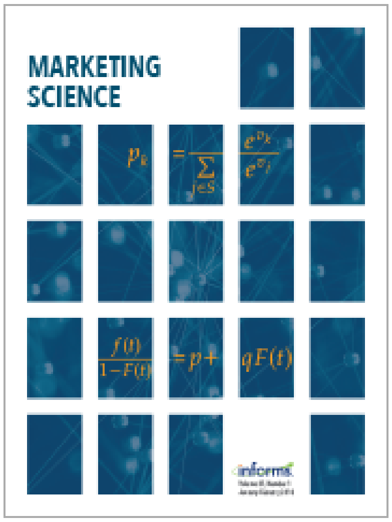

Marketing Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Social Advertising Effectiveness Across Products: A Large-Scale Field Experiment

Shan Huang, Sinan Aral, Yu Jeffrey Hu, Erik Brynjolfsson

MARKETING SCIENCE
Articles in Advance, pp. 1–24
ISSN 0732-2399 (print), ISSN 1526-548X (online)
http://pubsonline.informs.org/journal/mksc

Social Advertising Effectiveness Across Products: A Large-Scale Field Experiment

Shan Huang,[a] Sinan Aral,[b] Yu Jeffrey Hu,[c] Erik Brynjolfsson[b]

[a] Michael G. Foster School of Business, University of Washington, Seattle, Washington 98195; [b] Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142; [c] Scheller College of Business, Georgia Institute of Technology, Atlanta, Georgia 30332

Contact: shanh18@uw.edu, https://orcid.org/0000-0002-0276-1271 (SH); sinan@mit.edu, https://orcid.org/0000-0002-2762-058X (SA); jeffrey.hu@scheller.gatech.edu (YJH); erikb@mit.edu, https://orcid.org/0000-0002-8031-6990 (EB)

Received: September 24, 2017
Revised: October 17, 2018; October 2, 2019; February 24, 2020
Accepted: March 1, 2020
Published Online in Articles in Advance: October 5, 2020

https://doi.org/10.1287/mksc.2020.1240

Abstract. Most of the empirical evidence on social advertising effectiveness focuses on a single product at a time. As a result, little is known about how the effectiveness of social advertising varies across product categories or product characteristics. We therefore collaborated with a large online social network to conduct a randomized field experiment measuring social ad effectiveness across 71 products in 25 categories among more than 37 million users. We found some product categories, like clothing, cars, and food, exhibited significantly stronger social advertising effectiveness than other categories, like financial services, electrical appliances, and mobile games. More generally, we found that status goods, which rely on status-driven consumption, displayed strong social advertising effectiveness. Meanwhile, social ads for experience goods, which rely on informational social influence, did not perform any better or worse than social ads for their theoretical counterparts, search goods. Social advertising effectiveness also significantly varied across the relative characteristics of ad viewers and their friends shown in ads. Understanding the heterogeneous effects of social advertising across products can help marketers differentiate their social advertising strategies and lead researchers to more nuanced theories of social influence in product evaluation.

History: K. Sudhir served as the editor-in-chief and Christophe Van den Bulte served as associate editor for this article.This paper has been accepted for the *Marketing Science* Special Issue on Field Experiments.

Supplemental Material: Data replication files and the online appendix are available at https://doi.org/10.1287/mksc.2020.1240.

Keywords: social networks • social influence • social advertising • product • heterogeneity • field experiments

1. Introduction

Spending on social advertising is increasing dramatically, reflecting the high expectations advertisers place on this new form of advertising. Social advertising is a broad term, and its applications range widely, from network targeting, which targets those who are connected to previous adopters, to viral marketing, which encourages current adopters to spread positive word-of-mouth about products. But, the most widely used form of social advertising is arguably the placement of *social cues* in ads to encourage ad engagement through the power of social proof. For example, Facebook's social advertising places the images and names of Facebook friends who have liked a brand in their ads. Google's Shared Endorsement ads do the same thing, placing the names, images, and product ratings of others in product search results. These social ads rely on the power of social influence in product adoption and the value of social cues for social media engagement to encourage lift in ad effectiveness.

In this paper, we define *social advertising* as the placement of social cues or endorsements in ads shown to the friends of those who have engaged with a brand or product. Social influence, the effect of our behaviors and opinions on our peers (Turner 1991), is critical to the effectiveness of social ads and is one of the most important behavioral mechanisms driving the spread of products and behaviors through society (e.g., Van den Bulte 2000, Tucker 2008, Bakshy et al. 2009, Stephen and Toubia 2010, Aral 2011, Iyengar et al. 2011, Berger and Milkman 2012).

