



City Research Online

City St George's, University of London

Citation: Lambrecht, A., Tucker, C. M. & Wiertz, C. (2018). Advertising to Early Trend Propagators: Evidence from Twitter. *Marketing Science*, 37(2), pp. 177-199. doi: 10.1287/mksc.2017.1062

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/17540/>

Link to published version: <https://doi.org/10.1287/mksc.2017.1062>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).



City Research Online

City, University of London Institutional Repository

Citation: Lambrecht, A., Tucker, C. M. and Wiertz, C. (2018). Advertising to Early Trend Propagators: Evidence from Twitter. *Marketing Science*, 37(2), pp. 177-331. doi: 10.1287/mksc.2017.1062

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/17540/>

Link to published version: <http://dx.doi.org/10.1287/mksc.2017.1062>

Copyright and reuse: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Advertising to Early Trend Propagators: Evidence from Twitter

Anja Lambrecht, Catherine Tucker and Caroline Wiertz*

May 25, 2017

Abstract

In the digital economy, influencing and controlling the spread of information is a key concern for firms. One way firms try to achieve this is to target firm communications to consumers who embrace and propagate the spread of new information on emerging and ‘trending’ topics on social media. However, little is known about whether early trend propagators are indeed responsive to firm-sponsored messages. To explore whether early propagators of trending topics respond to advertising messages, we use data from two field tests conducted by a charity and an emerging fashion firm on the micro-blogging service Twitter. On Twitter, ‘promoted tweets’ allow advertisers to target individuals based on the content of their recent postings. Twitter continuously identifies in real time which topics are newly popular among Twitter users. In the field tests, we collaborated with a charity and a fashion firm to target ads at consumers who embraced a Twitter trend early in its life-cycle by posting about it, and compared their behavior to that of consumers who posted about the same topic only later on. Throughout both field tests, we consistently find that early propagators of trends are less responsive to advertising than consumers who embrace trends later.

Keywords: Online Advertising, Targeting, Twitter, User-Generated Content, Internet

JEL Codes: L86, M37

*Anja Lambrecht is an Associate Professor at London Business School, London NW1 4SA, UK. Catherine Tucker is the Sloan Distinguished Professor of Marketing at MIT Sloan School of Management, MIT, Cambridge, MA 02139, USA. Caroline Wiertz is a Professor of Marketing at Cass Business School, City, University of London, London EC1Y 8TZ, UK. This work was supported in part by the Marketing Science Institute and the Deloitte Institute of Innovation and Entrepreneurship. We thank Kevin Curran, Prasad Lakshmi Vana, Haemi Shin, Vaidheesh Venkatram, Andrew Leung, Leandro Galli, and Mario Campana for excellent research assistance, and an anonymous charity, Tim Bissett, Nicolai Schuemann, Sean Keyes and Catherine Bennett for invaluable help. We are grateful to participants at the 2014 Marketing Science conference, the 2014 SICS conference, the 2015 EMAC conference and the DIIE Research Symposium; seminar participants at Boston University, Bocconi University, ESSEC Business School, ETH Zurich, University of California San Diego, University of California Los Angeles, University of Chicago, University of Cologne, Rochester University, Vienna University of Economics and Business and Washington University; and to Amit Bhattacharjee, Avi Goldfarb, Puneet Manchanda and Hema Yoganarasimhan for their comments. All mistakes are ours alone.

1 Introduction

The digital economy is characterized by ever-faster information flows. This has two consequences for consumers: Information diffusion happens faster, and the sheer quantity of information means that consumers' attention is a scarce resource (Falkinger, 2008). Micro-blogging services like Twitter showcase these effects. Twitter is characterized by fast-paced and short-lived information flows; new topics continually emerge and fade.

As advertising platforms, micro-blogging services such as Twitter have two unique features. First, they allow advertisers to identify new topics that are gaining interest or 'trending' rapidly across the platform (Du and Kamakura, 2012). Trending topics are increasingly displayed on social media such as Twitter and Facebook, and on news sites such as the BBC.¹ Trending topics are usually identified by proprietary algorithms based on the recency and frequency of real-time mentions of or engagements with certain media content, showcasing the new topics that large numbers of people are (and are about to be) interested in. In addition, advertisers can target advertising to the users who propagate such newly trending topics (Vaynerchuck, 2013).²

As a result, firms increasingly attempt to mesh a product with an emerging trend, a strategy broadly referred to as 'trendjacking,'³ by inserting their branded messages into social media conversations around trending topics on Twitter - either organically or through advertising. Reaching out to a single large audience who participates in and spreads these trends may be beneficial to firms - but only if the audience indeed engages with the advertising message. Indeed, firms may hope that these users consume or disseminate information about their brand or product in similar ways as they consume and disseminate information

¹<http://www.bbc.co.uk/news/blogs/trending>

²For examples of firms attempting this, see <http://www.slideshare.net/razorfishmarketing/fluent-the-razorfish-social-influence-marketing-report>, http://www.huffingtonpost.com/jeff-cann/influencer-marketing_b_3786985.html and <http://www.inkybee.com/top-50-influencer-marketing-blogs>.

³<http://www.toptensocialmedia.com/social-media-social-buzz/ten-emerging-social-media-marketing-trends-for-2014/>

on new topics or trends, because of their demonstrated interest in the trending topic.

Since advertisers increasingly try and reach out to users who engage with trending topics, digital platforms have started to introduce advertising products that target consumers who show an interest in trending topics. On the New York Times' website, advertisers can choose to target individuals who read stories currently trending on Twitter.⁴ Similarly, the Guardian News and Media introduced a programmatic advertising offering that identifies trending topics on its sites and lets advertisers target users that engage with this content in real-time, 'at social media speed.'⁵ The entire technology of a new firm called Taykey focuses on the idea that by tracking what is 'trending' in media for a particular audience, a firm can better target its audience.⁶ All these actions point to an increased interest in how marketers can capitalize on the ability to identify 'trending' content in real time on social media.

However, underlying this shift towards advertising based on trending content is the assumption that individuals who are interested in trends and propagate them are good targets for advertising. Using data from two field tests, this research examines whether that assumption holds up to empirical analysis.

Our first field test was conducted in conjunction with a large charity that ran a campaign to create awareness of homelessness around Christmas. The second field test was conducted in cooperation with a new, upcoming fashion label. In both field tests, the organization targeted ads, in the form of 'promoted tweets', to Twitter users who had posted messages containing phrases related to trending topics. As such, the field tests closely mirrored firms' practice of targeting individuals who post on trends with advertising messages. What sets our field test apart from this managerial practice is that we continue targeting ads to users

⁴The NY Times calls this approach 'Sparking Stories' - see <http://www.psfk.com/2013/02/ny-times-trending-twitter-ads.html> and <http://paidcontent.org/2013/02/19/the-nyt-is-doing-something-smart-by-using-twitter-trends-to-target-ads>

⁵<http://www.thedrum.com/news/2016/07/25/guardian-news-and-media-will-allow-advertisers-tap-trending->

⁶<http://marketingland.com/taykey-trend-programmatic-181163>

who post on the same topic when it is no longer trending. We then compare the response to identical ads of early trend propagators who post on a topic when it was an emerging trend, or 'trending', to that of users who post on the same topic when it is no longer trending. We define early trend propagators as individuals predisposed to participate in an online conversation on a topic that is about to or has just started trending on social media. In our field studies, we operationalize early trend propagators as individuals who post on Twitter using a keyword or hashtag that is 'trending' on Twitter that day.

Throughout both field tests, engagement with the ads is lowest when targeting early trend propagators and higher when targeting individuals who embrace the trend on subsequent days. We conduct a battery of robustness checks to address concerns that differences unrelated to the individual might affect advertising response and find that our results hold.

We then present suggestive evidence about what individual-level characteristic might lead these early trend propagators to be less responsive to advertising. Existing research on Twitter has suggested that many users derive utility from status rewards (Stephen and Toubia, 2013), and it is plausible that early posting related to emerging trends may be driven by a desire to provide content that leads to recognition and acclaim from followers. In line with self-determination theory (Ryan and Deci, 2000; Deci and Ryan, 1980), we suggest that early trend propagators are extrinsically motivated and particularly care about such status rewards. They use trending discussions to conspicuously present themselves as knowing about the latest trends as the rapid pace of Twitter makes being on top of the latest trends one way to signal sub-cultural capital. Early trend propagators engage with and propagate content that serves this purpose and therefore have little reason to engage with advertising.

Further results of our field tests support this behavioral explanation. In both field tests, we find that only early propagators of organic (i.e., non-firm initiated) trends are unresponsive to advertising. Early propagators of trends that were initiated and sponsored by firms show no such pattern, suggesting that the effect is indeed due to individual-level differences.

In the second field test, we additionally find that when the product advertised is more likely to be relevant to consumers, based on the type of trend they posted on, and if the message conveys the image of a unique and less commercial brand, early trend propagators are as likely as others to engage with the advertising message. However, if the message is judged as being less unique but more commercial or the product category as being less relevant to the consumer, then early trend propagators continue to show lower engagement than consumers who post on later days. Taken together, these results suggest that early propagators of organic trends are generally not willing to engage with firm-sponsored messages, unless they see an opportunity that engaging with a message may allow them to raise their profile and gain the esteem of others. We also report survey evidence that further supports that early trend propagators are highly extrinsically motivated to post on Twitter.

This paper contributes to three streams of literature. The first is a growing literature on sharing behavior on micro-blogging services. Some of this literature has examined the role of Twitter data in improving the accuracy of forecasts of cultural, consumer and stock market trends (Asur and Huberman, 2010; Goel et al., 2010; Bollen et al., 2011). Other research has examined the motivation of Twitter users to post content and their responsiveness to it. Stephen and Toubia (2013) show that Twitter users contribute content both for intrinsic and image-related reasons. Watts and Dodds (2007) present an early model of potential influential behavior in social networks, emphasizing that influence not only depends on the ability of a user to influence others, but also on the susceptibility of a user to the influence of others. Accordingly, Bakshy et al. (2011) used Twitter user data to model the effects of Twitter users in spreading a message, and find that the size of an influencer’s network does not provide clear guidance on whom to compensate. This echoes work by Cha et al. (2010) who provide evidence that having a large number of Twitter followers is not that predictive of influence on Twitter. Perhaps an explanation of this result is Weng et al. (2010)’s finding that homophily explains followership among Twitter users. Zhang et al. (2016) add to this

by demonstrating that a user is more likely to retweet a social media message that fits to their interests. While these studies provide insights into the role of network characteristics in spreading messages and in whether users are likely to spread organic messages, they do not address whether individuals who post on trending topics are likely to be receptive to advertising.

The second stream is a literature on advertising and targeting online. The early literature on display advertising focused on the performance of non-targeted ads (Manchanda et al., 2006), but more recent articles have evaluated the effectiveness of new forms of targeting, including targeting based on search queries (Ghose and Yang, 2009; Rutz and Bucklin, 2011; Athey and Ellison, 2011; Goldfarb and Tucker, 2011b), content (Goldfarb and Tucker, 2011a), time (Sahni, 2015) and previous browsing behavior (Lambrecht and Tucker, 2013). To our knowledge, no papers have evaluated whether firms benefit from targeting consumers who embrace trends online. One constant theme of this literature is the trade-off between reach and effectiveness, as targeting improves effectiveness but also reduces reach. By contrast, our paper suggests that this limitation of reach may extend further in the digital economy, as attempts to target those who themselves might lead to reach are ineffective.

