with claims made in several industry reports (e.g. that photos have the highest EdgeRank weight). Further, the model fit appears to be good especially given that we have used generalized cross validation to guard against overfitting.




Figure 10: **Page-level Fixed effect Estimates from Generalized Additive Model Across 14 Demographic Bins:** This bar graph shows two randomly chosen page-level fixed effect estimates from the EdgeRank models. Only the statistically significant estimates are shown. New born clothing brands are positively significant for 18-24 female, 25-34 female, 35-44 female and 25-34 male. Protein bar brands have the highest fixed effect among 18-24 male demographics.




Figure 11: **Time Since Post Release (τ) Coefficients Box plot Across Demographics:** This box plot shows the coefficients on τ across all the demographics bin. $\tau = 1$ is the base case and every coefficients are significant at the highest level of $p < 0.001$.

Variable	Comment	Like
Constant	-6.913***(0.002)	-4.671***(0.001)
Persuasive	0.053***(0.001)	0.061***(0.000)
Informative	-0.143***(0.001)	-0.068***(0.000)
Persuasive \times Informative	0.012***(0.000)	0.003***(0.000)
McFadden R-sq.	0.015	0.009
Nagelkerke R-sq.	0.015	0.009
Log-likelihood	-4208220.431	-33678695.014
Deviance	8012471.987	66409947.187
AIC	8416448.861	67357398.028
N	665916	665916
Significance	'***' 0.001 '***' 0.01 '*' 0.05 '.' 0.1	

Table 8: **Persuasive vs Informative:** Logistic regression for {Comment, Like} with composite summary variables for persuasive and informative content.

4.2 Second-Stage

In the second-stage, we measure the effect of content characteristics on engagement using our selectivity-corrected model from the first-stage. All results in this section are based on an analysis of the entire set of over 100,000 messages (i.e. the 5000 AMT-tagged messages as well as the messages tagged using NLP). The results for only the 5,000 AMT-tagged messages are qualitatively similar and are presented in the appendix. To present the results in a simple way, we first create two composite summary variables corresponding to persuasive content and informative content. Persuasive (informative) composite variables are created by adding up the content variables categorized as persuasive (informative) in Table 2. To be clear, the persuasive variable is obtained by adding values of REMFACT, EMOTION, EMOTICON, HOLIDAYMENTION, HUMOR, PHILANTHROPIC, FRIENDLIKELY, and SMALLTALK resulting in a composite variable ranging from 0 to 8. The informative composite variable is obtained by adding values of BRANDMENTION, DEAL, PRICECOMPARE, PRICE, TARGET, PRODAVAIL, PRODLOCATION, and PRODMENTION resulting in a composite variable ranging from 0 to 8. Table 8 shows the result of logistic regression on engagement with these composite variables and interaction of those two variables as the x -s.

We find that persuasive content has a positive and statistically significant effect on both types of engagement; further, informative content reduces engagement. Interestingly, the interaction between persuasive and informative content is positive, implying that informative content increases engagement only in the presence of persuasive content in the message. This suggests that mixing persuasive and informative content should be made a basis of content engineering for improving engagement with consumers on this medium.

Table 9 presents the results of aggregate logistic regression with the full list of content variables. We present results for both engagement metrics (*Likes*/comments) as well as for models with and without the EdgeRank correction. We exclude the 16 estimated τ coefficients from the table since they are all negative and statistically significant just as in the EdgeRank model in Figure 11. Scanning through the results, we observe that the estimates are directionally similar, in most cases, with and without EdgeRank correction.

However, the magnitudes often change. For example, consider the coefficients for message type Photo. In the model without EdgeRank correction, Photos are very likely to get comments (coefficient = 0.844) and *Likes* (coefficient = 1.023). After EdgeRank correction, the results are similar but the magnitude of the effect drops. This makes sense because we know that EdgeRank prefers Photos. Similarly, Status Updates continue to be more likely (than apps) to get comments and *Likes* but the effect size is smaller after EdgeRank correction. In some instances, there are directional changes for some coefficients. For example, the result that links are more likely to get *Likes*/comments relative to apps changes sign after EdgeRank correction. This highlights the importance of EdgeRank correction. Several industry reports (e.g., Wildfire 2012) often evaluate user content preference without accounting for EdgeRank and we clearly find that the conclusions may often be changed (or sometimes even reversed) after EdgeRank correction. For example, most industry reports' ordering of engaging media type often list status update to be more engaging than videos. While we find this to be true before EdgeRank correction for *Likes*, we find that this is reversed after the EdgeRank correction.

We find that high reading complexity (SMOG) decreases both *Likes* and comments whereas shorter messages (MSGLEN) are *Liked* and commented on more, albeit with a small effect size. Having links (HTTP) is worse for engagement whereas asking questions (QUESTION) significantly increase comments but at the cost of *Likes*. Using blanks in the post to encourage comments has a similar effect of increasing comments but hurting *Likes*. Interestingly, while the odds ratio of comments increases by 75% if a post asks a question, it increases by 214% if blanks are included suggesting that blanks are more effective than questions if the goal is to increase comments. Asking for *Likes* increase both *Likes* and comments, whereas asking for comments increase comments but at the cost of *Likes*. It is clear that even these simple content variables impact user engagement.

The next 16 variables in the table are the persuasive and informative content variables. Figure 12 charts the coefficients for these variables in a bar graph and demonstrates the sharp difference between persuasive and informative content types. Looking at comments, a striking pattern is that most informative contents have a negative impact whereas persuasive contents have a positive impact. The informative content variables with the most negative impact are PRICE, DEAL and PRODMENTION. The persuasive content variables with the most positive impact are EMOTION and PHILANTHROPIC. Interestingly, HOLIDAYMENTION discourages comments.¹¹ One possible explanation is that near holidays, all Facebook pages indiscriminately mention holidays, leading to a dulled responses. For example, during Easter, the occurrence of holiday mention jumped to nearly 40% across all posts (of our data) released that day compared to the average occurrence of about 1%. Looking at *Likes*, fewer persuasive content variables have positive impact but the results are qualitatively similar to that for comments. Among persuasive contents, EMOTION has the most positive impact on *Likes* whereas HOLIDAYMENTION has the most negative impact. Most informative content variables continue to have a negative impact, with PRICE and DEAL having the most negative impact. The results also highlight that there exist some differences between impact on *Likes* versus Comments.

Figure 13 shows the results on content effects by industry. Only the statistically significant results are

¹¹We checked for correlation with other contents to investigate this matter but no correlation was over 0.02.