Although recent work has demonstrated that social ads achieve significant lift from the social proof in peer endorsements (e.g., Aral and Walker 2012, Bakshy et al. 2012, Taylor et al. 2013, Bapna and Umyarov 2015), almost all of the empirical evidence to date focuses on a single product at a time. Previous research has examined the impact of product types on ad effectiveness (e.g., Hanssens and Weitz 1980, Berger and Schwartz 2011, Bart et al. 2014, Colicev et al. 2017), but only limited research has systematically investigated the heterogeneity of social advertising effectiveness across products or how product characteristics moderate the impact of social influence on product adoption decisions (Bearden and Etzel 1982, Aral 2011).

The goal of our research is to identify the heterogeneous effects of social advertising across products and to investigate how social influence in product decisions varies across product characteristics. Are social ads more effective for electronics products or fashion accessories? Are we more likely to be swayed by the opinions of our friends when shopping for status goods or when we are seeking trusted information about a product? We simply do not know the answers to these questions, and it is difficult to generalize a theory of social influence in product evaluation, from one product to the next, while parameter estimates of influence in consumer decisions remain unknown or idiosyncratic (Friedman and Friedman 1979, Bearden and Etzel 1982, Kulviwat et al. 2009, Stephen and Galak 2012).

Although social influence is of central importance in marketing and social science more broadly (e.g., Deutsch and Gerard 1955, Burnkrant and Cousineau 1975, Sacerdote 2001, Cialdini and Goldstein 2004, Van den Bulte and Wuyts 2007, Trusov et al. 2009, Christakis and Fowler 2013), the causal estimation of social influence, especially in real business contexts, is a recent development (e.g., Bakshy et al. 2012, Muchnik et al. 2013, Aral and Walker 2014). Social influence is endogenous and randomized experiments improve influence identification by eliminating bias created by homophily, correlated effects, and confounding factors (Manski 1993). Online social networking platforms provide unprecedented opportunities for researchers to deploy such randomized field experiments at a population scale. The large-scale data that result from such experiments enable the detection of subtle but economically important effects across subpopulations. For example, previous work has employed large-scale field experiments to identify social influence (e.g., Aral and Walker 2011, Bakshy et al. 2012, Muchnik et al. 2013, Bond et al. 2012, Jones et al. 2017) and estimate the moderating effects of individual (Aral and Walker 2012), dyadic (Aral and Walker 2014, Taylor et al. 2015), and behavioral characteristics (Iyengar et al. 2015, Huang 2016). But, no experiment that we are aware of examines the heterogeneity in social influence across products.

We therefore designed and analyzed a randomized field experiment to measure social ad effectiveness across 71 products in 25 product categories and to examine how social influence in ad engagement varies across product characteristics. The experiment was conducted on a random sample of more than 37 million users of a large social network (WeChat) and focused on WeChat Moments ads, a type of social advertisement displayed in WeChat users' news feeds. WeChat is a world-leading mobile social networking platform with over a billion monthly active users. Our experiment involves user–ad–level randomization of social cues shown on WeChat Moments ads. By randomly assigning the presence and number of social cues displayed on otherwise identical ads, in a real-world context, we were able to obtain unbiased estimates of the impact of social influence on ad engagement across many different products simultaneously. Social influence and social advertising effectiveness in our experiment are measured by the degree to which social cues (i.e., friends' likes), representing friends' endorsements of products, affect users' engagement with social advertising (i.e., click-throughs).

Click-through rates are a critical measure of social advertising performance. Recent empirical evidence has shown that site visits, which are an upper-funnel outcome, can lead to lower-funnel outcomes, such as conversions (Johnson et al. 2017).[1] Furthermore, a growing literature emphasizes the importance of social endorsements in affecting individual decisions (e.g., Salganik et al. 2006, Bond et al. 2017), including ad click-throughs (e.g., Bakshy et al. 2012, Tucker 2016). Clicking on an ad represents a costly search for information that takes time. Users are likely to invest more time searching products through an ad if their peers endorse the ad through social engagement.