The third stream is a literature on the targeting of individuals within social networks. Much of this work focuses on the interaction between the social graph and targeting strategies (Kempe et al., 2003; Hinz et al., 2011; Stonedahl et al., 2010). A few studies examine campaigns that were explicitly designed to go ‘viral.’ Toubia et al. (2009) present evidence that a couponing campaign was more effective when transmitted using a ‘viral’ strategy on social media than when using more traditional offline methods. Chen et al. (2011) show that social influence is most important in the beginning of a product’s life. Ryan and Tucker (2012) model equilibrium outcomes of targeted seeding strategies on social networks. Gong et al. (2017) find that tweets by TV companies directly boost viewing but are less effective than retweets by influentials in bringing new followers to the company. At the heart of these

latter papers is the idea of advertising to individuals who are key in spreading content across a network. By contrast, we focus on the type of individuals who embrace and propagate trends early on, and ask whether firms can influence them through commercial messages. While our empirical setting is Twitter, we focus on the response to advertising and not on how individuals interact in social networks (Stephen and Toubia, 2010; Katona et al., 2011; Yoganarasimhan, 2012).

Our work has a number of managerial implications. In general, our findings suggest that individuals who propagate trends are not particularly receptive to marketing communications. Our results are also useful for firms thinking more broadly about how to translate concepts central to the early diffusion process for offline products to the fast-paced digital world of content. We conclude the paper by highlighting how early trend propagators on social media differ from other groups identified as being key to the diffusion processes such as early adopters, market mavens, and opinion leaders.

Our paper is organized as follows. Section 2 describes the empirical setting of Twitter, and also presents descriptive evidence about the nature of early trend propagators and the diffusion pattern of trends on Twitter. Section 3 describes our first field study conducted with a homeless charity. Section 4 describes our second field study conducted with a fashion brand. Section 5 compares early trend propagators with earlier studied individuals who are key in the diffusion process such as early adopters and market mavens and concludes.

2 Empirical Setting

2.1 Why are Trends and Early Trend Propagators Important for Marketers?

In the past, firms trying to disseminate communications were limited to purchasing ads on television and print media or using public relations to entice journalists to write about their product. Now, digital and social media buzz allow firms to attract publicity by involving consumers in spreading information on a large scale. Twitter, the platform at the center

of our study, has streamlined this process by allowing users to unite conversations around topics through ‘hashtags’. These hashtags can be used to identify new and evolving trends or topics of conversation.

Twitter itself promotes and emphasizes the importance of these trends to users. Figure 1 shows how the left-hand panel of the Twitter homepage presents ‘trends.’ While the exact algorithm Twitter uses to identify trends is proprietary, trends are detected partly based on the frequency and momentum of hashtagged posts, allowing Twitter to identify trends shortly after the topic becomes relevant to an unusually large audience. Research in computer science suggests that it takes Twitter only hours to recognize new trends (Chen et al., 2013; Kong et al., 2014).

We define early trend propagators as individuals predisposed to participate in an online conversation on a topic that is about to or has just started trending on social media. Early trend propagators play a pivotal role in disseminating information. Reaching out to social media users who are participating in and spreading these trends may be important to firms. Indeed, firms may hope that by linking their advertising message to a trend, these users may consume or disseminate information about their brand or product in similar ways as they consume and disseminate information on new topics or trends, because of their demonstrated interest in the trending topic.

This emphasis on people who engage with topics just as these are ‘trending’ echoes marketing practice. Big brands such as L’Oreal and Marriott have recently brought content marketers in-house to enable them to participate in trending topics quickly.⁷ Indeed, it has become an increasingly common practice among marketers to seed information with users who show an interest in trending topics in the hope they will engage with it. For example, the

⁷<http://digiday.com/brands/loreal-bringing-fast-twitch-content-house/>. The commercial director of That Lot, a company that specializes in writing for the Twitter feeds of large brands, states that: “A phrase I use a lot is ‘advertising at the speed of culture.’ Without doubt, topicality is the prevailing wind of social (especially on Twitter), and being on top of the latest trends is paramount.”

birth of Princess Charlotte in the UK sparked a flurry of messages by advertisers promoting their brand around the hashtag #royalbaby. Figure 2 displays some examples of advertising messages using the hashtag #royalbaby on the day Princess Charlotte was born.⁸ The advertising messages displayed in Figure 2 were displayed at no cost to the advertiser (other than the cost of employing marketers who would identify the occasion and create these ads) to Twitter users who read messages containing the hashtag #royalbaby.

However, marketers are also investing into paid media targeted towards individuals interested in trending topics. Figure 3 displays two examples of such tweets both marked with ‘Promoted’, indicating that the advertiser paid Twitter for an advertising impression targeted towards a specific set of individuals. In both instances the hashtag contained in the message was trending when the ad was displayed and the user who received the messages had just used this hashtag in their tweet (#DowntonAbbey refers to the start of the last season of a popular TV series and the ad is for the jewelry manufacturer DeBeers; #LFW refers to the start of London Fashion Week and the ad is for the employment website Monster). Other large brands such as Pepsi and Asda have used similar approaches in their advertising.⁹

2.2 Early Trend Propagators on Twitter

2.2.1 API Data Collection

We want to establish whether being an early trend propagator is indeed a relatively stable individual difference. One disadvantage of the data that Twitter provides to advertisers and that forms the body of our empirical analyses is that it does not include information on individual users exposed to the ads. Therefore, we collect an individual-level data set that is separate from the field experiments that constitute the main analysis of this paper.

On sixty consecutive days from May 25 to July 23 2016, we identified the ten organic

⁸See also <http://www.adweek.com/news/technology/here-are-top-10-branded-tweets-about-royal-baby-164496>, <http://www.adweek.com/news/advertising-branding/how-nissan-beat-everyone-royal-baby-tweet-just-seven-minutes-160001>

⁹See <https://business.twitter.com/success-stories/pepsi-max> and <https://business.twitter.com/success-stories/asda>.

trends listed on Twitter’s US website at 9 am EST and used the Twitter streaming API to collect all tweets on these trends as well as the user names of the individuals posting on this day and on the six subsequent days. We counted as the first day the first 24 hours after we identified the trends and started data collection. We collected a total of 35,574,598 tweets.¹⁰

Since recent research has demonstrated the importance of user-content fit in social media engagement (Zhang et al., 2016), we expect early trend propagators to be interested in particular topics, rather than all trending topics. Therefore, we manually categorize the 600 trends in our sample. Since the meaning of hashtags or keywords is not always obvious, we first asked a research assistant who was blind to the purpose of the study to research each of the 600 trending topics and develop a brief description of each trend. We then checked the validity of these descriptions. Next, we used three independent raters to classify all trends based on their descriptions into one of six major categories: general news (e.g. No Swimming which referred to a child being killed by an alligator in Disneyland), politics (e.g., #ClintonPassword which referred to Hillary Clinton’s email passwords), holiday (e.g. #MemorialDay2016), pop culture (e.g. #DemiOnGMA in reference to the appearance of Demi Lovato on the show Good Morning America), sport (e.g. #Formula1 which referred to a Formula 1 race that day) and Twitter-specific trends (comprising topics which recur on a weekly basis such as #MondayMotivation and topics that emerge only on Twitter and are unique to its culture such as #honestyhour, where Twitter users share their honest feelings

¹⁰From this data, we removed observations that are likely to be spam (Antonakaki et al., 2015). Reflecting the importance of trends in the Twitter ecosystem, various bots post (spam) to trending topics in order to attract interest. Spam posts are normally made via Twitter’s API and users do not log into the Twitter website, making them unlikely trend propagators. We identified and removed all accounts with more than ten posts on a single trend on a single day (12,263,002 tweets) or more than 50 posts on the same trend during the sixty-day window (1,433 tweets). For ease of analysis, we removed the 2.3% of tweets that refer to more than one trend (829,934 tweets). Lastly, we drop users with a very high number of followers (top 0.5% in terms of followers, resulting in the removal of accounts with more than 36,174 followers). While we are unable to ascertain on the level of individual accounts whether the number of followers reported are genuine, this allows us to eliminate observations where the number of followers seems unusually high and so less likely to be genuine. The main results are not sensitive to eliminating users with a high number of followers, or to a different cut-off.

or beliefs).¹¹

Note though that despite collecting over 35 million tweets, we are still far off from capturing the full activity during this period by all early trend propagators because we only capture tweets on topics that were trending at 9am EST. Other trends will have emerged during the day and already been replaced with new trends by the time we started collecting the next set of trends at 9am EST on the subsequent day. We also do not capture tweets on topics that were not trending.¹² Our results therefore have to be seen in light of this limitation of the data we were able to collect and are indicative rather than fully conclusive.

2.2.2 API Data Analysis

We analyze this data in two steps. First, we aim to establish whether being an early trend propagator is a relatively stable individual difference. For any user in our data who posted at least once during the first 30 days in our data (period 1) and at least once during the subsequent 30 days in our data (period 2), we compute for each period the share of their tweets that were made on the day a trend emerged. Table 1 demonstrates that 52.9% of users (i.e. those on the diagonal) have a similar share of tweets on the day a trend emerged across periods 1 and 2. A chi-square test rejects the null hypothesis that the share of tweets made on the day a topic is trending in period 1 is independent of that in period 2 ($p < 0.001$). This result suggests that early trend propagators indeed have a fairly stable tendency to post on trends the day they emerge.

Second, we analyze persistence in category-specific behavior. We wish to examine whether users who posted on a trend on the day it emerged are likely to post again later on a trend within a similar category. We focus our analysis on the 1,737,458 users in our data who post at least once in our data during period 1 and at least once during period 2, independently

¹¹We omit from the analyses tweets referring to 18 trends that we were unable to classify, such as #ed-campldr, which referred to a leadership conference for educators held in Hawaii.

¹²Note also that some users may have been less active during the summer period if they were on holidays and our servers were on some occasions temporarily disconnected from the Twitter API. In these instances we were unable to record the full universe of tweets on the keyword or hashtag.

of whether these posts were made on the day the trend emerged or on a later day. 2.8% of tweets by these users were about general news, 19.5% about a holiday, 42.5% about politics, 11.8% about pop culture, 10.4% about sports and 13.2% about a Twitter-specific topic.¹³

Here, we operationalize an early trend propagator in a category as someone who posted on the first day of a trend in a particular category. To investigate category persistence in tweeting behavior, we estimate a logit specification where the dependent variable captures whether a user ever tweeted on the first day of a trend in a particular category during period 2. The independent variables capture in which categories, during period 1, they tweeted on the first day of a trend. Column (1) of Table 2 demonstrates that a user who posted a first-day tweet on a general-news trend during period 1 is significantly more likely to make a first-day post on a trend within the same category in period 2. This positive coefficient for the general news category is markedly higher than for any of the other categories. Columns (2) to (6) present a similar picture for the remaining categories. Throughout, the same-category coefficient is significantly higher than that of other categories. The fact that coefficients relating to activity in categories other than that captured by the dependent variable can be positive, suggests that while users have particular focal interests, they can be active in other categories as well. Table 3 reports the results of a similar specification, this time using a seemingly unrelated regression of linear probability models. The results confirm those in Table 2.¹⁴ Overall, these results indicate the potential for a distinct group of early trend propagators on Twitter who are likely to post early on trending topics within a category of their interest.