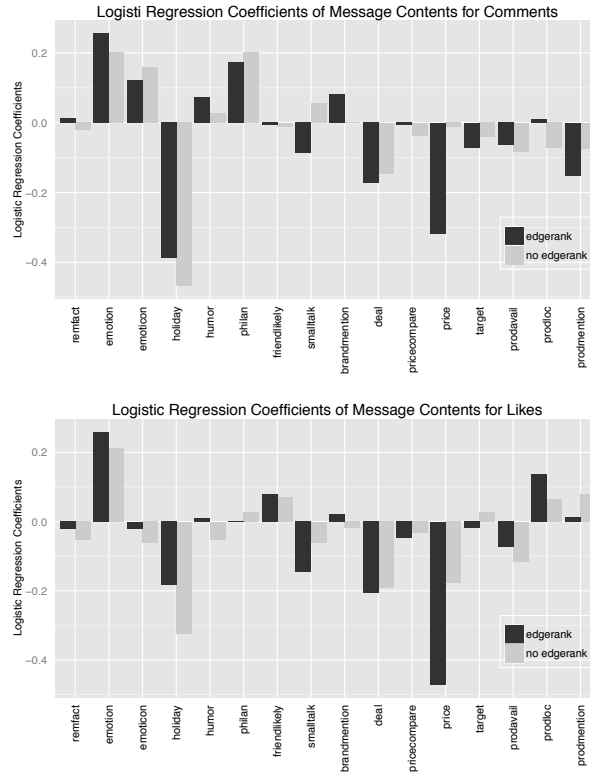


Figure 12: **Post Characteristic Coefficients for Comments and Likes:** These bar graphs show the coefficients of logistic regression for both EdgeRank corrected and uncorrected models. Only the significant coefficients are plotted.

graphed and all results are EdgeRank-corrected. The coefficients are very different across industries both in magnitude and, for some variables, in direction. For example, emotional and philanthropic content has the most positive impact on Facebook pages of type “Organizations” which include non-profits, educational organizations and religious groups. Further, while mentioning holidays has a negative impact on engagement for most industry types, it has a positive impact on engagement for Organizations. Similarly, looking at informative contents, we observe that variables such as Price, Product Availability and Product Mention generally have a negative impact on engagement for most industry types but they have a positive impact for industry type “Celebrity.” Users seem to more forgiving of celebrity pages endorsing products and sharing price information. Similarly, the message type coefficients also vary by industry. Coefficients for message type Link are negative for celebrity, consumer product, and entertainment pages whereas they are positive for organization, places and business, and websites.

Overall, we find that persuasive content engages users better than informative content but their effectiveness varies across industries. Results from alternative model specifications as well as for our original specification applied to only the set of 5000 AMT-tagged messages, which are shared in the appendix, show that the main results are robust across these different specifications and datasets.

	NO ER COMMENT	OR	ER COMMENT	OR	NO ER LIKE	OR	ER LIKE	OR
Constant	-8.232***(0.004)	0.000	-6.889***(0.004)	0.001	-5.206***(0.001)	0.005	-3.892***(0.001)	0.020
SMOG	-0.042***(0.000)	0.959	-0.068***(0.000)	0.934	-0.029***(0.000)	0.971	-0.061***(0.000)	0.941
MSGLEN	0.000***(0.000)	1.000	-0.000***(0.000)	1.000	-0.000***(0.000)	1.000	-0.000***(0.000)	1.000
HTTP	-0.545***(0.002)	0.580	-0.355***(0.002)	0.701	-0.388***(0.000)	0.678	-0.189***(0.000)	0.828
QUESTION	0.497***(0.001)	1.644	0.564***(0.001)	1.758	-0.284***(0.000)	0.753	-0.175***(0.000)	0.839
BLANK	1.020***(0.003)	2.773	1.146***(0.003)	3.146	-0.701***(0.002)	0.496	-0.596***(0.002)	0.551
ASKLIKE	0.101***(0.010)	1.106	0.221***(0.010)	1.247	0.502***(0.003)	1.652	0.565***(0.003)	1.759
ASKCOMMENT	0.502***(0.021)	1.652	0.469***(0.021)	1.598	-0.252***(0.011)	0.777	-0.465***(0.011)	0.628
Persuasive								
REMACT	-0.021***(0.002)	0.979	0.014***(0.002)	1.014	-0.054***(0.001)	0.947	-0.021***(0.001)	0.979
EMOTION	0.203***(0.002)	1.225	0.256***(0.002)	1.292	0.213***(0.001)	1.237	0.260***(0.001)	1.297
EMOTICON	0.159***(0.004)	1.172	0.121***(0.004)	1.129	-0.062***(0.001)	0.940	-0.020***(0.001)	0.980
HOLIDAYMENTION	-0.468***(0.014)	0.626	-0.388***(0.014)	0.678	-0.323***(0.004)	0.724	-0.183***(0.004)	0.833
HUMOR	0.028***(0.002)	1.028	0.072***(0.002)	1.075	-0.052***(0.000)	0.949	0.009***(0.000)	1.009
PHILANTHROPIC	0.202***(0.002)	1.224	0.174***(0.002)	1.190	0.028***(0.001)	1.028	0.002*(0.001)	1.002
FRIENDLY	-0.011***(0.002)	0.989	-0.006***(0.002)	0.994	0.070***(0.001)	1.073	0.080***(0.001)	1.083
SMALLTALK	0.057***(0.002)	1.059	-0.086***(0.002)	0.918	-0.060***(0.001)	0.942	-0.146***(0.001)	0.864
Informative								
BRANDMENTION	0.001(0.002)	1.001	0.081***(0.002)	1.084	-0.018***(0.000)	0.982	0.021***(0.000)	1.021
DEAL	-0.146***(0.002)	0.864	-0.172***(0.002)	0.842	-0.192***(0.001)	0.825	-0.207***(0.001)	0.813
PRICECOMPARE	-0.036***(0.001)	0.965	-0.006***(0.001)	0.994	-0.032***(0.000)	0.969	-0.047***(0.000)	0.954
PRICE	-0.013***(0.005)	0.990	-0.317***(0.005)	0.728	-0.178***(0.001)	0.837	-0.471***(0.001)	0.624
TARGET	-0.041***(0.002)	0.960	-0.071***(0.002)	0.931	0.027***(0.001)	1.027	-0.019***(0.001)	0.981
PRODAVAIL	-0.083***(0.002)	0.920	-0.064***(0.002)	0.938	-0.115***(0.001)	0.891	-0.073***(0.001)	0.930
PRODLOCATION	-0.072***(0.002)	0.931	0.011***(0.002)	1.011	0.064***(0.001)	1.066	0.138***(0.001)	1.148
PRODMENTION	-0.074***(0.002)	0.929	-0.151***(0.002)	0.860	0.080***(0.001)	1.083	0.012***(0.001)	1.012
Message Type - App is the base								
-Link	0.221***(0.003)	1.247	-0.370***(0.003)	0.691	0.125***(0.001)	1.133	-0.505***(0.001)	0.604
-Photo	0.844***(0.003)	2.326	0.373***(0.003)	1.452	1.023***(0.001)	2.782	0.561***(0.001)	1.752
-Status Update	1.117***(0.003)	3.056	0.641***(0.003)	1.898	0.462***(0.001)	1.587	-0.083***(0.001)	0.920
-Video	0.042***(0.009)	1.043	0.398***(0.009)	1.489	-0.061***(0.003)	0.941	0.231***(0.003)	1.260
Industry Category - Celebrity is the base								
-ConsumerProduct	0.177***(0.002)	1.194	-0.347***(0.002)	0.707	-0.409***(0.001)	0.664	-0.932***(0.001)	0.394
-Entertainment	0.470***(0.002)	1.600	0.529***(0.002)	1.697	-0.250***(0.001)	0.779	-0.193***(0.001)	0.824
-Organization	0.598***(0.002)	1.818	0.408***(0.002)	1.504	0.028***(0.001)	1.028	-0.171***(0.001)	0.843
-PlaceBusiness	0.500***(0.005)	1.649	-0.021***(0.005)	0.979	-0.685***(0.002)	0.504	-1.275***(0.002)	0.279
-Websites	0.307***(0.003)	1.359	0.182***(0.003)	1.200	0.214***(0.001)	1.239	0.041***(0.001)	1.042
McFadden R-sq.	0.223		0.171		0.296		0.201	
Nagelkerke R-sq.	0.223		0.172		0.297		0.203	
Log-likelihood	-2594461.972		-3570184.37		-14596455.538		-27221231.997	
Deviance	4784243.846		6736399.867		28244853.214		53495021.154	
AIC	5189021.945		7140466.74		29193009.075		54442561.994	
N	665916		665916		665916		665916	
Significance **** 0.001 *** 0.01 ** 0.05 * . 0.1								