Bakshy et al. (2012) identify the *average* ad engagement effects of social advertising using large-scale randomized field experiments on Facebook. But, our paper provides the first large-scale experimental evidence of the *heterogeneous* effects of social advertising across products. We found that some product categories, like food, clothing, and cars, exhibited significantly stronger social advertising effectiveness than other categories like financial services, electrical appliances, and mobile games. More generally, we found that status goods, which displayed status signals, exhibited strong social advertising effectiveness, but social ads for experience goods with greater product uncertainties did not perform any better than social ads for their theoretical counterparts, search goods. The relative status and product involvement of the friends displayed in ads, compared with the ad viewers, were also critical to social ad effectiveness for different products. Understanding the heterogeneous effects of social ads across products will help researchers build more nuanced theories of social influence in product decisions and help marketers target their social advertising strategies more effectively.

2. Theory

2.1. Social Advertising Effectiveness Across Products

Multiple theories make clear predictions about when social influence will be salient for consumer decisions. For example, social influence may be a consequence of learning (Cai and Chen 2009, Zhang 2010). Consumers may seek out their friends' experience with products to infer their quality or evaluate peers' product adoption decisions to infer their value (Burnkrant and Cousineau 1975, Lin et al. 2015), especially when the products exhibit greater quality uncertainty, such as with experience goods (Nelson 1970, Van den Bulte and Lilien 2001, Zhu and Zhang 2010). On the other hand, consumers also use products to build, signal, and maintain their social status. Consumption and status are likely related for some products but not for others (Veblen 1899, Bernheim 1994, O'Cass and McEwen 2004, Van den Bulte and Wuyts 2007, Li 2018). Understanding how the effects of social ads vary across status and nonstatus goods will provide deeper insights into why social advertising operates differently across products. We therefore compared the effects of social influence and social advertising for experience (or search) goods and status (or nonstatus) goods (see the online appendix for definitions).

2.2. Experience vs. Search Goods

The distinction between search goods and experience goods is based on consumers' ability to evaluate product attributes before deciding to purchase (Nelson 1970, Schmalensee 1978). This distinction is related to informational social influence, which is reflected in customers' desire to use others' preferences or behaviors to characterize a product (Burnkrant and Cousineau 1975, Van den Bulte and Lilien 2001). The social influence process for experience goods may involve more information transfer between friends than search goods because the quality of search goods can be evaluated before purchase (after a costly search), whereas the quality of experience goods can be evaluated only by experiencing them or being exposed to the experience of others. Faced with a lack of product information, individuals tend to rely more on the experience of their trusted peers to evaluate experience goods than search goods, which are easier to evaluate using nonsocial information about the product's characteristics, such as simple information found online. This can affect the distribution of product sales, disproportionately benefiting popular products (Brynjolfsson et al. 2011) or creating an echo-chamber effect of positive feedback (Van Alstyne and Brynjolfsson 2005). Experience goods, which exhibit quality uncertainty and performance risks, are therefore more likely to be informed by social influence than search goods, because customers are motivated to avoid uncertainty and risk (Jacoby and Kaplan 1972, Friedman and Friedman 1979). Friends' endorsements provide additional information for ad viewers and reduce uncertainty when deciding whether to engage with a social advertisement. As a result, friends' endorsements, in the form of social cues, may have a greater effect on ad engagement for experience goods than for search goods.

2.3. Status vs. Nonstatus Goods

Social influence is more relevant for status goods, not because we learn about the product and its quality from our friends, but because we evaluate the utility of our purchases by making relative comparisons to our friends in constructing our social identity. Social status, defined as "a position in a social structure based on esteem that is bestowed by others" (Hu and Van den Bulte 2014, p. 510), has been theorized as a powerful driver of consumption choices (Veblen 1899, Bagwell and Bernheim 1996, Corneo and Jeanne 1997, Wang and Griskevicius 2014). People purchase or consume a product not only to directly enjoy it, but also to create and support status differences with others in society. O'Cass and McEwen (2004, p. 14) define status-driven consumption as "the behavioral tendency to value status and acquire and consume products that provide status to the individual."

Possession of material goods signals individuals' social status, and social influence plays a particularly important role in status-driven consumption (Bernheim 1994, Pesendorfer 1995). Consumers are motivated to identify themselves with individuals in their status group, or those with superior status, to maintain or improve their own social standing. Status symbols facilitate the identification process, in which consumers identify themselves with others of a desired social status by consuming the same status goods. People may conform to the product decisions of others to avoid social risks and in exchange for social returns. Consumers care about how others perceive their use of products (Jacoby and Kaplan 1972, Friedman and Friedman 1979). Consuming status goods that higher-status individuals consume can also establish common ground for communication and thus stronger relationships (Kuksov and Xie 2012). For these reasons, friends' endorsements, in the form of social cues, are likely to have a greater effect on ad engagement for status goods than nonstatus goods.