¹³This split is not necessarily representative of the overall by-category volume of tweets on Twitter since Twitter activity with current events.

¹⁴We find similar results, albeit with larger effect sizes, if we limit the analyses in Tables 2 and 3 to individuals who posted on day 1 in both periods 1 and 2.

2.2.3 Twitter-related Characteristics

We then compare several Twitter network characteristics of early trend propagators who posted on a day a trend emerged in both periods 1 and 2 with those of users who posted only later on the same trends. Specifically, these users posted during periods 1 and 2, but either never posted on the day a trend emerged or did so only in one of the two periods. Table 4 illustrates that early trend propagators have somewhat more followers. They also have more friends (i.e., accounts they follow), and have posted more tweets in the past even though their accounts have been somewhat more recently created. In sum, they appear to be more active. In Section 5.1 we discuss how early trend propagators relate to other marketing concepts such as early adopters, market mavens, and opinion leaders.

Table 4 indicates a large number of early trend propagators relative to the number of users who posted only later on the trends in our sample. However, it is important to remember that this is conditional on having used a trending keyword or hashtag in a post (as we do not collect data on keywords or hashtags that were not trending) and that early trend propagators represent only a relatively small share of the total number of about 65 million monthly Twitter users.¹⁵

2.3 Diffusion of Trends on Twitter

Next, we document the aggregate pattern by which trends diffuse on Twitter. To this end, we use a different sliver of the previously described data. For hashtags or keywords that were trending on the first 14 days when we identified trends on Twitter, we collected the tweets using those hashtags or keywords over a period of 30 days (instead of just 7 days). Figure 4 shows the diffusion patterns for these trends, based on the subset of 4,204,825 tweets which used the trending keywords or hashtags.¹⁶ Additionally, in Figure 5 we only look at data for

¹⁵Statistics based on <http://www.statista.com/statistics/274564/monthly-active-twitter-users-in-the-united-states/>

¹⁶We exclude the weekly recurring hashtags such as #WednesdayWisdom.

the top 5 trends on any one day.

Throughout, Figures 4 to 5 suggest that the volume of tweets for trending topics is initially high and then sharply declines. As such, interest in trends does not follow the sales cycle implied by the initial curvature of the s-curve (Rogers, 1962; Bass, 1969) where there are slow initial sales followed by a period of higher sales but instead suggests that cumulative adoption slows down rather than speeds up. This pattern is similar to the diffusion pattern of digital information or entertainment goods such as video games where the sales of new copies sharply declines after the release week (Ishihara and Ching, 2012). It also echoes an earlier literature such as Rangaswamy and Gupta (2000) and Bughin (2003), which documented that digital goods exhibit faster diffusion than traditional categories of goods when using calibrations such as the Bass model.

2.4 Advertising on Twitter

On Twitter, advertisers can target users in different ways, including based on a hashtag or keyword mentioned in a post, interests or demographics (gender, geography). The focus of our field studies is targeting trend propagators, i.e., individuals who mentioned in a post hashtags or keywords that were listed as trending on Twitter. In each study, we identify on multiple days trends on the Twitter website and then target promoted tweets, that is advertising messages, to users who use a trending hashtag or keyword in their post. We target these messages to users who use the hashtag or keyword on the day it was identified as a trend and on subsequent days. In each field test, we use several advertising messages and refer to each targeting and message combination as a ‘campaign.’ Table 5 summarizes the similarities and differences across the two field tests which we will explain in detail in the respective sections.

Twitter distinguishes between trends which emerge organically and ‘sponsored trends’ where a firm pays for a hashtagged term to be first on the list of trends as advertising,

independently of its appeal to users. As such, it does not reflect the popularity of the topic. Figure 1 shows a sponsored trend. Here, Microsoft promotes its software and charitable contributions. Below this sponsored trend are the organic trends of the day.

Content targeting on Twitter is based on recency. This means that users are targeted by a promoted tweet based on whether they used the targeted words or phrases in one of their tweets within the last 24 hours.¹⁷ In theory, once the target audience that posted on this topic within the past 24 hours is exhausted, Twitter extends this moving window backwards to cover users who had posted on the topic at earlier dates. However, given the relatively small size of the campaigns in our sample (on average 411 daily impressions per campaign in field test 1 and 935 in field test 2), compared to the large size of conversations about the top trends that were targeted (on average about 100,000 postings within the preceding 24 hours in both field tests 1 and 2), this seems unlikely to have happened.

On the day a topic is identified as a ‘top trend’ the advertising campaign targets early trend propagators, and thereafter it targets late propagators. The 24-hour targeting window means that our approach includes as early trend propagators both the users who were posting about the trend at the time that Twitter identified it as an emerging trend and users propagating the trend immediately after it was listed. This behavior contrasts with people tweeting about the topic on a following day when it is not trending anymore. Therefore, the key variation in our data is how early people started talking about this new (and popular) topic on Twitter.

After determining targeting criteria and message, an advertiser bids on impressions to individual Twitter users by submitting a maximum price per engagement. In both studies this maximum bid was £1.00. The auction mechanism used to allocate ad impressions resembles a second-price auction where the bidder pays an amount similar to the second

¹⁷Only users that access the Twitter website and not accounts that exclusively post through the Twitter API, such as spam accounts, are eligible targets for advertising through promoted tweets.

highest price bid at the auction for that particular impression. The promoted tweet is then displayed just once at or near the top of the user’s timeline.¹⁸

In the next two Sections, we turn to the core of the empirical analysis and examine in two field studies how early trend propagators respond to advertising messages.

3 Field Test 1: Charity

3.1 Campaign Setup

The field test was designed and implemented in December 2013 in cooperation with a large UK charity. The aim was to attract publicity for their annual Christmas Appeal to help the homeless. The test relating to targeting trends that we focus on in this paper was part of a broader exploration of different message conditions and techniques on Twitter that we report fully in the Web Appendix.

For 19 days at the start of December 2013 at 8am, a person in London examined the top trends on the website of Twitter UK.¹⁹ In selecting the new trend to be targeted on any specific day, an exogenously determined procedure was implemented: On the first day, the charity targeted the first trend independently of whether or not it was organic or sponsored. On the second day, the charity targeted the first trend if that was organic but the second trend if the first trend was sponsored. The third day followed the pattern of day one and the fourth day the pattern of day two, and so on. We use this quasi-randomization between targeting organic and sponsored trends to ensure we have a sufficient number of both types of trends, since this will help shed light later on the behavioral mechanism. The charity then targeted advertising messages to individuals who had posted on these trends, the day the trend emerged and for three subsequent days (Table 5). Additional data from topsy.com, a website that tracks Twitter data and makes it searchable, documents that on average there

¹⁸<https://business.twitter.com/help/what-are-promoted-tweets?lang=en&location=emea>

¹⁹The individual was logged into their Twitter account. Due to the slight customization of trends to users, the individual recording the trends in study 2 logged out of their Twitter account.

were close to 5,600 postings containing a trend targeted by the charity within one hour and more than 100,000 on-topic postings in the 24 hours before campaign launch.

The trends targeted covered a wide range of topics, from pop culture, to music, to public affairs (Table 6). This range indicates that organic Twitter trends indeed capture momentary spikes of interest by a wide variety of users, rather than being exclusively driven by a specific sub-group of users such as journalists. In addition, they illustrate why trends emerge suddenly. For example, offering blessings upon the death of Nelson Mandela (`#RIPNelsonMandela`) could only practically emerge as a trend on the day of his death. Similarly, the announcement about Rebecca Black’s single (`#RebeccaBlack`) could only become a trend the day it was announced, and people could only tweet about the Apollo Theatre’s collapsed roof (`#ApolloTheatre`) after this incident had occurred.

Each trend was targeted by 16 wording variants for the sponsored tweet that we control for but do not focus on in this research (see Web Appendix). Each campaign was set up with a constant daily budget and a maximum bid per engagement of £1.

3.2 Data

To advertisers, Twitter reports daily campaign performance. As Table 7 summarizes, our data contain a total of 1216 campaign-day observations (19 trends x 16 message variations x 4 days per campaign), with a mean of 414 daily impressions. This leads to a large number of impressions and a sample size of over 2 million views. This large sample size is in line with studies that document the need for large samples to precisely measure online advertising effectiveness (Lewis and Reiley, 2014; Lewis et al., 2011).

To protect user privacy, Twitter does not disclose details about users exposed to the promoted messages, their social networks or number of followers. Twitter refers to any measurable response to a promoted message as an ‘engagement.’ Here, a campaign-day has on average 3.81 engagements. Engagements can be clicks (mostly on the ad or the charity’s

URL)²⁰ or ‘retweets’.²¹ A ‘retweet’ occurs when a Twitter user deliberately rebroadcasts the message to their followers. In our study, most engagements are clicks, with only 0.16 retweets per campaign day. Clicks are important since they signal a user’s interest in the message and relate to brand awareness even if the message is not immediately disseminated.

Retweets are a subset of behaviors that can result from clicks and generate further impressions for free. Since users may also ‘manually retweet’ a message by copying a tweet and pasting it into their own tweet, which would be counted as a click but not a retweet, Twitter may potentially underreport retweets. Because advertisers worry about both clicks and retweets and since the number of retweets is potentially imprecise, the majority of our analyses focuses on engagement as outcome variable, subsuming both clicks and retweets.

On average, the charity spent £1.64 per campaign-day, that is a cost per engagement of £0.35 (around 50 cents). This reflects that the charity did not always pay its maximum bid and also reflects a 20% increase in advertising spend granted to charities.

3.3 Model-Free Evidence

Figure 6 presents model-free evidence for our main finding. It reports the success rate of each campaign, using engagements relative to the number of impressions by number of days passed since the trend peaked. It shows that engagement is significantly lower when ads are targeted at early trend propagators, that is users who posted on a trend the day it emerged, than on any of the following days, and increases as days pass. On day 4, the rate is more than twice as high as on day 1. Figures 7 and 8 show that this pattern holds when splitting the data into clicks and retweets.

This analysis does not control for differences across trends or days. For example, as Christmas approaches, responsiveness to charitable ads might increase. Likewise, the model-

²⁰Twitter does not separate out these two types of actions in its reports to advertisers.

²¹Twitter also reports the decision to ‘follow’ an advertiser’s account in order to receive future messages. This metric is relevant for campaigns promoting a Twitter account rather than a message as was the case in our studies. Our campaigns attracted only 99 follows as a result. There were too few observations to obtain precise results or to estimate the full complement of main and fixed effects in a regression analysis.

free analysis equally weighs each campaign and day and so does not account for different numbers of impressions across campaigns. We next use an econometric framework to capture the effectiveness of promoted tweets when targeting early trend propagators.

3.4 Main Results

Our analysis exploits the fact that we have campaign-day level data of consumer responses to a campaign targeted at a particular trend over multiple days. The analysis measures the relative uplift of the campaign on subsequent days relative to the first day. Without advertising there would have been no clicks or retweets, since the account used for the experiment had only a handful of followers and little history of advertising or promotions.

The standardized design of the field test makes our empirical analysis relatively straightforward. Since our data is at the campaign-day level, we build our main empirical specification at the individual level and then use aggregated estimation techniques to reflect the fact that we only have campaign-day level data.