Table 9: **Aggregate Logistic Regression Results For Comments and Likes:** This table presents the aggregate logistic regression on comments and *Likes* for both EdgeRank-corrected (ER) and uncorrected (NO ER) for all data. OR means Odds ratio and shows the odds ratio for the estimates left of the column.

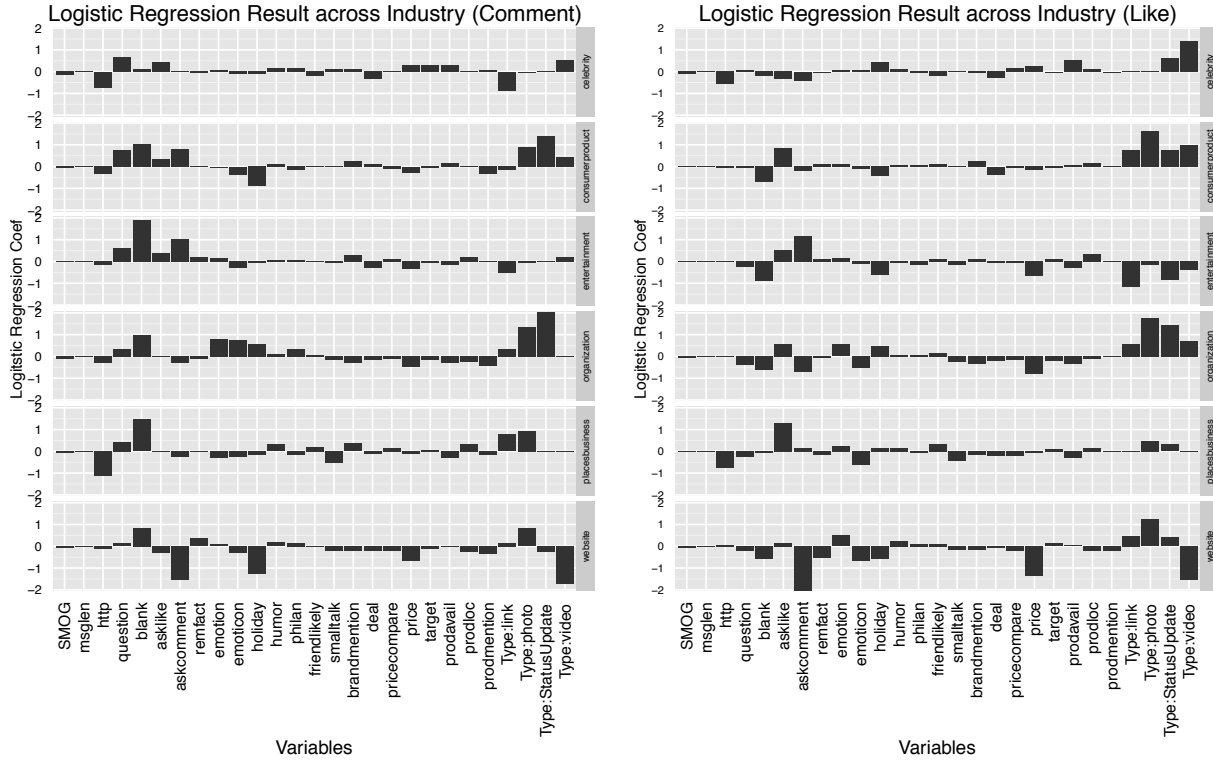


Figure 13: **Logistic Regression by Industry (Comments and Likes)**: This bar graphs show the coefficients of logistic regression for EdgeRank-corrected model. Only the significant (at $p < 0.05$ - but most are $p < 0.001$) coefficients are graphed. In *Like* (right) graph, ASKCOMMENT for website is at -4.8 but zoomed in to make the graph look better.

4.3 Out-of-Sample Prediction & Managerial Implications

To conclude the paper, we assess the extent to which the models we develop may be used as an aid to content engineering, and to predict the expected levels of engagement for various hypothetical content profiles a firm may consider for a potential message it could serve to users. First, we present an illustration set of out-of-sample prediction of engagement with real posts. Then we discuss a back-of-the-envelope calculation to show how adding or removing particular content profiles may affect engagement outcomes for typical posts in our data. Our intent is to give the reader a rough sense of the predictive validity of our estimates as a tool to assess expected engagement for hypothetical content bundles.

To illustrate the above, we choose three posts released around the same time outside our sample. To emphasize that our second-stage model of engagement has predictive power, we choose these to be for the same firm, of the same message type and having roughly the same number of impressions (i.e., we are taking out the effect of *EdgeRank*). Table 10 shows the messages, the actual lifetime engagement realized for those messages, and the content tags we generated for those messages. For each message, we use the coefficients from our second stage to predict the expected engagement. In the “Content Coef:” column, we present the latent linear index of the logistic function of the respective engagement probabilities obtained by multiplying

the coefficients for the engagement model (Table 9, EdgeRank corrected) with indicator variables for whether each type of content attribute is present in these messages, and then adding these up. In the last two columns we present the predicted and actual ranks for the three messages in terms of their engagement. We see that the match is very good and that the model can effectively help a marketer select the post that maximizes engagement.

Now imagine that a firm starts with the second message as the base creative. Content engineering for this message using the model is straightforward. For instance, if the marketer is assessing the potential impact of adding a philanthropic aspect to message two, we can determine that it will increase the latent linear index for comments from 0.941 to $0.941 + 0.174 = 1.115$, which increases the predicted comments rank of this message to 1, and increases the predicted odds ratio for comments by 19%. Similarly, if the marketer is considering asking for comments explicitly, this will increase the number of comments for the message obtained by increasing the latent linear index from 0.941 to $0.941 + 0.469 = 1.410$. In this sense, the model is able to aid the assessment of the anticipated engagement from various possible content bundles.