2.4. The Moderating Effects of User Status and Product Involvement

Psychologists and communication scholars have studied the role of source effects in persuasion and attitude change for decades (Chaiken 1980, 1987; Hass 1981; Goldenberg et al. 2001). More recently, research on online social interactions has uncovered individual level heterogeneity in social influence and the role of identity in persuasion on social media (Peters et al. 2013, Berger 2014). The basic argument is that different people have varying degrees of influence over those who follow their social media posts (e.g., Aral and Walker 2012, 2014). These studies show how individual-level and relationship-level characteristics of the sources of online messages can cause viewers to change their opinions about and subsequently change their behavior in response to those messages.

We therefore explore the moderating effects of users' characteristics on social advertising effectiveness. We theorize that the status and prior product involvement of users shown in social ads should have a meaningful effect on the performance of those ads. Levina and Arriaga (2014) argue that individuals develop status online by accumulating social and cultural capital through differential streams of content contributions and favorable social network positions, like network centrality. They further argue that having higher status online leads to various types of power, for instance, preferential treatment of content, more attention, or the ability to influence others. We believe this type of status is likely to affect the persuasive power of users' endorsements of products in social advertising (e.g., Goldenberg et al. 2009). Consumers are not likely to be concerned with the absolute social status of their friends but instead with their relative social status. A product endorsement is likely to be more influential if it comes from a higher-status individual than if it comes from a lower-status individual. However, it is likely that status is not equally meaningful for all products, but instead that it is more meaningful for some products than for others. Status-based consumption is driven by relative comparisons to peers or friends. An endorser's status is likely to be more meaningful if he or she is endorsing a status good than if he or she is endorsing a nonstatus good.

Product involvement, the degree to which one researches or reads about a product or product category, is another critical enabler of source persuasiveness. Previous research finds that product involvement significantly impacts how effective one is as an information source (Godes and Mayzlin 2009, Iyengar et al. 2011). Peers who are more involved with products and have gathered more product information tend to be regarded as credible information sources. It is therefore reasonable to assume that the prior product involvement of an information source will impact the persuasiveness of his or her endorsements.

Consumers are not likely to be concerned with the absolute product involvement of their friends, but instead with their relative product involvement. Knowing that a friend has more experience and information about a product than oneself will likely increase the persuasive power of the friend's endorsement of that product. Whether product involvement enhances or depresses social influence is also likely to depend on product characteristics. Search goods are easily researched by seeking readily available factual information and present little ambiguity of product quality. Experience goods, on the other hand, are only really evaluated through hands-on experience or knowledge of others' hands-on experience. Product involvement may thus be meaningful for evaluations of experience goods, where the experience of peers is more useful and relevant to the product evaluation process. Friends' endorsements, in the form of social cues, may therefore have a greater effect on ad engagement for experience goods than for search goods when the user whose cue is used in the ad has greater involvement with the product than the ad viewer.

3. Experiment

3.1. Experimental Design

We utilized a large-scale randomized field experiment to estimate social influence effects across products. Previous research has endeavored to identify social influence empirically (Tucker 2008, Aral 2009, Bramoullé et al. 2009). However, randomized experiments are the gold standard for causal inference in advertising and social influence research (e.g., Bakshy et al. 2012; Sahni 2015, 2016; Tucker 2016).

Our experiment was conducted on WeChat Moments ads, which are displayed in users' Moments. WeChat Moments is equivalent to the Facebook news feed, which supports posting images and text, as well as sharing music and short videos. WeChat delivers ads to the Moments of targeted users as Facebook does in their news feeds. Users see those ads while scrolling through their feeds and can express their attitudes toward the ads and show their preferences and opinions by liking and commenting on them. They can also engage directly with the ads by clicking through on the links to the advertisers' profile pages, landing pages, and product photos (see Figure 1 for an example of a WeChat Moments ad).[2]