For an individual who engages on day i with a trend j and is exposed to message k on date t , the likelihood of engaging with the promoted tweet is a function of:

$$Engagement_{ijkt} = \beta_1 DaysSinceTrend_{ijt} + \beta_2 Targeting_j + \beta_3 Message_k + \delta_t + \epsilon_{jk} \quad (1)$$

Here, β_1 captures the key coefficient of interest, which is the extent to which someone posting on a trend is an early propagator of this trend, measured in terms of days elapsed since the trend was listed as a top trend on Twitter. β_2 is a vector of coefficients for each of the different trends that were targeted, to control for differences in behavior associated with the kind of person who would tweet on that trend. This means that β_1 captures the differences in behavior of the individuals who had the same propensity to use a certain set of words or phrases in their tweets. β_3 is a vector of coefficients that capture the effect of the sixteen different message conditions to control for differences in response to slight differences

in the wording of messages. Last, the vector δ_t is a series of day fixed effects that control for heterogeneity in baseline behavior over time. Initially, we assume that β_1 is constant across trends. But in our later analysis we estimate this equation separately by different groups of trends and so capture how users who post on different types of trends early, respectively late, respond differently to advertising messages. We cluster standard errors at the campaign level in accordance with the simulation results presented by Bertrand et al. (2004).

Twitter reports data by grouping all successes and failures on each day. This means that while the consumer’s decision is a binary choice, our data is aggregated across consumers, and we observe a number of successes (engagements) out of a number of trials (impressions) for each campaign-day. Since the different trends we targeted had varying degrees of use, our campaigns had different numbers of impressions. For example, there was one impression for a campaign targeting #AirportCommission on days 1-3 but 3,334 impressions for a campaign targeting #Spoty on day three. Unlike other research using field tests for online advertising, where in a straightforward ordinary least squares approach the click-through rate is the dependent variable, we need to account for such differences in daily impressions. To see why this is important, imagine two campaigns, one which received 100 impressions and the other which received 10,000 impressions, where both received zero clicks. Simply using the click-through rate as a dependent variable would effectively treat these instances as the same, though they convey very different information. As a result, we estimate an aggregate logit model using maximum likelihood (Flath and Leonard, 1979).

Aggregate discrete choice models need to account for both heterogeneity and endogeneity (Chintagunta, 2001). Here, fixed effects account for heterogeneity across different campaigns and days. We tackle the question of endogeneity by focusing on variation which occurs within the set of individuals who embrace the same trend - we study the variation in timing of when the Twitter user started posting about a particular trend.

Let F denote the logistic likelihood function. Due to the aggregate nature of the data

provided by Twitter, which does not have user-level variables, all individuals i exposed to a particular campaign j with message m on day t have the same vector of x control variables. The likelihood of observing each observation of the sum of positive engagements as a function of the sum of impressions in the data is then:

$$F(\beta x)^s \{1 - F(\beta x)\}^{r-s} \quad (2)$$

where s is the number of engagements and r is the population exposed to the messages.

Table 8 shows the initial results from our estimation, controlling for trend and date fixed effects. Column (1) indicates that the likelihood for someone to click on a tweet increases the more time has elapsed since the trend emerged. Put differently, early trend propagators, that is users that are targeted on the day a trend emerges, are least inclined to engage with a promoted tweet, and users that are targeted three days after a trend emerges are most likely to engage with the promoted message. The effect holds when controlling for message fixed effects in Column (2). Column (3) confirms the robustness of the effect when the effectiveness of advertising varies by days passed since the trend emerged. This more flexible specification means we do not force a linear time trend on responsiveness. Instead, we take the behavior on day 4, the final day of the campaign and the day furthest from the trend emerging, as a baseline. Therefore all estimates are relative to the slowest trend propagators in our data. Since this specification echoes the non-linear increase in responsiveness over time of Figure 6, we subsequently emphasize this more flexible functional form.

In sum, our empirical specification shows that early trend propagators are less likely to respond to promoted messages than later propagators.

3.5 Robustness

We conduct a battery of robustness checks to make sure our results hold across different controls, functional forms, samples and dependent variables. One first concern is that, despite the use of the aggregate-logit model, our results could be driven by outliers because on some days, campaigns had very few impressions. We re-estimate our model focusing only on campaigns for which Twitter initially categorized the trend as having an above-average number of tweets. Our findings are robust (Column (4), Table 8). Second, in Column (5) we confirm that the results hold for an audience that conversed on a trend related to that of the promoted tweet (`#gooddeeds`, `#foodbankdebate`). As these were spread out, we are unable to identify day fixed effects and instead use week fixed effects. Still, the results hold.

Third, our results could reflect a competitive effect. If early trend propagators are more attractive to advertisers and receive more sponsored ads, then our focal ads may attract less attention. To investigate whether this drives our results, we control for the average amount the charity spent per engagement in Column (6) of Table 8. If competition is driving our results, the charity would pay more on highly competitive days of a campaign due to the nature of the auction for keywords and an increased spend would act as a proxy for competition for that keyword. However, our results hold.

Fourth, we check robustness to other functional forms to model the aggregate click and impression data. The results are similar when using maximum likelihood grouped estimation with a probit functional form (Column (7)). As discussed by Flath and Leonard (1979), one can estimate a logit on aggregate data using maximum likelihood or weighted least squares. Though their evidence tends to favor the maximum likelihood approach we focus on, we checked robustness to weighted least squares. Again, the results are similar (Column (8)).

Our results in Table 8 focus on engagement as independent variable, the key outcome variable for Twitter since it determines how much an advertiser is charged for a promoted

tweet. Table 9 displays the results separately for clicks and retweets for both the linear and non-parametric specification of the key explanatory variable. The results hold when the dependent variable captures only clicks (Columns (1) and (3)). In line with our previous results, early trend propagators appear to have less interest in exploring and pursuing it than later trend propagators. Additionally, early trend propagators are less likely than later propagators to retweet a promoted message to their social network (Columns (2) and (4)).

3.6 Mechanism

Our results so far show that early trend propagators are a distinct group of Twitter users who are more likely than others to post on trending topics (see Section 2.2), but who are also less responsive to advertising than consumers who post on the same trend at a later point in time. We now explore a possible behavioral mechanism underlying our results.

We compare the behavior of users who embraced organic trends and of those who embraced sponsored trends. This comparison is possible because of the quasi-randomized algorithm that the charity used in deciding whether to target a sponsored or organic trend. Early propagators of sponsored trends, similarly to early propagators of organic trends, post on topics the day these topics are listed as trends, but sponsored trends are typically *not* newsworthy items and their listing does not relate to their popularity.

In Table 10, Column (1) reports the results for sponsored trends and Column (2) for organic trends.²² A comparison across the two columns indicates that early propagators of organic trends indeed respond negatively to promoted messages. By contrast, Twitter users who embraced firm-sponsored trends early on do not show this pattern. This result shows that the effect is related to individual-level characteristics of early propagators of organic trends and not to a mechanical reflection of the way Twitter serves advertising. It suggests that there is a key distinction between Twitter users who engage with firm-sponsored trends

²²We stratify rather than add an interaction term, due to the difficulty of interpreting interaction terms in non-linear models (Ai and Norton, 2003).

or messages, and those who use Twitter to embrace organic trends early on.

To explain why early propagators of organic trends are less responsive to firm-sponsored messages, it is useful to consider what may motivate someone to become an early trend propagator on Twitter. Earlier research has documented that social media users post content because they derive utility from status or prestige associated with the activity and desire to acquire a larger followership (Stephen and Toubia, 2013). This is in line with self-determination theory (Ryan and Deci, 2000; Deci and Ryan, 1980), which identifies the extrinsic motivation to affirm one’s self-worth as one of the basic motivations of human action. On Twitter, feelings of self-worth stem from having an audience who values the user’s postings. The number of followers, likes, favorites and retweets are prominently displayed and send a strong signal about the Twitter users’ value to their audience, triggering extrinsic motivation in those susceptible to it. We suggest that early trend propagators may have a specific regulatory style characterized by a ‘control orientation,’ which means that they are particularly susceptible to the pressures of extrinsic motivation (Ryan, 1982; Deci and Ryan, 1985). To ensure that their postings are valued by their followers and to signal attractiveness as a content provider, early trend propagators thus carefully curate which content to engage with and post. In a fast-moving social network such as Twitter (Arvidsson and Caliandro, 2016), new content on trends - topics considered relevant by a large number of people - attracts the attention and potential acclaim of others.

Previous research has documented that people engage in audience tuning (Berger, 2014; Krauss and Fussell, 1991) by tailoring what they share to the interests of their audience, particularly when this audience and their reactions are made salient as on Twitter. Early trend propagators may attach particular importance to consuming and posting timely information of interest to potential followers to signal they are on top of the latest trends. As a result, unless advertising content is helpful for audience tuning, it is uninteresting for such users. This conceptualization is consistent with the results in Table 10 which document a difference

in responses to advertising between Twitter users who have engaged with a sponsored trend, and are likely to be more open to advertising, rather than with a non-sponsored trend.

To provide support for this conceptualization, which we cannot test directly in our field data,²³ we collect survey data from 251 US Twitter users via Amazon’s Mechanical Turk.²⁴ Since Twitter does not share details of individual users who were exposed to ads to advertisers, this sample of Mechanical Turk users is different and at best a selected subsample of users in the Twitter advertising data which is the focus of our main empirical analyses.

The survey adapted existing scales to measure early engagement with trends (Goldsmith and Hofacker, 1991) and extrinsic motivation (Kankanhalli et al., 2005). To measure intrinsic motivation, we took inspiration from Mathwick and Rigdon (2004) and adapted items to the Twitter context. All items were evaluated with 5-point Likert scales (‘strongly disagree’ to ‘strongly agree’). In addition, general Twitter involvement was measured on a semantic differential scale, based on Mathwick and Rigdon (2004). All scales display good psychometric properties (see Web Appendix for details).

Rather than splitting the survey respondents into early trend propagators and other segments, we use a continuous measure of early engagement with trends and correlate it with other continuous variable scales. The results show that early engagement with trends significantly correlates with extrinsic motivation (correlation 0.412, $p < 0.001$), but not intrinsic motivation (correlation -0.043, $p = 0.500$). This result suggests that early trend propagators indeed care about the image they project to their audience and are mostly motivated by status rewards. Furthermore, early engagement with trends significantly correlates with general Twitter involvement (correlation 0.394, $p < 0.001$), indicating that early trend propagators generally pay close attention to their Twitter activity.

²³While research such as Dubé et al. (2017) was able to vary incentives to test for motivation, this is more difficult in our context which involves advertising and where users are not identifiable to the advertiser.

²⁴At the start of the survey, we ask participants “Do you have a Twitter account?” We then only conduct the survey for participants who respond with yes. All participants are paid the full amount regardless of their response to the screening question.

Finally, we present respondents with seven strategies to build a larger following, and asked them to pick their top three. We find that these are to 'post on trending topics or trending hashtags' (n=156), 'post on current topics that others care about' (n=141); and 'post frequently' (n=120). These results confirm that posting on trending topics is indeed valued by Twitter users and leads to status rewards through acclaim from followers.