Next, we discuss a back of the envelope calculation to assess the effect of adding new content to typical posts in our data. The average post in the data obtains 10,000 impressions, 5 comments and 50 *Likes*. Thus, the probability of a comment is 0.0005 and the probability of a *Like* is 0.005. Suppose one can engineer a change in content that increases the probability of comments and *Likes* respectively to (0.00065,0.0065) – this increase, for instance, is the increase predicted by our estimates generated by adding emotional content to a post with no emotional content. The new predicted comments and *Likes* are 6.5 and 65 respectively, i.e., a 30% increase in comments and *Likes*. Now note that the standard deviation of the number of impressions is 129,874. For a message two standard deviations from the mean number of impressions, i.e., at $10,000 + 2 \times 129,874 = 269,748$ impressions, a 30% increase in comments and *Likes* translates to roughly an increase of 41 comments and 405 *Likes*, suggesting that content engineering can produce a fairly substantial increase in engagement for many posts.

Sample Messages {Actual Comments, Actual Likes}	Content Tags	Content Coef {Com,Likes}	Comments Rank	Likes Rank
<i>Don't forget in celebration of hitting over 70,000 we are giving all our awesome fans {exclusively} the "employee discount" take 20% off your entire order on our website {http://anonymized.com} with the code: SOMECODE and it is good until 3/16/12. Enjoy some shopping on us :) {12, 83}</i>	HTTP, DEAL, PRODLOCATION, PRODAVAIL, EMOTICON	{-0.459,-0.351}	Actual:3 Predicted:3	Actual:2 Predicted:2
<i>Who is ready for a givvvveeawayyyy?! :) :) (35 mins from now!) {132, 438}</i>	EMOTION, EMOTICON, QUESTION	{0.941,0.065}	Actual:2 Predicted:2	Actual:1 Predicted:1
<i>COMPLETE THIS SENTENCE: Crafting is best with ----- {416, 72}</i>	BLANK, SMALLTALK	{1.060,-0.742}	Actual:1 Predicted:1	Actual:3 Predicted:3

Table 10: Predicted versus Actual Engagement Ranking for Three Illustrative Posts: Note: we anonymized some parts of the messages for presentation.

5 Conclusions and Implications

We show through a large-scale study that content engineering in social media has a significant impact on user engagement as measured by *Likes* and comments for posts. Our analysis shows that persuasive content, such as emotional and philanthropic content, has a positive impact on engagement. This suggests that firms gain from sharing their brand personality and information about their social initiatives in social media. Further, we find that product informative content has a negative impact on user engagement. This presents a challenge to marketers who seek to build a large following on social media and who seek to leverage that following to disseminate information about new products and promotions. One takeaway from our study is that these strategies work when product informative content is combined with persuasive content. In addition, our results are moderated by industry type suggesting there is no one-size-fits-all content strategy and that firms need to test multiple content strategies. These results account for post selection by *EdgeRank*, Facebook’s filtering algorithm, which to our knowledge, has been largely ignored in the literature.

Because of the scale of our study (over 800 firms and 100,000 messages analyzed), we believe our results generalize and have broad applicability. Nonetheless, it is important to recognize the results from any study on consumer response to content depend on the mix of content used in the study. For example, we find that posts mentioning holidays, especially by consumer product companies, have a negative effect on engagement. This may be due to excessive use of holiday messages by firms. It is possible that the effect may be positive if firms use these kinds of posts in moderation. Similarly, we find that emotional messages have a positive impact on engagement. Here again, it is possible this effect may reduce in the future if firms start using emotional content excessively pushing consumer response to the region of declining returns. Hence, it is important to interpret our results in the context of the content mix used by firms and redo the analysis in the event of large-scale changes in the content mix used by firms.

We used two metrics for user engagement, namely *Likes* and comments on posts. There may be other measures worth considering, including whether users share posts with friends, visit the websites of firms posting messages, or buy more products from these firms. Our use of *Likes* and comments is motivated both by the widespread use of these metrics in social media settings, and also the availability of data. Future studies that evaluate other measures of interest can add value, particularly in validating the generalizability of our findings and in exploring mechanisms underpinning the effects we describe. On a related note, as we acknowledge upfront in the introduction of the paper, we do not address the question of how engagement affects product demand and firm’s profits so as to complete the link between ad-attributes and those outcome measures. Such data are still not widely available at the scale needed for this study. Further, advertisers are often interested in social media engagement *per se* on the maintained assumption that such engagement can often to translate into brand-loyalty or purchases in the long-term. Although it is not the focus of our study, it is worth highlighting that several extant studies have studied the link between Facebook advertising and engagement (albeit at a smaller scale). For example, based on randomized studies, comScore (2012) reports a 38% lift in purchase for fans exposed to Starbucks advertising on Facebook through Facebook Pages or Facebook paid advertising (in the same study, compared to the control group, fans of Target were 19% more likely to purchase at Target in the four weeks following exposure to Facebook messages). Chadwick-Martin-

Bailey (2010) document that users who sign-up as fans of the Facebook Page of a firm are more likely to recommend and buy the product of the page than before. Kumar et al. (2013) show that social media can be used to generate growth in sales, and ROI, connecting social media metrics such as “comments” to financial metrics.

The competition for consumer attention across media outlets is intense, especially on social media platforms. Consumers, in turn, are overwhelmed by the proliferation of online content, and it seems clear that marketers will not succeed without engineering this content for their audience. We hope this study contributes to improve content engineering by firms on social media sites and, more generally, creates interest in evaluating the effects of advertising content on consumer engagement.

References

- Abernethy, A. M. and G. R. Franke: 1996, ‘The Information Content of Advertising: A Meta-Analysis’. *Journal of Advertising* **25**(2), 1–17.
- Anand, B. and R. Shachar: 2009, ‘Targeted advertising as a signal’. *Quantitative Marketing and Economics* **7**(3), 237–266.
- Anand, B. and R. Shachar: 2011, ‘Advertising the Matchmaker’. *Rand Journal of Economics* **42**(2), 205–245.
- Anderson, S. and R. Renault: 2006, ‘Advertising Content’. *American Economic Review* **March**, 93–113.
- Aral, S., C. Dellarocas, and D. Godes: 2013, “Social Media and Business Transformation: A Framework for Research”. *Information Systems Research* **Article in Advance**(1-11).
- Archak, N., A. Ghose, and P. G. Ipeirotis: 2011, ‘Deriving the Pricing Power of Product Features by Mining Consumer Reviews’. *Management Science* **57**(8), 1485–1509.
- Armstrong, S.: 2010, *Persuasive Advertising*. Palgrave Macmillan.
- Bagwell, K.: 2007, ‘Chapter 28: The Economic Analysis of Advertising’. *Handbook of IO* **3**, 1701–1844.
- Bennett, P. N.: 2006, ‘Building Reliable Metaclassifiers for Text Learning’. Ph.D. thesis, CMU.
- Berger, J. and K. L. Milkman: 2012, ‘What Makes online Content Viral?’. *Journal of Marketing Research* **49**(2), 192–205.
- Bertrand, M., D. Karlan, S. Mullianathan, E. Shafir, and J. Zinman: 2010, ‘What’s Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment’. pp. 263–306.
- Bird, S., E. Klein, and E. Loper: 2009, *Natural Language Processing with Python*. O’Reilly Media.
- Blundell, R. and J. Powell: 2003, *Endogeneity in Nonparametric and Semiparametric Regression Models*, Vol. II of *Econometric Monograph Series* 36. Cambridge University Press.
- Brieman, L.: 1996, ‘Bagging predictors’. *Machine Learning* **24**, 123–140.
- Buhrmester, M., T. Kwang, and S. D. Gosling: 2011, ‘Amazon’s Mechanical Turk: a New Source of Inexpensive, Yet High-Quality, Data?’. *Psychological Science* **6**(1).
- Butters, G.: 1977, ‘Equilibrium Distributions of Sales and Advertising Prices’. *Review of Economic Studies* **44**(3), 465–91.
- Chadwick-Martin-Bailey: 2010, ‘Consumers Engaged Via Social Media Are More Likely To Buy, Recommend’.
- Chandy, R., G. Tellis, D. Macinnis, and P. Thaivanich: 2001, ‘What to Say When: Advertising Appeals in Evolving Markets’. *Journal of Marketing Research* **38**, 399–414.