Social influence and social advertising effectiveness in our experiment are measured as the degree to which social cues (i.e., friends' likes), representing friends' endorsements of products, impact users' engagement with social advertising (i.e., click-throughs). There are two types of social cues in Moments ads: friends' likes and comments (see Figure 1). Because comments vary widely in their content (they may be positive or negative about the ads or products), and to cleanly estimate the effects of friends' endorsements, we focused, in this paper, exclusively on the effect of likes and hid all friends' comments on ads in all the

Figure 1. (Color online) An Example WeChat Moments Ad

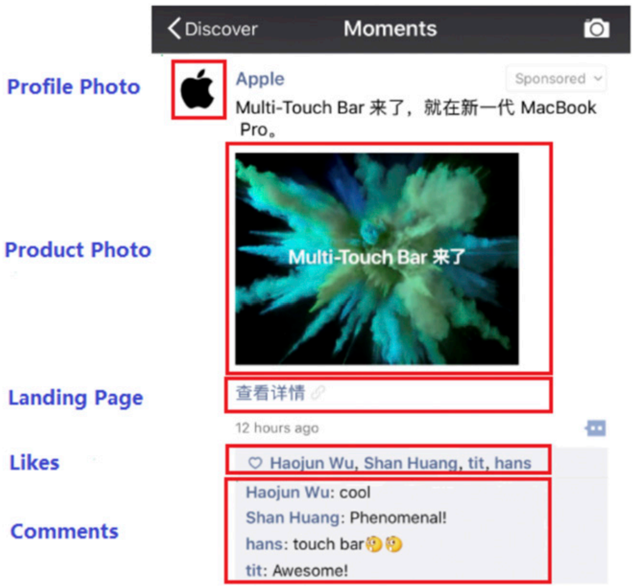

Note. This figure displays an example WeChat Moments ad, including the brand profile photo, product photo, the link to brand landing page, and social likes and comments.

experimental groups during the experiment. If a user clicks "like" below an ad in WeChat Moments, they show their endorsement of the ad to their friends, who will see the like immediately afterward through WeChat Moments. Social likes could be interpreted by viewers as endorsements of the ad, the product, or both. Unlike in Facebook, only *first-degree* friends' likes or comments are visible to WeChat users. As a result, social ad effectiveness and social influence in our experiment have a consistent meaning that reflects the effect of a social cue (i.e., a friend's like), representing first-degree friends' endorsements of an ad, on users' engagement with the ad (i.e., click-throughs).[3]

WeChat Moments targets different user segments for different products before delivering ads. However, our experiment was conducted in the very early stages of WeChat Moments ads.[4] During our experiment, WeChat Moments targeted users solely based on their age, gender, and city. Although different targeting strategies for ads can lead to unbalanced user pools for ads for different products and product types (e.g., women may be more likely to see ads for cosmetics, whereas car ads are often targeted more toward men), the simple three-variable targeting strategies used by WeChat in the early stages of Moments ads allowed us to reliably control for the variables and conditions used in ad targeting.

During the experiment, as users were served new ads, they were *randomly* assigned to three experimental groups, each with an equal probability of 8%—one without any social cue (the control group),[5] one with a maximum of one displayed like (Treatment Group 1), and one with the organic number of likes displayed on the ads (Treatment Group 2)—or they remained outside the experiment with a 76% probability (see Figure 2). In Treatment Group 1, when there was more than one organic like, the first organic like was displayed.[6] Every time a new ad was served to a user, the randomization reoccurred for the new ad. In this way, randomization occurred each time a user received a new ad. Users could be assigned to a different group for each different ad they saw, but stayed in the same group for the same ad. The randomization, therefore, took place at the user–ad level. Every ad stayed in a user's news feed for a maximum of 48 hours. After 48 hours, the old ad would disappear. Users were allowed to see only one ad at a time in WeChat Moments, which also reduced our susceptibility to statistical interference between different ads. Our treatments varied the existence of social cues and the number of social cues separately,

Figure 2. (Color online) Experimental Treatments

Note. This figure illustrates the control condition (without any social cue), the first treatment group (with a maximum of one like), and the second treatment group (with the organic number of likes), from left to right, respectively.

allowing us to independently estimate the effects of exposure to a social cue and multiple social cues displayed on ad engagement.