Overall, the first field test demonstrates that early trend propagators are less likely to engage with promoted tweets than individuals who post on the same trend later. It suggests that they may be extrinsically motivated by a self-presentation goal that influences which social media content they engage with. The survey data we collected support this conceptualization. However, a number of concerns related to the specific setup of this first field test remain. First, the results might be specific to the charitable context: If early trend propagators care particularly about themselves rather than about others, they may be less responsive to messages about helping the homeless. Second, early trend propagators might not have engaged with the message because the charity and the topic of homelessness did not feel relevant to them. While our results hold for trends related to charitable actions, it would be reassuring to find that they continue to hold when an advertiser targets trends more closely related to their offering. Third, it would be useful to test the assertion that early trend propagators filter social media content with a view to how they can satisfy their extrinsic motivation to attract the attention and acclaim of others. We next conduct a second field test with a commercial brand.

4 Field Test 2: Fashion Firm

We aim to replicate our results for a non-charitable enterprise where the advertised product category is more closely related to the targeted trend. Many trends on Twitter are related to pop culture, TV and music. Leading consumer goods brands focus on such topics in their Twitter campaigns. For example, PepsiCo partnered with Beyoncé for a Twitter campaign.

Fashion brands seem to have similar aims, so in our second field study, we partner with a relatively new, up and coming fashion label.

4.1 Campaign Setup

The field test was carried out in the US over 26 days in June and July 2014 jointly with a small UK-based fashion label unknown in the US, but which sell products to US consumers. The aim of the Twitter advertising campaign was to attract attention to the fashion label’s website.

The setup was similar to that of field test 1, with the exception that we target not only the first but all ten trends listed (see Table 5). Over a period of 20 days, we identify every day at 9 am Eastern Summer Time the ten trends listed on Twitter. The fashion firm then targets advertising to consumers posting on any of these trends for seven successive days.

As before, trends vary considerably from day to day. Out of a total of 200 trends, only nine appeared more than once.²⁵ Three independent raters blind to the study’s hypothesis classified all trends into whether they broadly related to pop culture (including TV or music, such as BET Awards, The Walking Dead, #FinallyA5sosAlbum), meaning that the ad by the fashion firm was more likely to be broadly related to these topics, or not (including trends about sport events, politics or weather events such as #BrazilvsGermany, Gaza, #HurricaneArthur). Our data include 77 related and 123 unrelated trends.

We additionally identify and target sponsored trends. Since trends displayed on Twitter can vary very slightly depending on the individual logged in, we logged out of Twitter to collect sponsored trends. As a result, no sponsored trends were displayed. To identify sponsored trends, we separately logged into a US Twitter account. Sponsored trends are not displayed to an individual user every day - presumably to not ‘spam’ them with advertising. During the time period of our data there were nine sponsored trends.

We also target, over the entire period of the experiment, advertising messages to Twitter

²⁵These were across a range of topics (e.g., #GOT7COMEBACK, 4th of July, Israel, Starbucks).

users who use the hashtag or keyword '#shopping' or 'shopping' in their post. This gives us a baseline of Twitter usage to allow us to control for time trends that relate to the intensity of Twitter usage by users who are broadly interested in the category.

If indeed early trend propagators respond more negatively because they are extrinsically motivated and have little reason to respond to content that does not further their goal of self-presentation, then a message that is framed as more unique may more successfully engage them. Such a message would be more likely to attract the interest of others. Similarly, a message that is framed as less commercial would be more likely to signal attractiveness as a content provider. To explore this, we used two different advertising messages, phrased within the limits of what the firm considered appropriate. The first message is designed to convey the image of a more unique brand ('Discover an original Brixton-based fashion brand: [brand name] - where crazy is the new normal [URL]'). The second message is designed to convey the image of a more commercial firm ('Discover a new global fashion brand: [brand name] - where crazy is the new normal [URL]'). We tested these messages with a sample of 196 US participants at Amazon's Mechanical Turk and find that the 'Brixton'-message indeed ranks higher on being unique (3.22 vs. 2.71, $p < 0.001$) and is perceived as less commercial (3.02 vs. 3.63, $p < 0.001$) (see Web Appendix for details). We likewise tested whether participants felt the message captured their attention or was surprising and found neither of these variables to be significantly different.

In sum, we target users posting on a total of 200 organic trends (10 trends identified on each of 20 days), 9 sponsored trends and 'shopping', each with two different messages. This makes for a total of 420 campaigns. All trend-based campaigns run over 7 days. The campaigns targeting 'shopping' run continuously throughout the 26 days of the field study as this was not a trend, but instead just data we collect for a baseline.

All campaigns had the same daily budget. The maximum bid per engagement was £1. Table 11 summarizes basic descriptives. The campaigns had an average of 935 impressions

and 8 engagements per day, the great majority of which were clicks. The firm did not always pay its maximum bid and the average cost per engagement was £0.92.

4.2 Model-Free Evidence

Figure 9 presents initial model-free evidence for our main finding, focusing on data from campaigns targeted towards users who posted on a top ten trend. It reports the success rate of each campaign, as measured by the number of engagements relative to the number of impressions, by how many days have passed since the trend peaked on Twitter. It illustrates that the engagement rate consistently increases over the first four days of a campaign, in line with our results from field test 1, after which it levels off. This pattern is similar when splitting up the data by clicks and retweets in Figures 10 and 11.

4.3 Main Results

Columns (1) to (3) of Table 12 focus on data for users who posted on one of the top ten trends on any day in our data. Column (1) shows that users are more likely to engage with a tweet as days elapse since the trend emerged. Column (2) confirms that the results hold when controlling for a message fixed effect and spend per engagement. In Column (3) we instead allow the effectiveness of advertising to vary by days elapsed, using engagement on the day the trend emerged as a baseline. We find that advertising effectiveness increases, though the pattern does not fully reflect the monotonic trend as in field test 1.

4.4 Robustness

Columns (1) to (3) use the last day of a campaign as the baseline for engagement. In Column (4) we use as an alternative baseline the engagement of campaigns that target ‘shopping’ on any day in our data. As explained earlier, the idea of this baseline is that by targeting campaigns towards people whose tweet contains the phrase ‘shopping’ we can control for changes over time in posting behavior which may provide an alternative explanation. Recall that shopping itself is not a trend, which is why we do not have a metric of ‘Days Since

Emerg'd' for this data and instead use it as a baseline. The results hold. They indicate that early trend propagators are less responsive to firm-sponsored messages than individuals who post on generic topics, such as 'shopping,' while individuals who embrace trends later on respond similarly to individuals who post on generic topics.

Column (5) reports the results for sponsored trends only. Similar to field test 1, we do not find an increase in engagement over time as we target users who post on sponsored trends. One concern is that some other difference between sponsored and organic trends, such as whether it is easy to foresee they would arise or whether they are more easily understood out of context, is driving the effect. To ensure that this is not the case, we had all trends rated by ease of comprehension or the degree to which such a trend was easy to foresee. We stratify the results across these variables and find that the results hold.

Next, to rule out that the results are an artifact of variation in the number of postings on a trend over time, we use data we collected on a daily level on the size of each trend on Twitter as measured by the number of postings within the past 24 hours on topsy.com. Column (1) in Table 13 illustrates that the results are robust to the inclusion of this variable.

Similarly, the results might be specific to the position in which the trend was displayed on Twitter. Columns (2) and (3) in Table 13 show that the results continue to hold if we separately look at trends that were in positions 1 to 5 on the day they emerged, versus in positions 6 to 10. Importantly, the pattern is more pronounced for the top trends and the first four days, which closely reflects our findings of field test 1.

As before, we check whether the results are robust to separating out clicks and retweets in the dependent variable. Table 14 shows that the main pattern holds: Clicks and retweets are lowest on the day a trend emerged, relative to later days.

4.5 Mechanism

We analyze the variation in message content and trend category, which are the two dimensions in the design of the field test that may shed light on the behavioral mechanism. We stratify the data by whether the message referred to a more unique/less commercial brand ('original Brixton-based') or a less unique/more commercial brand ('new global') and by whether the trend is more or less relevant to consumers interested in fashion. A comparison of Columns (1)-(4) in Table 15 suggests that independently, neither the phrasing of the advertising message as more or less unique, nor the closeness of the trend to fashion, influence the effect considerably: The effect holds across all stratifications.

We then investigate how early trend propagators respond to firm-sponsored content if the message is both more unique/less commercial and advertises a product category that is more closely related to the targeted trend and so more relevant to them. We focus on trends that are broadly related to fashion and stratify the data by message type. Column (5) indicates that the previous results do not hold when a message that refers to an 'original Brixton-based' fashion brand is targeted to users who posted on related trends. Instead, in this instance, targeting early trend propagators is similarly effective as targeting users on the two days after the trend emerged, and more effective than targeting users who post at even a later point in time. By contrast, Column (6) demonstrates that when the advertising message refers to a 'new global' fashion brand, the previously established pattern still holds: Targeting early trend propagators is least effective. These findings are in line with our theory that early trend propagators differ from other users in regulatory style, as they are extrinsically motivated and seek status rewards from their Twitter activities. Only messages that are both relevant and unique help them to attract the attention and acclaim of others.

The variation in advertising content that leads to a different response by early trend propagators also suggests that inattention is an unlikely explanation for the effect we observe:

If early trend propagators simply never read promoted tweets, then changing the phrasing should not affect their response.

To sum up, the results of field test 2 suggest that while advertising to early trend propagators is mostly not effective, it can - under a restrictive set of conditions - be similarly effective as advertising to users who post on trends later on: First, the product category advertised must be related to the trend the targeted users are posting on, and so be of likely interest to early trend propagators. Second, the advertising message should be unique and not-commercial to enhance a user's self-presentation and signal greater sub-cultural capital.

More broadly, the results of field test 2 confirm that the results hold outside the charitable context, for different types of advertising messages, and for related trends - unless the trends are related and the advertising message is more unique and less commercial. Likewise, the results hold both in the UK and in the US.

5 Discussion

5.1 Relationship of Early Trend Propagators to Earlier Conceptions of Key Individuals in Diffusion Processes

Reaching out to social media users who are participating in and spreading trends early on may be attractive to firms. Indeed, firms may hope that such early trend propagators may consume or disseminate information about their brand or product in similar ways as they consume and disseminate information on new topics or trends.

This hopefulness may stem from earlier insights into other types of individuals who have an unusually high social influence within a social contagion process. The marketing literature in fact distinguishes between three main types of influencers - early adopters, market mavens, and opinion leaders - who at first glance may seem similar to early trend propagators. Like early trend propagators, they disseminate information and appear to be somewhat more connected and more active. However, unlike early trend propagators,

they are responsive to advertising and other marketing messages. In light of our finding that early trend propagators are not easy targets for advertising despite their important role in the diffusion of trending topics on Twitter, we attempt to address the similarities and differences of early trend propagators with these other key individuals in the diffusion process.

Early adopters are product category experts interested in innovations within that product category, and the first to know about and to own a new product (Rogers, 1962; Bass, 1969). To maintain and affirm their position as pioneers, they actively seek out information and are interested in marketing and advertising communication about their product category. Early adopters influence the adoption decisions of a large group of other consumers indirectly, by owning and displaying the new product so that it is visible to others, and directly, by disseminating new product information and evaluations to other consumers. This leads to the classic s-curve diffusion pattern (Rogers, 1962; Bass, 1969).