- Cialdini, R.: 2001, *Influence: Science and Practice*. Needham Heights, MA: Allyn and Bacon.
- comScore: 2012, ‘The Power of Like 2’. Technical report, comScore.
- Creamer, M.: 2012, ‘Study Only 1 percent of Facebook ‘Fans’ Engage With Brands’. *AdAge*.
- Freund, Y. and R. E. Schapire: 1995, ‘A decision-theoretic Generalization of on-line learning and an application to boosting’. *Computational Learning Theory: EUROCOLT* pp. 23–37.
- Gardete, P. M.: 2013, ‘Cheap-Talk Advertising and Misrepresentation in Vertically Differentiated Markets’. *Marketing Science*.
- Gentzkow, M. and J. Shapiro: 2010, ‘What Drives Media Slant? Evidence from U.S. Newspapers’. *Econometrica* **78**(1).
- Ghose, A., P. G. Ipeirotis, and B. Li: 2012, ‘Designing Ranking Systems for Hotels on Travel Search Engines by Mining User-Generated and Crowd-Sourced Content’. *Marketing Science* **31**(3).
- Grossman, G. and C. Shapiro: 1984, ‘Informative Advertising with Differentiated Products’. *Review of Economic Studies* **51**(1), 63–81.
- Hassan, M. H., B. S. Vikas, and H. Fiorletta: 2011, ‘Accurate Information Extraction for Quantitative Financial Events’. *CIKM*.
- Hastie, T., R. Tibshirani, and J. Friedman: 2009, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer.
- Hastie, T. and R. J. Tibshirani: 1990, *Generalized Additive Models*. Chapman and Hall CRC.
- Ipeirotis, P. G., F. Provost, and J. Wang: 2010, ‘Quality Management on Amazon Mechanical Turk’. *HCOMP*.
- Jurafsky, D. and J. H. Martin: 2008, *Speech and Language Processing*. Prentice Hall.
- Kihlstrom, R. and M. Riordan: 1984, ‘Advertising as a Signal’. *Journal of Political Economy* **92**(3), 427–50.
- Kumar, V., V. Bhaskara, R. Mirchandani, and M. Shah: 2013, ‘Creating a Measurable Social Media Marketing Strategy: Increasing the Value and ROI of Intangibles and Tangibles for Hokey Pokey’. *Marketing Science* **32**(2), 194–212.
- Liaukonyte, J., T. Teixeira, and K. Wilbur: 2013, ‘TV Advertising Content and Online Shopping’. (*in progress*).
- Liu, B.: 2011, *Sentiment Analysis and Opinion Mining*. Morgan and Claypool.
- Lodish, L., M. Abraham, S. Kalmenson, J. Livelsberger, B. Lubetkin, B. Richardson, and M. E. Stevens: 1995, ‘How T.V. Advertising Works: a Meta-Analysis of 389 Real World Split Cable T.V. Advertising Experiments’. *Journal of Marketing Research* **32**(2), 125–139.
- Mason, W. A. and S. Suri: 2012, ‘Conducting behavioral research on Amazon’s Mechanical Turk’. *Behavior Research Methods* **44**(1), 1–23.
- Mayzlin, D. and J. Shin: 2011, ‘Uninformative Advertising as an Invitation to Search’. *Marketing Science* **30**(4), 666–685.
- Milgrom, P. and J. Roberts: 1986, ‘Price and Advertising Signals of Product Quality’. *Journal of Political Economy* **94**(4), 796–821.
- Nan, X. and R. J. Faber: 2004, ‘Advertising Theory: Reconceptualizing the Building Blocks’. *Marketing Theory* **4**, 7–30.
- Nelson, P.: 1974, ‘Advertising as Information’. *Journal of Political Economy* **82**(4), 729–54.
- Netzer, O., R. Feldman, J. Goldenberg, and M. Fresko: 2012, ‘Mine Your Own Business: Market-Structure Surveillance Through Text Mining’. *Marketing Science* **31**(3), 521–543.

- Paolacci, G., J. Chandler, and P. G. Ipeirotis: 2010, ‘Running Experiments on Amazon Mechanical Turk’. *Judgement and Decision Making* **5**(5).
- Resnik, A. and B. L. Stern: 1977, ‘An Analysis of Informative Content in Television Advertising’. *Journal of Marketing* **41**(1).
- Sheng, V., F. Provost, and P. G. Ipeirotis: 2007, ‘Get another label? Improving data quality and data mining using multiple, noisy labelers’. *14th ACM SIGKDD* pp. 614–622.
- Snow, R., B. O’Connor, D. Jurafsky, and A. Y. Ng: 2008, ‘Cheap and fast - but is it good? Evaluating non-expert annotations for natural language tasks’. *Empirical Methods in Natural Language Processing*.
- Sundararajan, A., F. Provost, G. Oestreicher-Singer, and S. Aral: 2013, ‘Information in Digital, Economic, and Social Networks’. *Information Systems Research Article in Advance*, 1–23.
- Tucker, C.: 2012a, ‘Ad Virality and Ad Persuasiveness’. *MIT Working paper*.
- Tucker, C.: 2012b, ‘Social Advertising’. *Working paper, MIT*.
- Vakratsas, D. and T. Ambler: 1999, ‘How Advertising Works: What do we really know?’. *Journal of Marketing* **63**(1).
- Wildfire: 2012, ‘Edgerank and Your Brand: An Analysis of Messaging Engagement on Facebook by Business Category’. Technical report, Wildfire.
- Wood, S. N.: 2006, *Generalized Additive Models: An Introduction with R*. CRC press.
- Xu, L. and A. Krzyzak: 1992, ‘Methods of Combining Multiple Classifiers and Their Applications to Handwriting Recognition’. *IEEE Transactions on systems*.
- Zadrozny, B. and C. Elkan: 2002, ‘Transforming Classifier Scores into Accurate Multiclass Probability Estimates’. *KDD*.

Appendix 1: Survey Instrument

We are interested in identifying the content of social media messages. Please read the following short message posted on a social media site and answer the following questions. Thank you.

{Message Content}

Indicate the message characteristics by (Yes/No).