Our experimental design avoids many known sources of bias in influence identification and networked experiments. First, by randomly assigning social cues, it eliminates bias created by homophily. Homophily can bias estimates of social influence when similarities between nodes create correlated behavioral patterns among friends without direct peer influence (Manski 1993, Aral et al. 2009). Without the experimental manipulation, the ad viewers who are affiliated with more friend endorsers for an ad and are more likely to engage with the ad because of homophily would also be exposed to more social cues. Randomly manipulating the presence and number of social cues shown in ads for user–ad pairs in control and treatment groups separates influence-driven contagion and homophilous diffusion, and breaks the correlation between the number of social cues displayed on ads and the number of friend endorsers affiliated with ad viewers. Observed and unobserved attributes of users are equally distributed across control and treatment groups in our experiment.

Second, randomization controls for external confounding factors because users are equally likely to be exposed to external stimuli that could affect engagement across treatment groups. Third, all of the ads involved in the experiment were new and distinctive, so users could not have been exposed to the ads through any external sources before or outside the experiment. Fourth, likes from different users were shown in identical formats in Moments and are different only in friends' names or profile pictures, eliminating the heterogeneity of immeasurable characteristics of social cues. Fifth, because of the one-ad limit every 48 hours, users would not receive different treatments from different ads at the same time. Because randomization reoccurred every 48 hours, it is unlikely that users noticed they were being treated during the experiment. The treatment effects, therefore, were not confounded by habituation or users who suspected they were in an experiment. Our design also avoids statistical interference and guarantees that the stable unit treatment value assumption is met (Rubin 1990).

3.2. Data Collection

During data collection, we recorded the number of organic social cues (organic likes), the number of social cues displayed on the ads due to treatment (displayed likes), the exact friends who were shown in the ads and to which viewers, ad viewers' responses to the ads (whether they clicked and their response times), and ad viewers' and their friends' demographics (age, gender, and city), network degree (i.e., number of WeChat friends), and behavioral characteristics on WeChat (see Sections 3.2.2 and 3.2.3). For our main analysis, we considered users' responses only during their first impressions of a new ad and measured social influence as the effect of displaying exactly one social cue on users' likelihood of clicking on an ad. Most of the ad responses happened during users' first impressions, which is also the least confounded measure of user ad engagement. We later relaxed the first impression constraint and found that our results are robust to estimates beyond the first impression. We counted any click on an ad as long as the ad viewer clicked on the brand's profile page, the ad landing page, or the product photos (all of which were displayed on the ad). We also collected data about the ads, including product and brand names and the product's category. We adopted the product categorization used by the WeChat ad department, which is a standard one used in the advertising industry.

3.2.1. Measuring Product Types. We used multiple raters to classify the 71 experimental products on the two theoretically motivated dimensions in our study: experience versus search goods and status versus nonstatus goods, based on definitions that we provided them (see Appendix A for these definitions). Four independent judges, all undergraduate economics students at a prestigious Chinese university, separately classified the 71 products. Their interrater agreement ranged from 0.87 to 0.99 as measured by the intraclass correlation and between 0.62 to 0.83 by Fleiss' kappa. The four judges resolved all disagreements by consensus.

3.2.2. Measuring Users' Status. We use degree centrality, which, in our case, is the number of WeChat friends, to measure the social status of a user (Hu and Van den Bulte 2014). WeChat facilitates undirected network ties requiring both parties to mutually approve one another before becoming friends. The undirected network degree is similar to the in-degree centrality of directed networks, which has been a popular measure of status in the networked marketing literature (e.g., Iyengar et al. 2011, Hu and Van den Bulte 2014). In-degree centrality is measured as the number of incoming ties in directed social networks (Jackson 2008), which reflects the extent to which one is desired and respected by others (social status). Because network degree follows a power-law distribution and has a long tail, we log transformed network degree to reduce the skew of its distribution. Relative social status is measured as the friends' status minus the ad viewer's status. To make the related coefficients more interpretable, we also standardized the measure of friends' relative social status in our analysis.

3.2.3. Measuring Users' Product Involvement. Product involvement is measured by the accumulated number of articles that an individual has read about a given product or category on WeChat over the past half year (May 2015 to November 2015).[7] The WeChat team used latent Dirichlet allocation to model users' involvement in different fields, such as finance, technology, and fashion, according to their historical reading behavior. The inputs to the model were the articles that a user read on WeChat, and the output was a vector of scores that measure the user's involvement in various fields. If a product category matches a field, we use the user's score in that field to represent his or her involvement in that product category on WeChat. Relative product involvement is measured as the friends' involvement minus the ad viewers' involvement. To make the related coefficients more interpretable, we also standardized the measure of friends' erelative product involvement in our analysis.