Early trend propagators are similar to early adopters in that they want to be the first to know about a new development in their domain of interest. However, unlike early adopters they are not motivated by their involvement in a certain product category but by their involvement in trending topics within their domain of interest. Early trend propagators seek out information to be on top of the latest developments and to provide content that leads to acclaim from followers. Because they are interested in trending topics, advertising is not generally a source of valuable new information for early trend propagators since advertising conveys product information, not information that is trending. Also unlike early adopters, early trend propagators do not contribute to the adoption of a trending topic by an even larger group of people - trends are too short-lived for that to happen. Our data obtained through the Twitter API show a high initial volume of tweets about trending topics that very quickly tapers off.

A second type of influencer is the market maven, who is defined as a ‘diffuser of market-

place information' (Feick and Price, 1987) with a broad interest in and knowledge of markets. A market maven is not a product specialist but has a broader marketplace interest, and embraces the role of being the first to diffuse information about new products or information content to their network. At the heart of the very definition of market mavens is the hypothesis that in order to seek out relevant marketplace information, they will 'demonstrate higher levels of general market interest through...attention to advertising' (Feick and Price, 1987). This hypothesis was supported in subsequent studies which found that market mavens were good targets for marketing communications (Abratt et al., 1995). Traditional advertising content can be of use to market mavens as a source of the latest marketplace information that they pay attention to in private, and choose to share whenever a social situation arises in which this knowledge might be useful. Early trend propagators have a different motivation and derive prestige from sharing timely information of interest to potential followers to signal they are on top of the latest trends. Therefore, marketplace information such as advertising is only of interest to early trend propagators if it helps to achieve this goal.

A third related concept is that of an opinion leader. Opinion leaders can be monomorphic influentials that are specialized in a limited field (e.g., a product category) or polymorphic influentials exerting influence in a variety of fields (Merton, 1968). As such, the conceptual boundaries between opinion leaders and early adopters and market mavens are not absolute: If an opinion leader's field of interest pertains to a particular product category or marketplace, they could also be early adopters or market mavens. But opinion leadership is not limited to a product or market focus and could pertain to any field. Whatever their field(s) of interest, opinion leaders are highly involved in and knowledgeable about it, and it is this deep knowledge and expertise that motivates them to share information and influence others (Childers, 1986). Due to their knowledge and expertise they are often sought out as information sources by other consumers (King and Summers, 1970). As a result, opinion leaders have been described as information brokers who intervene between mass media and general

consumers (Feick and Price, 1987). If their domain of interest pertains to products and markets, opinion leaders may be interested in and responsive to advertising, as advertising and other marketing communication can be a valuable information source to further enhance their knowledge and expertise. In contrast, early trend propagators are not primarily motivated by a desire to enhance their knowledge and expertise, but by being among the first to know about a trending topic or development.

In sum, the main difference between different types of influencers previously discussed in the marketing literature and early trend propagators is the type of information they find interesting to further their objectives, and this difference impacts how responsive they are to marketing communication.

5.2 Conclusion

Online marketing, and micro-blogging services such as Twitter, have made it possible to target a large number of individual consumers in a timely manner based on their self-expressed interest in a topic. Reaching out to early trend propagators is becoming an increasingly popular marketing practice as marketing managers hope they will engage with their message and potentially help spread the word about their products or brands. While interest-based social media platforms like Twitter open up the possibility to target individuals who propagate trends early on with advertising messages, it is unclear whether doing so is effective. Using data from two field tests conducted on Twitter with a charity and a fashion label, we examine the effectiveness of promoted tweets, that is advertising messages sent to Twitter users, in engaging early trend propagators.

In our field tests, we identify each day the trending topics and target advertising to individuals who post on these trends that day and the following days. Consistently across both field tests, we find that early trend propagators (i.e., Twitter users who post on the trend the day it emerged) are significantly less likely to respond positively to the ad than

users who post on the trend during the following days. A series of robustness checks show that our results consistently hold.

We draw on self-determination theory (Deci and Ryan, 1980; Ryan and Deci, 2000) to advance a possible explanation why early trend propagators are not very responsive to advertising messages. We conceptualize early trend propagation as a fairly stable individual difference. Specifically, we suggest that early trend propagators have a regulatory style characterized by a ‘control orientation’, which means that they are motivated extrinsically through status rewards they receive when their posts are valued by their followers. They are therefore concerned with self-presentation, and use hashtagged trending discussions to present themselves to the Twitter publics as users who are knowledgeable about the latest trends. Because of the rapid pace of Twitter, being on top of latest trends is one way to signal sub-cultural capital. Since early trend propagators use Twitter to feel that their posts receive acclaim from their followers, they will engage with content and propagate content that serves this purpose. As such, they have little motivation to engage with messages by advertisers.

Our two field tests and additional survey data provide some support for this potential behavioral mechanism. Throughout both field tests, we do not find a negative effect of targeting individuals who post on commercially sponsored trends and who consequently seem generally open to firm-sponsored messages. Additionally, our second field test illustrates that early trend propagators are similarly responsive to advertising messages as others if the advertising message is perceived as more unique and less commercial, that is a type of content that more likely furthers their goal of self-presentation. Importantly, early trend propagators are still less responsive to advertising if only one of these conditions is met. In addition, the survey data demonstrates that early engagement with trends is strongly correlated with extrinsic, but not intrinsic motivation.

Our results have important implications for firms. First, many guidelines for sparking

contagion speak of the advantages of encouraging early trend propagators to spread word of mouth, but our results suggest that this might be difficult to achieve via advertising. Second, our results suggest that early trend propagators who are willing to adopt a commercial trend respond differently to advertising than early propagators of organic trends. Engaging the former group through advertising is likely to be more successful, but loses the organic quality of word-of-mouth communication. Third, our results point to a difference between early trend propagators and other known influencers in the diffusion process, such as early adopters, opinion leaders, or market mavens. If firms considered only the network characteristics of early trend propagators, they would appear similar to such influencers because early trend propagators have more followers and friends and are more active than users who do not adopt trends early on. However, the domain of interest and motivation of early trend propagators is very different from that of known influencers, and these differences explain why early trend propagators are less responsive to commercial messages and therefore do not make for good advertising targets. In addition, the fast-paced and ephemeral nature of social media trends implies that the influence of early trend propagators is at best short-lived. Finally, our results have implications for the future of advertising on micro-blogging sites. It is tempting to think that what makes sites such as Twitter distinctive from alternative advertising platforms is the fact that they originate trends and relay timely information, but our results suggest that these features are likely to distinguish them as useful advertising platforms only under a limited set of conditions. Indeed, recent announcements by Twitter such as the decision to not show ads to top users, suggest that indeed Twitter does not see the benefits of showing ads to users who potentially could have influence on other Twitter users.²⁶

There are of course limitations to our research. First, while our data provides evidence leading us to suggest the difference in motivation for Twitter use as a behavioral mechanism,

²⁶<http://www.fastcompany.com/3055910/fast-feed/twitter-has-stopped-showing-ads-to-its-most-valuable-u>

we are unable to directly measure and test for this effect. Therefore, our study presents suggestive rather than conclusive evidence for our suggested mechanism. A number of other potential mechanisms at play include a general lack of sensitivity to advertising both on Twitter and outside Twitter, or perhaps some form of reactance which we can not tease out in this study and we leave to future research. Second, we emphasize that there are aspects of our definition of an early trend propagator which are not definitive. Future research could explore in more detail the definition and boundaries of early trend propagators such as the extent to which this is a persistent trait over a consumer's lifetime. Third, our empirical focus is on the micro-blogging site Twitter. While Twitter is an increasingly important medium and an attractive platform for us to study because of its increasing importance in the identification of trends, we recognize that it allows a very specific kind of ad format and that the results may not fully generalize to other formats. Last, we use data from two empirical tests conducted by a charity in the UK and by a fashion firm in the US. This gives us confidence that our results hold across product categories. However, it is still possible that there are yet undiscovered ways by which targeting early trend propagators may be successful in other product domains. Notwithstanding these limitations, we believe that our results offer a first insight into the challenges that firms may face when trying to use online behavior to identify early trend propagators to target advertising messages.

References

- Abratt, R., D. Nel, and C. Nezer (1995). Role of the market maven in retailing: A general marketplace influencer. *Journal of Business and Psychology* 10(1), 31–55.
- Ai, C. and E. C. Norton (2003, July). Interaction terms in logit and probit models. *Economics Letters* 80(1), 123–129.
- Antonakaki, D., I. Polakis, E. Athanasopoulos, S. Ioannidis, and P. Fragopoulou (2015). Think before rt: An experimental study of abusing Twitter trends. In L. M. Aiello and D. McFarland (Eds.), *Social Informatics*, Volume 8852 of *Lecture Notes in Computer Science*, pp. 402–413. Springer International Publishing.
- Arvidsson, A. and A. Caliendo (2016). Brand public. *Journal of Consumer Research* 42(5), 727–748.
- Asur, S. and B. A. Huberman (2010). Predicting the future with social media. In *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on*, Volume 1, pp. 492–499. IEEE.
- Athey, S. and G. Ellison (2011). Position auctions with consumer search. *The Quarterly Journal of Economics* 126(3), 1213–1270.
- Bakshy, E., J. M. Hofman, W. A. Mason, and D. J. Watts (2011). Everyone’s an influencer: Quantifying influence on Twitter. In *Proceedings of the fourth ACM international conference on Web search and data mining*, WSDM ’11, New York, NY, USA, pp. 65–74. ACM.
- Bass, F. M. (1969). A new product growth for model consumer durables. *Management science* 15(5), 215–227.
- Berger, J. (2014). Word of mouth and interpersonal communication: A review and directions for future research. *Journal of Consumer Psychology* 24(4), 586–607.
- Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics* 119(1), 249–275.
- Bollen, J., H. Mao, and X. Zeng (2011). Twitter mood predicts the stock market. *Journal of Computational Science* 2(1), 1–8.
- Bughin, J. (2003). The diffusion of internet banking in western europe. *Electronic Markets* 13(3), 251–258.
- Cha, M., H. Haddadi, F. Benevenuto, and P. K. Gummadi (2010). Measuring user influence in Twitter: The million follower fallacy. In W. W. Cohen and S. Gosling (Eds.), *ICWSM*. The AAAI Press.
- Chen, G. H., S. Nikolov, and D. Shah (2013). A latent source model for nonparametric time series classification. In *Advances in Neural Information Processing Systems*, pp. 1088–1096.