The Message,

1: contains an interesting or remarkable fact.
2: includes emotional content.
3: uses humor.
4: includes **philanthropic, awareness or activist content** (promotes social welfare, social change, or informs the reader about certain facts to improve people's lives).
5: **mentions a company (organization) name or brand name** (e.g. Nike, Red Cross, VitaminWater or musicians such as The Beatles).
If yes, please type the company or brand name as it appears in the message in the following box (only one per box)
Most relevant company/brand mentioned _____
Second company/brand mentioned (if any) _____
6: provides information about any type of **discounts or freebies** (e.g., deals, coupons, promotional offers, sweepstakes, rewards, free items).
7: makes a **price comparison** or a **price-match guarantee** against competitors' product(s) or service(s).
8: contains content **other than** about a product or company business (e.g. small talk, social pleasantries, casual banter, "Happy Halloween", "what was your day like?").
9: **is targeted towards an audience segment**. E.g., a particular demographic (race, gender, age, location) or people with certain qualifications or characteristics. (e.g. "all moms, do you like Gerber?", "All chocolate lovers, get this deal").
10: provides information about the **availability of a product or service** (e.g. Only 5 deals remain, available only until tomorrow, new product coming tomorrow).
11: provides information about **where or how to obtain a product or service** (e.g. link, physical location to buy, general location)?
12: does the message have a dollar sign (\$)?
13: Are your friends on social media likely to post messages such as the above?
14: The message is specific to a **particular product or service** (e.g. electronics, drinks, music, etc.).
If yes, please type the mentioned product or service:
Most relevant product mentioned _____

Figure 14: Survey Form Used in Amazon Mechanical Turk

Appendix 2: Discussion of the Number of Impressions

We discuss our procedure for constructing $n_{kjt}^{(d)}$, the number of impressions for each message k of firm j in day t split by demographic bin d for use in the EdgeRank correction model. As mentioned above, $n_{kjt}^{(d)}$ is not directly reported by Company X (or made available to page-owners by Facebook *Insights*). Instead, Company X reports $n_{jt}^{(d)}$, the number of impressions for all posts associated with firm j in demographic bucket d on day t , which is essentially $n_{kjt}^{(d)}$ summed across all k associated with j . To assess how we may split this across the various demographic buckets, we checked the extent to which pages release different types of posts over time. The bulk of impressions for a post occur within the first week of its release. Hence, the total impressions for a page on a given day out of a specific demographic bucket, $n_{jt}^{(d)}$, reflects the aggregate impressions to users in that bucket of all posts released by that firm over the past one week. Since EdgeRank allocates posts to users by post-type, if the firm releases the same *type of posts* (i.e., photos, videos, status updates, apps or links) over a week's duration, then the split of $n_{jt}^{(d)}$ across the various k posts released by firm j within the past week may be roughly the same. In other words, the distribution of demographics of the impressions of all posts released by a firm in the past week should be the same if all those released posts are similar.

To check informally if this is the case, we picked a random sample of 10,000 page-7-day combinations from our data. For each combination, we collated all the posts released by that page during that 7-day

window and tabulated the type of these posts (i.e., photos, videos, status updates, apps or links). We then construct two concentration metrics, \mathbb{C}_1 , the proportion of posts released by that page during that 7-day window that belong to the highest type bucket, and \mathbb{C}_2 , the proportion of posts released by that page during that 7-day window that belong to the highest and second-highest type bucket. \mathbb{C}_1 and \mathbb{C}_2 are analogous to top-firm and top-two-firm concentration ratios used in industry-concentration studies, and measure the extent to which the posts released by a page in a given 7-day period are spread across types. If all posts released by a page during that 7-day window are of the same type, \mathbb{C}_1 and \mathbb{C}_2 will be both 1. The spread away from 1 thus indicates higher variation in post-types released by an average Facebook page over a week’s duration. Table 11 reports on the distribution of \mathbb{C}_1 and \mathbb{C}_2 we computed in this manner. Looking at Table 11, we find that the median \mathbb{C}_1 is .71 (mean .72) and the median \mathbb{C}_2 is 1.0 (mean .94). Most pages seem to be releasing at-most 2-types of posts within a week window, and more than $2/3^{rd}$ of posts released by an average page in an average week are of the same type. Given this, we assume that $n_{jt}^{(d)}$ is equally split across all k associated with a firm over the past 7-day period. We construct the variable $n_{kjt}^{(d)}$ in the left hand-side of the EdgeRank correction equation 2 in this manner.

	<i>Min.</i>	<i>1st Qu.</i>	<i>Median</i>	<i>Mean</i>	<i>3rd Qu.</i>	<i>Max.</i>
\mathbb{C}_1	0.250	0.535	0.706	0.719	0.915	1.000
\mathbb{C}_2	0.500	0.889	1.000	0.942	1.000	1.000

Table 11: **Distribution of the Top (\mathbb{C}_1) and Top-two (\mathbb{C}_2) Concentration Ratios of the type of Posts Served by a Facebook Page over a Randomly picked 7-day period**

The method is not without its limitations. We view it as a practical way to deal with the lack of data-reporting by Facebook, while exploiting the variation embedded in the observed impressions and to correlate it with the observed variation in the the potential market for each post in each demographic bucket ($\mathbb{N}_{jt}^{(d)}$). The method produces potential measurement error in the dependent variable, $n_{kjt}^{(d)}$ in the EdgeRank correction stage. Measurement error in the dependent variable is absorbed into the RHS unobservables and is usually less of a concern. The fact that we include page-fixed effects separately for each demographic ($\theta_{1j}^{(d)}$ in Equation 3) also mitigates concerns that these unobservables may systematically be correlated with included characteristics. More broadly, to the best of our knowledge, the full details of EdgeRank are not known to any firm or researcher. In our view, a “perfect” solution to the selection problem is unlikely to be achieved without knowledge of Facebook’s targeting rule.