4. Analysis

4.1. Model Specification and Estimation: Product Types

In the main analysis, we estimated social ad effectiveness as the effect of displaying one social cue to an ad on users' ad engagement during their first ad impressions. This is one of the most relevant measures of social advertising effectiveness because it is not confounded by, for instance, variation in the number of organic likes a product or brand receives, or by the multiple impressions of an ad to which a user may be exposed. In robustness checks, we extended our analysis using the alternative measures, such as how social cues impacted ad viewers' responses during the entire span of their ad impressions and how the organic number of social cues an ad received impacted ad engagement.

We specified a logistic regression model to estimate the heterogeneous effects of social ads across product types at the user–ad level, as shown in Equation (1).[8] Each observation represents a user–ad pair in the control group and Treatment Group 1. All the model-based analyses in both the main analysis and robustness checks always pertain to two experimental groups: the control and a particular treatment group. The model simultaneously estimates the impacts of two dimensions of product types: search/experience goods and status/nonstatus goods, as shown here:

$$\log\left(\frac{\Pr(Y_{ij}=1)}{1-\Pr(Y_{ij}=1)}\right) = \alpha_0 + \alpha_j + \eta_t + \beta_1 S_{ij} + \gamma_1(S_{ij}\times StG_j) + \gamma_2 \times (S_{ij}\times ExG_j) + \boldsymbol{C}'_{\boldsymbol{ij}}\boldsymbol{\theta_1} + (S_{ij}\times \boldsymbol{C_{ij}})'\boldsymbol{\theta_2}, \quad (1)$$

where Y_{ij} is a dummy variable indicating whether user i clicked on ad j during his or her first ad impression. The term t indicates the week in which ad j was delivered. In our main analysis, S_{ij} is a dummy variable that indicates whether a user–ad pair (i, j) is in Treatment Group 1 for ad_j. The variables StG_j and ExG_j indicate whether the product of ad j is a status good or a nonstatus good and whether it is an experience good or a search good. The coefficient β_1 captures the marginal effect of social cues on ad engagement. The variables γ_1 and γ_2, which capture the impact of product types on the effectiveness of social ads, are our main interest.

The term C_{ij} is a vector of control variables that represent the age, gender, and city[9] of user i and his or her affiliated friend who generated the first social cue (i.e., like) for user–ad pair (i, j), as well as whether an ad is associated with a big brand. The social cue of the affiliated friend was displayed for user–ad pair (i, j) in Treatment Group 1 and was hidden for user–ad pair (i, j) in the control group. We include not only C_{ij}, but also their interactions with the treatment group of a user–ad pair (whether they are treated with a social cue displayed in the ad), $S_{ij}\times C_{ij}$, to account for their effects on both clicking and social influence.

Previous studies suggest that individual demographic characteristics significantly affect the magnitude of social influence (e.g., Aral and Walker 2012). Because different ads target different users, it is necessary to control for the variables used for ad targeting. Our experiment was conducted at the very early experimental stages of WeChat Moments ads. The ad targeting conditions used during our experiment were based simply on users' age, gender, and city, which we controlled for to further reduce the confounding effects of ad targeting on social ad effectiveness. We controlled for characteristics of the affiliated friends whose social cues were displayed in the ads. Targeting conditions also affect the demographics of the affiliated friends. Both influence and susceptibility are key factors that drive social influence and social ad effectiveness. As a result, it is necessary to take affiliated friends' demographic characteristics into account. Finally, brand characteristics have also been shown to affect word of mouth (Lovett et al. 2013) and may affect social influence. We therefore used a dummy variable to indicate whether a brand was among the 100 Best Global Brands, as rated by Interband, to control for big brand effects on clicking and influence.

We added ad-specific fixed effects, α_j (ad dummies), to control for variation in users' engagement caused by ad characteristics. We did not include any person-specific effect, because the average number of observations per user in the control group and Treatment

Group 1 is only 1.096. Users' adoption outcomes may be correlated for the same ads, which share the same design and were delivered during the same time. To account for this, we specified clustered standard errors at the ad level (Cameron and Miller 2015). Finally, because our experimental period covered the Christmas and New Year holidays, we included week fixed effects, η_t (week dummies that indicate the week user *i* expose to ad *j*), to control for time effects. The coefficient θ_1 captures the variation in ad engagement explained by this vector of control variables. The coefficient θ_2 represents the effects of the age, gender, and city of users and their affiliated friends, and brand characteristics on social ad effectiveness.