- Chen, Y., Q. Wang, and J. Xie (2011). Online social interactions: A natural experiment on word of mouth versus observational learning. *Journal of Marketing Research* 48(2), 238–254.
- Childers, T. (1986). An assessment of the psychometric properties of an opinion leadership scale. *Journal of Marketing Research* 23(2), 184–188.
- Chintagunta, P. K. (2001, December). Endogeneity and heterogeneity in a probit demand model: Estimation using aggregate data. *Marketing Science* 20(4), 442–456.
- Deci, E. L. and R. M. Ryan (1980). The empirical exploration of intrinsic motivational processes. *Advances in experimental social psychology* 13(2), 39–80.
- Deci, E. L. and R. M. Ryan (1985). The general causality orientations scale: Self-determination in personality. *Journal of research in personality* 19(2), 109–134.
- Du, R. Y. and W. A. Kamakura (2012). Quantitative trendspotting. *Journal of Marketing Research* 49(4), 514–536.
- Dubé, J.-P., X. Luo, and Z. Fang (2017). Self-signaling and prosocial behavior: A cause marketing experiment. *Marketing Science* 36(2), 161–186.
- Falkinger, J. (2008). Limited attention as a scarce resource in information-rich economies*. *The Economic Journal* 118(532), 1596–1620.
- Feick, L. F. and L. L. Price (1987). The market maven: A diffuser of marketplace information. *The Journal of Marketing* 1(51), 83–97.
- Flath, D. and E. W. Leonard (1979). A comparison of two logit models in the analysis of qualitative marketing data. *Journal of Marketing Research* 16(4), pp. 533–538.
- Ghose, A. and S. Yang (2009). An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Science* 55(10), 1605–1622.
- Goel, S., J. M. Hofman, S. Lahaie, D. M. Pennock, and D. J. Watts (2010). Predicting consumer behavior with web search. *Proceedings of the National Academy of Sciences* 107(41), 17486–17490.
- Goldfarb, A. and C. Tucker (2011a, May). Online display advertising: Targeting and obtrusiveness. *Marketing Science* 30, 389–404.
- Goldfarb, A. and C. Tucker (2011b). Search engine advertising: Channel substitution when pricing ads to context. *Management Science* 57(3), 458–470.
- Goldsmith, R. E. and C. F. Hofacker (1991). Measuring consumer innovativeness. *Journal of the Academy of Marketing Science* 19(3), 209–221.
- Gong, S., J. Zhang, P. Zhao, and X. Jiang (2017). Tweeting as a marketing tool ? field experiment in the tv industry. *Journal of Marketing Research, Forthcoming* ?(?), ?

- Hinz, O., B. Skiera, C. Barrot, and J. U. Becker (2011). Seeding strategies for viral marketing: an empirical comparison. *Journal of Marketing* 75(6), 55–71.
- Ishihara, M. and A. Ching (2012). Dynamic demand for new and used durable goods without physical depreciation: The case of japanese video games. Rotman School of Management Working Paper 2189871.
- Kankanhalli, A., B. C. Tan, and K.-K. Wei (2005). Contributing knowledge to electronic knowledge repositories: an empirical investigation. *MIS Quarterly* 29(1), 113–143.
- Katona, Z., P. P. Zubcsek, and M. Sarvary (2011). Network effects and personal influences: The diffusion of an online social network. *Journal of Marketing Research* 48(3), 425–443.
- Kempe, D., J. Kleinberg, and É. Tardos (2003). Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146. ACM.
- King, C. W. and J. O. Summers (1970, February). Overlap of opinion leadership across product categories. *Marketing Research* 7(1), 43–50.
- Kong, S., Q. Mei, L. Feng, and Z. Zhao (2014). Real-time predicting bursting hashtags on twitter. In *Web-Age Information Management*, pp. 268–271. Springer.
- Krauss, R. M. and S. R. Fussell (1991). Perspective-taking in communication: Representations of others’ knowledge in reference. *Social Cognition* 9, 2–24.
- Lambrecht, A. and C. Tucker (2013). When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research* 50(5), 561–576.
- Lewis, R. A., J. Rao, and D. Reiley (2011). Here, there, everywhere: Correlated online behaviors can lead to overestimates of the effects of advertising. In *International World Wide Web Conference*.
- Lewis, R. A. and D. H. Reiley (2014). Online ads and offline sales: measuring the effect of retail advertising via a controlled experiment on yahoo! *Quantitative Marketing and Economics* 12(3), 235–266.
- Manchanda, P., J.-P. Dube, K. Y. Goh, and P. K. Chintagunta (2006). The effect of banner advertising on internet purchasing. *Journal of Marketing Research* 43(1), 98 – 108.
- Mathwick, C. and E. Rigdon (2004). Play, flow, and the online search experience. *Journal of consumer research* 31(2), 324–332.
- Merton, R. K. (1968). *Social theory and social structure*. Simon and Schuster.
- Rangaswamy, A. and S. Gupta (2000). *New Product Diffusion Models*, Chapter Innovation adoption and diffusion in the digital environment: some research opportunities, pp. 75–96. Kluwer Academic Publishers Hingham, Massachusetts.
- Rogers, E. (1962). *Diffusion of innovations*. New York,: Free Press.

- Rutz, O. J. and R. E. Bucklin (2011). From Generic to Branded: A Model of Spillover Dynamics in Paid Search Advertising. *Journal of Marketing Research* 48(1), 87–102.
- Ryan, R. M. (1982). Control and information in the intrapersonal sphere: An extension of cognitive evaluation theory. *Journal of personality and social psychology* 43(3), 450.
- Ryan, R. M. and E. L. Deci (2000). Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary educational psychology* 25(1), 54–67.
- Ryan, S. and C. Tucker (2012). Heterogeneity and the dynamics of technology adoption. *Quantitative Marketing and Economics* 10(1), 63–109.
- Sahni, N. (2015). Effect of temporal spacing between advertising exposures: Evidence from an online field experiment. *Quantitative Marketing and Economics* 13(3), 203–247.
- Stephen, A. T. and O. Toubia (2010). Deriving Value from Social Commerce Networks. *Journal of Marketing Research* 47(2), 215–228.
- Stephen, A. T. and O. Toubia (2013). Intrinsic vs. image-related utility in social media: Why do people contribute content to Twitter? *Marketing Science* 32(3), 368–392.
- Stonedahl, F., W. Rand, and U. Wilensky (2010). Evolving viral marketing strategies. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pp. 1195–1202. ACM.
- Toubia, O., A. T. Stephen, and A. Freud (2009). Viral Marketing: A Large-Scale Field Experiment. Mimeo, INSEAD.
- Vaynerchuck, G. (2013). How to master the 4 big social-media platforms. Technical report, Inc. Magazine.
- Watts, D. J. and P. S. Dodds (2007). Influentials, networks, and public opinion formation. *Journal of Consumer Research* 34, 441–458.
- Weng, J., E.-P. Lim, J. Jiang, and Q. He (2010). Twiterrank: finding topic-sensitive influential twitterers. In *Proceedings of the third ACM international conference on Web search and data mining*, pp. 261–270. ACM.
- Yoganarasimhan, H. (2012). Impact of social network structure on content propagation: A study using YouTube data. *Quantitative Marketing and Economics* 10(1), 111–150.
- Zhang, Y., W. W. Moe, and D. A. Schweidel (2016). Modeling the role of message content and influencers in social media rebroadcasting. *International Journal of Research in Marketing* 34(1), 100–119.




Figure 1: Screenshot of Twitter Trends and Promoted Tweets




Figure 2: Tweets Targeting Trending Topic




Figure 3: Promoted Tweet Targeting Trending Topic




Figure 4: Histogram showing diffusion pattern of trends




Figure 5: Histogram showing diffusion pattern of trends in position 1 to 5




Figure 6: Engagement rate by day after trend peaked, Study 1




Figure 7: Click rate by day after trend peaked, Study 1




Figure 8: Retweet rate by day after trend peaked, Study 1




Figure 9: Engagement Rate by Day After Trend Peaked, Study 2




Figure 10: Click Rate by Day After Trend Peaked, Study 2




Figure 11: Retweet Rate by Day After Trend Peaked, Study 2

Table 1: Relationship between share of tweets made on day trend emerged across periods 1 and 2. Numbers are percentage of users in data

Share of Tweets Made on Day Trend Emerged	Period 2: 0 - 0.33	Period 2: 0.33 - 0.67	Period 2: 0.67 - 1
Period 1: 0 - 0.33	5.1%	3.3%	8.0%
Period 1: 0.33 - 0.97	2.9%	4.0%	6.7%
Period 1: 0.67 - 1	13.0%	12.6%	43.8%

Table 2: Relationship between tweeting in a category on day 1 in first 30 days on tweeting in the category on day 1 in second 30 days

	(1) General News Days 31-60	(2) Holiday Days 31-60	(3) Politics Days 31-60	(4) Pop Culture Days 31-60	(5) Sport Days 31-60	(6) Twitter-Specific Days 31-60
General News, Days 1 - 30	1.319*** (0.014)	0.350*** (0.008)	0.474*** (0.008)	0.607*** (0.009)	0.444*** (0.012)	0.167*** (0.009)
Holiday, Days 1 - 30	-0.312*** (0.010)	0.851*** (0.004)	0.066*** (0.003)	0.281*** (0.005)	-0.506*** (0.007)	0.712*** (0.004)
Politics, Days 1 - 30	0.717*** (0.009)	0.210*** (0.004)	0.800*** (0.003)	0.437*** (0.005)	0.085*** (0.006)	0.076*** (0.004)
Pop Culture, Days 1 - 30	0.324*** (0.011)	0.445*** (0.005)	0.197*** (0.004)	1.017*** (0.005)	0.459*** (0.007)	0.049*** (0.005)
Sport, Days 1 - 30	0.676*** (0.011)	-0.147*** (0.006)	0.091*** (0.005)	0.406*** (0.006)	1.541*** (0.006)	-0.007 (0.006)
Twitter-Specific, Days 1 - 30	0.484*** (0.012)	0.399*** (0.005)	0.252*** (0.005)	0.298*** (0.006)	-0.138*** (0.009)	1.093*** (0.005)
Constant	-4.039*** (0.008)	-1.517*** (0.003)	-0.762*** (0.003)	-2.399*** (0.004)	-2.675*** (0.005)	-1.725*** (0.003)
Observations	1737458	1737458	1737458	1737458	1737458	1737458

Logit estimates. Dependent variable is whether or not a user tweeted in that category on day 1 of a trend in the second 30 days of the data. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Relationship between tweeting in a category on day 1 in first 30 days on tweeting in the category on day 1 in second 30 days - Seemingly Unrelated Regression

	(1) General News Days 31-60	(2) Holiday Days 31-60	(3) Politics Days 31-60	(4) Pop Culture Days 31-60	(5) Sport Days 31-60	(6) Twitter-Specific Days 31-60	
General News, Days 1 - 30	0.071*** (0.001)	0.074*** (0.002)	0.113*** (0.002)	0.094*** (0.001)	0.038*** (0.001)	0.032*** (0.002)	
Holiday, Days 1 - 30	-0.005*** (0.000)	0.172*** (0.001)	0.016*** (0.001)	0.039*** (0.001)	-0.032*** (0.000)	0.126*** (0.001)	
Politics, Days 1 - 30	0.023*** (0.000)	0.043*** (0.001)	0.193*** (0.001)	0.058*** (0.001)	0.005*** (0.000)	0.016*** (0.001)	
Pop Culture, Days 1 - 30	0.012*** (0.000)	0.090*** (0.001)	0.046*** (0.001)	0.151*** (0.001)	0.036*** (0.001)	0.011*** (0.001)	Seemingly
Sport, Days 1 - 30	0.024*** (0.000)	-0.020*** (0.001)	0.021*** (0.001)	0.056*** (0.001)	0.175*** (0.001)	0.003*** (0.001)	
Twitter-Specific, Days 1 - 30	0.017*** (0.000)	0.083*** (0.001)	0.059*** (0.001)	0.043*** (0.001)	-0.010*** (0.001)	0.221*** (0.001)	
Constant	0.013*** (0.000)	0.172*** (0.001)	0.317*** (0.001)	0.068*** (0.000)	0.066*** (0.000)	0.144*** (0.001)	
Observations	1737458						

unrelated regression with linear probability models. Dependent variable is whether or not a user tweeted in that category on day 1 of a trend in the second 30 days of the data. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Summary Statistics for Characteristics of Early Trend Propagators relative to Others