Appendix 3: Result for AMT-Tagged & Different Models

	NO ER COMMENT	OR	ER COMMENT	OR	NO ER LIKE	OR	ER LIKE	OR
Constant	-7.099***(0.010)	0.001	-6.042***(0.010)	0.002	-4.715***(0.003)	0.009	-3.779***(0.003)	0.023
SMOG	-0.100***(0.001)	0.905	-0.123***(0.001)	0.884	-0.047***(0.000)	0.954	-0.062***(0.000)	0.940
MSGLEN	-0.001***(0.000)	0.999	-0.000***(0.000)	1.000	-0.000***(0.000)	1.000	-0.000***(0.000)	1.000
HTTP	-0.304***(0.005)	0.738	-0.183***(0.005)	0.833	-0.111***(0.002)	0.895	-0.026***(0.001)	0.974
QUESTION	0.228***(0.004)	1.256	0.253***(0.004)	1.288	-0.303***(0.001)	0.739	-0.197***(0.001)	0.821
BLANK	0.932***(0.012)	2.540	0.914***(0.012)	2.494	-0.957***(0.010)	0.384	-1.076***(0.010)	0.341
ASKLIKE	-0.351***(0.033)	0.704	-0.240***(0.033)	0.787	-0.066***(0.009)	0.936	0.085***(0.009)	1.089
ASKCOMMENT	0.678***(0.048)	1.970	0.202***(0.048)	1.224	0.123***(0.020)	1.131	-0.380***(0.020)	0.684
Persuasive								
REMACT	0.169***(0.006)	1.184	0.203***(0.005)	1.225	-0.017***(0.002)	0.983	0.012***(0.002)	1.012
EMOTION	0.119***(0.006)	1.126	0.184***(0.006)	1.202	0.152***(0.002)	1.164	0.203***(0.002)	1.225
EMOTICON	-0.523***(0.026)	0.593	-0.594***(0.026)	0.552	-0.623***(0.008)	0.536	-0.528***(0.008)	0.590
HOLIDAYMENTION	-1.483***(0.033)	0.227	-1.328***(0.033)	0.265	-0.428***(0.006)	0.652	-0.277***(0.006)	0.758
HUMOR	-0.131***(0.012)	0.877	0.022(0.012)	1.022	-0.358***(0.004)	0.699	-0.143***(0.004)	0.867
PHILANTHROPIC	0.351***(0.007)	1.420	0.217***(0.007)	1.242	0.226***(0.002)	1.254	0.108***(0.002)	1.114
FRIENDLIKELY	-0.168***(0.005)	0.845	-0.104***(0.005)	0.901	0.205***(0.002)	1.228	0.185***(0.002)	1.203
SMALLTALK	-0.003(0.004)	0.997	-0.099***(0.004)	0.906	-0.001(0.001)	0.999	-0.070***(0.001)	0.932
Informative								
BRANDMENTION	-0.195***(0.004)	0.823	-0.184***(0.004)	0.832	-0.118***(0.001)	0.889	-0.143***(0.001)	0.867
DEAL	-0.023***(0.007)	0.977	0.086***(0.006)	1.090	-0.426***(0.002)	0.653	-0.226***(0.002)	0.798
PRICECOMPARE	0.826***(0.193)	2.284	0.385*(0.193)	1.470	-0.452***(0.103)	0.636	-0.839***(0.103)	0.432
PRICE	-0.098***(0.014)	0.907	-0.611***(0.014)	0.543	-0.320***(0.005)	0.726	-0.826***(0.005)	0.438
TARGET	0.104***(0.010)	1.110	-0.109***(0.010)	0.897	0.035***(0.003)	1.036	-0.089***(0.003)	0.915
PRODAVAIL	-0.259***(0.007)	0.772	-0.366***(0.007)	0.694	0.069***(0.002)	1.071	-0.051***(0.002)	0.950
PRODLOCATION	-0.111***(0.006)	0.895	0.008(0.005)	1.008	-0.208***(0.002)	0.812	-0.078***(0.002)	0.925
PRODMENTION	-0.194***(0.005)	0.824	-0.392***(0.005)	0.676	0.272***(0.001)	1.313	0.054***(0.001)	1.055
Message Type - App is the base								
-Link	0.191***(0.008)	1.210	-0.354***(0.008)	0.702	-0.100***(0.003)	0.905	-0.640***(0.003)	0.527
-Photo	0.577***(0.008)	1.781	0.050***(0.008)	1.051	0.653***(0.002)	1.921	0.102***(0.002)	1.107
-Status Update	1.273***(0.008)	3.572	0.752***(0.008)	2.121	0.588***(0.003)	1.800	0.168***(0.003)	1.183
-Video	-0.240***(0.027)	0.787	-0.665***(0.027)	0.514	-0.650***(0.011)	0.522	-1.211***(0.011)	0.298
Industry Category - Celebrity is the base								
-ConsumerProduct	-0.040***(0.006)	0.961	-0.316***(0.006)	0.729	-0.506***(0.002)	0.603	-0.700***(0.002)	0.497
-Entertainment	0.166***(0.006)	1.181	0.386***(0.006)	1.471	-0.344***(0.002)	0.709	-0.001(0.002)	0.999
-Organization	0.365***(0.006)	1.441	0.307***(0.006)	1.359	-0.026***(0.002)	0.974	0.039***(0.002)	1.040
-PlaceBusiness	0.563***(0.014)	1.756	0.275***(0.014)	1.317	-0.703***(0.007)	0.495	-0.967***(0.007)	0.380
-Websites	0.164***(0.008)	1.178	0.224***(0.007)	1.251	0.214***(0.002)	1.239	0.336***(0.002)	1.399
McFadden R-sq.	0.244		0.198		0.262		0.16	
Nagelkerke R-sq.	0.244		0.198		0.262		0.161	
Log-likelihood	-315945.323		-449901.889		-1808652.064		-3333155.941	
Deviance	586061.615		853991.131		3524714.113		6573726.255	
AIC	631988.646		899901.778		3617402.127		6666409.881	
N	38706		38706		38706		38706	
Significance '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1								

Table 12: **Aggregate Logistic Regression Results For Comments and Likes (5000 Messages):** This table presents the aggregate logistic regression on comments and *Likes* for both EdgeRank-corrected (ER) and uncorrected (NO ER) for 5000 messages data tagged by Turkers. OR means Odds ratio and shows the odds ratio for the estimates left of the column.