(a) Early Trend Propagators			
	Mean	Std Dev	Median
Followers	1208.22	2906.37	415
Friends	805.22	1574.60	386
Number of posts	19697.44	37519.55	7481
Days since account created	1305.09	797.49	1262
Observations	1217612		

(b) Others			
	Mean	Std Dev	Median
Followers	1088.01	2747.61	357
Friends	760.77	1515.56	366
Number of posts	15260.04	30783.07	5497
Days since account created	1328.19	811.72	1310
Observations	519846		

Table 5: Summary of different field tests

Variation	Field Test 1	Field Test 2
Domain	Homelessness Charity	Fashion Firm
Geographical target	UK	US Females
Number of days trends were identified	19	20
Number of trends targeted each day	Only 1	Top 10
Number of sponsored trends targeted	7 (incl. in daily trends)	9 (additional trends)
Number of days each trend-based campaign runs	4 days	7 days
Additional keywords targeted	None	#shopping, shopping
Message variations	16	2
Length of total test	22 days	26 days

Table 6: Top Trends on Twitter by Day, Study 1

Day	Trend	Type
1.	#Gooddeeds	Sponsored
2.	#TheTribez	Organic
3.	#WreckTheHalls	Sponsored
4.	#RIPNelsonMandela	Organic
5.	#ProjectUpgrade	Sponsored
6.	#RebeccaBlack	Organic
7.	#MPs	Organic
8.	#RAW	Organic
9.	#Rewind2013	Sponsored
10.	#PopATrace	Organic
11.	#ExtinctionDay	Sponsored
12.	#Cook	Organic
13.	#ScherzingHair	Sponsored
14.	#Spoty	Organic
15.	#AirportsCommission	Organic
16.	#Foodbankdebate	Organic
17.	#TryYourGearOn	Sponsored
18.	#ApolloTheatre	Organic
19.	#HasJustineLandedYet	Organic

Table 7: Summary Statistics per Campaign, Study 1

	Mean	Std Dev	Min	Max
Daily impressions	413.66	544.50	0	3850
Daily engagements	3.81	4.70	0	26
Daily retweets	0.16	0.52	0	6
Daily clicks	3.56	4.37	0	26
Daily spend	1.64	1.73	0	8
Cost per Engagement	0.35	0.18	0	1
Observations	1216			

Table 8: Main Results, Study 1

	Main (1)	(2)	(3)	Large Campaigns (4)	Charities (5)	Spend Controls (6)	ML Probit (7)	WLS Logit (8)
Days elapsed since trend emerges	1.171*** (0.045)	1.192*** (0.048)						
Targeted day trend emerges			-2.911*** (0.230)	-3.023*** (0.521)	-1.223*** (0.152)	-2.418*** (0.384)	-1.154*** (0.084)	-0.995*** (0.146)
Targeted day 1 after trend emerges			-1.936*** (0.154)	-1.918*** (0.286)	-0.587*** (0.102)	-1.600*** (0.257)	-0.768*** (0.056)	-0.667*** (0.113)
Targeted day 2 after trend emerges			-1.194*** (0.045)	-1.209*** (0.046)	-0.416*** (0.034)	-1.026*** (0.114)	-0.469*** (0.015)	-0.580*** (0.102)
Cost per Engagement						-0.321 (0.210)		
Date Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Message Fixed Effects	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Trend Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	503012	503012	503012	243567	55352	503012	503012	676

Estimates from Aggregate Logit Estimation. Dependent variable is engagement. Robust standard errors clustered at the campaign level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Alternative Dependent Variables, Study 1

	(1)	(2)	(3)	(4)
	Daily clicks	Daily retweets	Daily clicks	Daily retweets
Days elapsed since trend emerges	1.239*** (0.033)	1.456*** (0.422)		
Targeted day trend emerges			-3.133*** (0.211)	-4.781*** (1.428)
Targeted day 1 after trend emerges			-2.076*** (0.143)	-3.795*** (1.039)
Targeted day 2 after trend emerges			-1.240*** (0.031)	-2.143*** (0.672)
Date Fixed Effects	Yes	Yes	Yes	Yes
Message Fixed Effects	Yes	Yes	Yes	Yes
Trend Fixed Effects	Yes	Yes	Yes	Yes
Observations	503012	501291	503012	501291

Estimates from aggregate logit estimation. Dependent variable as shown. Robust standard errors clustered at the campaign level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Exploring Different Explanations, Study 1

	(1) Sponsored Trend	(2) Non-Sponsored Trend
Targeted day trend emerges	0.224 (0.149)	-2.811*** (0.221)
Targeted day 1 after trend emerges	1.244*** (0.130)	-1.804*** (0.136)
Targeted day 2 after trend emerges	-1.050*** (0.165)	-1.196*** (0.045)
Date Fixed Effects	Yes	Yes
Message Fixed Effects	Yes	Yes
Trend Fixed Effects	Yes	Yes
Observations	135966	366931

Estimates from aggregate logit estimation. Dependent variable is engagements. Robust standard errors clustered at the campaign level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Summary Statistics per Campaign, Study 2

	Mean	Std Dev	Min	Max
Daily impressions	935.35	796.89	0	8481
Daily engagements	8.40	6.62	0	66
Daily retweets	0.53	0.95	0	8
Daily clicks	8.04	6.31	0	64
Daily spend	6.59	3.95	0	27
Cost per Engagement	0.92	0.33	0	2
Observations	2978			

Table 12: Main Results Study 2

	Main (1)	(2)	(3)	Shopping Baseline (4)	Sponsored (5)
Days elapsed since trend emerges	0.027*** (0.006)	0.044*** (0.005)			
Targeted day trend emerges			-0.239*** (0.037)		1.297 (1.683)
Targeted day 1 after trend emerges			-0.044 (0.035)		0.988 (1.404)
Targeted day 2 after trend emerges			-0.012 (0.034)		0.720 (1.122)
Targeted day 3 after trend emerges			0.070** (0.034)		0.909 (0.840)
Targeted day 4 after trend emerges			0.091*** (0.034)		0.624 (0.563)
Targeted day 5 after trend emerges			0.085*** (0.029)		0.275 (0.283)
Targeted day trend emerges				-0.322*** (0.063)	
Targeted day 1 after trend emerges				-0.129** (0.064)	
Targeted day 2 after trend emerges				-0.098 (0.063)	
Targeted day 3 after trend emerges				-0.018 (0.057)	
Targeted day 4 after trend emerges				0.003 (0.062)	
Targeted day 5 after trend emerges				-0.005 (0.059)	
Targeted day 6 after trend emerges				-0.089* (0.053)	
Date Fixed Effects	Yes	Yes	Yes	Yes	Yes
Trend Fixed Effects	Yes	Yes	Yes	Yes	Yes
Trend Pos Fixed Effects	Yes	Yes	Yes	Yes	No
Cost per Engagement Control	No	Yes	Yes	Yes	Yes
Message Control	No	Yes	Yes	Yes	Yes
Observations	2605415	2581132	2581132	2635320	125381

Estimates from aggregate logit estimation. Dependent variable is engagements. Robust standard errors clustered at the campaign level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 13: Robustness, Study 2

	Control for Trend Size (1)	Top Trends (2)	Bottom Trends (3)
Targeted day trend emerges	-0.239*** (0.039)	-0.388*** (0.066)	-0.215*** (0.047)
Targeted day 1 after trend emerges	-0.044 (0.036)	-0.195*** (0.064)	-0.015 (0.041)
Targeted day 2 after trend emerges	-0.012 (0.035)	-0.125** (0.060)	0.007 (0.041)
Targeted day 3 after trend emerges	0.070** (0.034)	-0.015 (0.059)	0.081** (0.041)
Targeted day 4 after trend emerges	0.091*** (0.034)	0.019 (0.061)	0.101** (0.041)
Targeted day 5 after trend emerges	0.085*** (0.030)	-0.011 (0.054)	0.114*** (0.033)
Trend size	-0.000 (0.000)		
Date Fixed Effects	Yes	Yes	Yes
Trend Fixed Effects	Yes	Yes	Yes
Trend Pos Fixed Effects	Yes	Yes	Yes
Cost per Engagement Control	Yes	Yes	Yes
Message Control	Yes	Yes	Yes
Observations	2581132	774464	1806668

Estimates from aggregate logit estimation. Dependent variable is engagements. Robust standard errors clustered at the campaign level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14: Alternative Dependent Variables, Study 2

	Daily clicks (1)	Daily retweets (2)	Daily clicks (3)	Daily retweets (4)
Days elapsed since trend emerges	0.045*** (0.005)	0.053*** (0.016)		
Targeted day trend emerges			-0.245*** (0.036)	-0.279** (0.117)
Targeted day 1 after trend emerges			-0.044 (0.036)	-0.175 (0.119)
Targeted day 2 after trend emerges			-0.000 (0.034)	-0.169 (0.148)
Targeted day 3 after trend emerges			0.066* (0.035)	0.033 (0.131)
Targeted day 4 after trend emerges			0.093*** (0.036)	0.075 (0.097)
Targeted day 5 after trend emerges			0.088*** (0.030)	0.019 (0.123)
Date Fixed Effects	Yes	Yes	Yes	Yes
Trend Fixed Effects	Yes	Yes	Yes	Yes
Cost per Engagement Control	Yes	Yes	Yes	Yes
Message Control	Yes	Yes	Yes	Yes
Observations	2581132	2557006	2581132	2557006

Estimates from aggregate logit estimation. Dependent variable as shown. Robust standard errors clustered at the campaign level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 15: Explaining the Behavioral Mechanism, Study 2

	Original Message (1)	Non-Original Message (2)	Related Category (3)	Unrelated Category (4)	Related+Original (5)	Related+Not Original (6)
Targeted day trend emerges	-0.306*** (0.046)	-0.185*** (0.051)	-0.331*** (0.066)	-0.242*** (0.046)	0.619*** (0.111)	-0.981*** (0.099)
Targeted day 1 after trend emerges	-0.087** (0.042)	0.007 (0.054)	-0.205*** (0.074)	-0.019 (0.043)	0.627*** (0.113)	-0.770*** (0.104)
Targeted day 2 after trend emerges	-0.038 (0.042)	0.025 (0.057)	-0.086 (0.063)	-0.019 (0.041)	0.528*** (0.104)	-0.488*** (0.089)
Targeted day 3 after trend emerges	-0.034 (0.047)	0.182*** (0.053)	-0.037 (0.054)	0.087** (0.044)	0.398*** (0.086)	-0.305*** (0.077)
Targeted day 4 after trend emerges	0.058 (0.036)	0.131** (0.055)	-0.025 (0.057)	0.129*** (0.040)	0.267*** (0.064)	-0.206** (0.092)
Targeted day 5 after trend emerges	0.014 (0.039)	0.162*** (0.047)	0.024 (0.055)	0.100*** (0.034)	0.224*** (0.079)	-0.092 (0.081)
Date Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Trend Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Trend Pos Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Cost per Engagement Control	Yes	Yes	Yes	Yes	Yes	Yes
Message Control	No	No	Yes	Yes	No	No
Observations	1331456	1249676	858342	1722790	443897	414445

Estimates from aggregate logit estimation. Dependent variable is engagements. Robust standard errors clustered at the campaign level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.