Variable	Intercept only	Controls	Friendly	Persuasive	Informative	All
Constant	-7.052***(0.001)	-6.906***(0.003)	-6.952***(0.004)	-6.975***(0.004)	-6.861***(0.004)	-6.889***(0.004)
SMOG		-0.065***(0.000)	-0.065***(0.000)	-0.065***(0.000)	-0.067***(0.000)	-0.068***(0.000)
MSGLEN		-0.000***(0.000)	-0.000***(0.000)	-0.000***(0.000)	-0.000***(0.000)	-0.000***(0.000)
HTTP		-0.406***(0.001)	-0.400***(0.001)	-0.393***(0.001)	-0.390***(0.002)	-0.355***(0.002)
QUESTION		0.543***(0.001)	0.546***(0.001)	0.557***(0.001)	0.541***(0.001)	0.564***(0.001)
BLANK		1.172***(0.003)	1.168***(0.003)	1.187***(0.003)	1.144***(0.003)	1.146***(0.003)
ASKLIKE		0.222***(0.010)	0.229***(0.010)	0.220***(0.010)	0.223***(0.010)	0.221***(0.010)
ASKCOMMENT		0.508***(0.021)	0.529***(0.021)	0.465***(0.021)	0.506***(0.021)	0.469***(0.021)
Message Type - App is the base						
-Link		-0.388***(0.003)	-0.389***(0.003)	-0.394***(0.003)	-0.376***(0.003)	-0.370***(0.003)
-Photo		0.376***(0.003)	0.371***(0.003)	0.365***(0.003)	0.372***(0.003)	0.373***(0.003)
-Status Update		0.656***(0.003)	0.654***(0.003)	0.644***(0.003)	0.645***(0.003)	0.641***(0.003)
-Video		0.406***(0.009)	0.406***(0.009)	0.410***(0.009)	0.395***(0.009)	0.398***(0.009)
Industry Category - Celebrity is the base						
-ConsumerProduct		-0.402***(0.002)	-0.392***(0.002)	-0.371***(0.002)	-0.374***(0.002)	-0.347***(0.002)
-Entertainment		0.514***(0.002)	0.519***(0.002)	0.537***(0.002)	0.514***(0.002)	0.529***(0.002)
-Organization		0.391***(0.002)	0.393***(0.002)	0.409***(0.002)	0.400***(0.002)	0.408***(0.002)
-PlaceBusiness		-0.087***(0.005)	-0.082***(0.005)	-0.052***(0.005)	-0.065***(0.005)	-0.021***(0.005)
-Websites		0.123***(0.002)	0.134***(0.002)	0.150***(0.003)	0.136***(0.003)	0.182***(0.003)
FRIENDLIKELY			0.064***(0.001)	-0.051***(0.002)		-0.006***(0.002)
REMFAC				-0.065***(0.002)		0.014***(0.002)
EMOTION				0.217***(0.002)		0.256***(0.002)
EMOTICON				0.109***(0.004)		0.121***(0.004)
HOLIDAYMENTION				-0.391***(0.014)		-0.388***(0.014)
HUMOR				0.022***(0.001)		0.072***(0.002)
PHILANTHROPIC				0.054***(0.002)		0.174***(0.002)
SMALLTALK				-0.055***(0.002)		-0.086***(0.002)
BRANDMENTION					0.071***(0.002)	0.081***(0.002)
DEAL					-0.124***(0.002)	-0.172***(0.002)
PRICECOMPARE					0.043***(0.001)	-0.006***(0.001)
PRICE					-0.368***(0.005)	-0.317***(0.005)
TARGET					0.049***(0.002)	-0.071***(0.002)
PRODAVAIL					0.028***(0.002)	-0.064***(0.002)
PRODLOCATION					0.023***(0.002)	0.011***(0.002)
PRODMENTION					-0.186***(0.002)	-0.151***(0.002)
McFadden R-sq.		0.161	0.161	0.165	0.164	0.171
Nagelkerke R-sq.		0.161	0.162	0.165	0.165	0.172
Log-likelihood	-4267283.759	-3612205.117	-3610817.72	-3597891.814	-3598824.088	-3570184.37
Deviance	8130598.643	6820441.36	6817666.566	6791814.753	6793679.301	6736399.867
AIC	8534569.517	7224474.234	7221701.44	7195863.627	7197730.175	7140466.74
N	665916	665916	665916	665916	665916	665916
Significance '****' 0.001 '***' 0.01 '**' 0.05 '.' 0.1						

Table 13: Logistic Regression EdgeRank-Corrected Estimates Model Comparison (Comments)

Variable	Intercept only	Controls	Friendly	Persuasive	Informative	All
Constant	-4.662***(0.000)	-3.929***(0.001)	-3.979***(0.001)	-3.931***(0.001)	-3.921***(0.001)	-3.892***(0.001)
SMOG		-0.057***(0.000)	-0.057***(0.000)	-0.057***(0.000)	-0.060***(0.000)	-0.061***(0.000)
MSGLEN		-0.001***(0.000)	-0.000***(0.000)	-0.001***(0.000)	-0.000***(0.000)	-0.000***(0.000)
HTTP		-0.214***(0.000)	-0.207***(0.000)	-0.205***(0.000)	-0.205***(0.000)	-0.189***(0.000)
QUESTION		-0.198***(0.000)	-0.195***(0.000)	-0.181***(0.000)	-0.198***(0.000)	-0.175***(0.000)
BLANK		-0.603***(0.002)	-0.612***(0.002)	-0.592***(0.002)	-0.602***(0.002)	-0.596***(0.002)
ASKLIKE		0.576***(0.003)	0.583***(0.003)	0.576***(0.003)	0.564***(0.003)	0.565***(0.003)
ASKCOMMENT		-0.379***(0.011)	-0.359***(0.011)	-0.444***(0.011)	-0.404***(0.011)	-0.465***(0.011)
Message Type - App is the base						
-Link		-0.517***(0.001)	-0.519***(0.001)	-0.518***(0.001)	-0.511***(0.001)	-0.505***(0.001)
-Photo		0.572***(0.001)	0.567***(0.001)	0.561***(0.001)	0.566***(0.001)	0.561***(0.001)
-Status Update		-0.062***(0.001)	-0.067***(0.001)	-0.083***(0.001)	-0.067***(0.001)	-0.083***(0.001)
-Video		0.244***(0.003)	0.245***(0.003)	0.252***(0.003)	0.225***(0.003)	0.231***(0.003)
Industry Category - Celebrity is the base						
-ConsumerProduct		-0.948***(0.001)	-0.937***(0.001)	-0.941***(0.001)	-0.939***(0.001)	-0.932***(0.001)
-Entertainment		-0.172***(0.001)	-0.168***(0.001)	-0.174***(0.001)	-0.189***(0.001)	-0.193***(0.001)
-Organization		-0.165***(0.001)	-0.164***(0.001)	-0.160***(0.001)	-0.171***(0.001)	-0.171***(0.001)
-PlaceBusiness		-1.305***(0.002)	-1.301***(0.002)	-1.292***(0.002)	-1.294***(0.002)	-1.275***(0.002)
-Websites		0.028***(0.001)	0.038***(0.001)	0.019***(0.001)	0.042***(0.001)	0.041***(0.001)
FRIENDLIKELY			0.072***(0.000)	0.051***(0.001)		0.080***(0.001)
REMFAC				-0.065***(0.001)		-0.021***(0.001)
EMOTION				0.250***(0.001)		0.260***(0.001)
EMOTICON				-0.049***(0.001)		-0.020***(0.001)
HOLIDAYMENTION				-0.187***(0.004)		-0.183***(0.004)
HUMOR				-0.009***(0.000)		0.009***(0.000)
PHILANTHROPIC				-0.072***(0.001)		0.002*(0.001)
SMALLTALK				-0.112***(0.001)		-0.146***(0.001)
BRANDMENTION					0.012***(0.000)	0.021***(0.000)
DEAL					-0.166***(0.001)	-0.207***(0.001)
PRICECOMPARE					-0.036***(0.000)	-0.047***(0.000)
PRICE					-0.504***(0.001)	-0.471***(0.001)
TARGET					0.050***(0.001)	-0.019***(0.001)
PRODAVAIL					-0.027***(0.001)	-0.073***(0.001)
PRODLOCATION					0.134***(0.001)	0.138***(0.001)
PRODMENTION					-0.004***(0.001)	0.012***(0.001)
McFadden R-sq.		0.191	0.191	0.195	0.196	0.201
Nagelkerke R-sq.		0.192	0.193	0.197	0.197	0.203
Log-likelihood	-33968732.078	-27584573.527	-27566219.315	-27424516.699	-27419099.696	-27221231.997
Deviance	66990021.316	54221704.214	54184995.789	53901590.557	53890756.552	53495021.154
AIC	67937466.156	55169211.054	55132504.629	54849113.398	54838281.393	54442561.994
N	665916	665916	665916	665916	665916	665916
Significance '****' 0.001 '***' 0.01 '**' 0.05 '.' 0.1						

Table 14: Logistic Regression EdgeRank-Corrected Estimates Model Comparison (*Likes*)