4.2. Modeling the Moderating Effects of Users' Status and Product Involvement

We first specify a logistic regression model for experience, search, status and non-status goods separately, as in Equation (2). The model estimates the impacts of the relative social status and product involvement between the user shown in the ad and the ad viewer on ones' tendency to engage with the ads for different types of products as follows:

$$\log\left(\frac{\Pr(Y_{ij}=1)}{1-\Pr(Y_{ij}=1)}\right) = \alpha_0 + \alpha_j + \eta_t + \beta_1 S_{ij} + \beta_2 St_{ij} + \beta_3 In_{ij} + \gamma_1(S_{ij}\times St_{ij}) + \gamma_2(S_{ij}\times In_{ij}) + \boldsymbol{C}'_{\boldsymbol{ij}}\boldsymbol{\theta_1} + (S_{ij}\times \boldsymbol{C_{ij}})'\boldsymbol{\theta_2}, \quad (2)$$

where St_{ij} and In_{ij} indicate the social status of the affiliated friend *j* relative to user *i* (i.e., the status of friend *j* minus the status of user *i*) and the product involvement of the affiliated friend *j* relative to user *i* (i.e., friend *j*'s involvement minus user *i*'s involvement in that product category). The coefficients γ_1 and γ_2 in the interaction terms capture the impact of relative social status and involvement on social advertising effectiveness. All the control variables used in Equation (1) are included in C_{ij}. We also control for whether the product in the ad is an experience good (ExG_j) for the sample of status and nonstatus goods, and control for whether the product in the ad is a status good (StG_j) for the sample of experience and search goods.

We then estimate the logistic regression model in Equation (3) to test whether the γ_1 and γ_2 coefficients in Equation (2), the impact of the relative social status and product involvement of the friend featured in the ad compared with the status and product involvement of the ad viewer on ad performance, statistically significantly vary between experience and search goods, and between status and nonstatus goods:

$$\begin{aligned}\log\left(\frac{\Pr(Y_{ij}=1)}{1-\Pr(Y_{ij}=1)}\right) &= \alpha_0 + \alpha_j + \eta_t + \beta_1 S_{ij} + \beta_2 St_{ij} + \beta_3 In_{ij} + \beta_4 P_j \\ &+ \gamma_1(S_{ij}\times St_{ij}) + \gamma_2(S_{ij}\times In_{ij}) + \gamma_3(S_{ij}\times P_j) \\ &+ \gamma_4(St_{ij}\times P_j) + \gamma_5(In_{ij}\times P_j) + \pi_1(S_{ij}\times St_{ij}\times P_j) \\ &+ \pi_2(S_{ij}\times In_{ij}\times P_j) + \boldsymbol{C}'_{\boldsymbol{ij}}\boldsymbol{\theta_1} + (S_{ij}\times \boldsymbol{C_{ij}})'\boldsymbol{\theta_2} \\ &+ (P_j\times \boldsymbol{C_{ij}})'\boldsymbol{\theta_3} + (P_j\times S_{ij}\times \boldsymbol{C_{ij}})'\boldsymbol{\theta_4}, \quad (3)\end{aligned}$$

where P_j represents ExG_j when we compare the effects between experience and search goods, and P_j represents StG_j when we compare the effects between status and nonstatus goods. The coefficients π_1 and π_2 on the three-way interaction terms capture the difference in γ_1 and γ_2, the impact of relative social status and involvement on social advertising effectiveness, across product types.

5. Empirical Results

5.1. Descriptive Statistics

The experiment was conducted over a 21-day period starting in December of 2015, during which 57,605,029 user–ad pairs, 37,985,501 distinct users, and 99 ads participated in the experiment. A total of 19,198,166 user–ad pairs were randomly assigned to the control group with no social cues, 19,201,745 user–ad pairs were randomly assigned to the treatment group displaying a maximum of one like, and 19,205,118 user–ad pairs were randomly assigned to the treatment group displaying the organic number of likes. Assignment to experimental groups was random, with no statistically significant mean differences between the three experimental groups in terms of users' age, gender, city, network degree (i.e., number of WeChat friends), and level of WeChat Moments activity (i.e., log-in days; *F*-tests, $p > 0.1$). WeChat rolled out its Moments ads for a limited number of products, brands, and users in 2015, and on average, each user was exposed to fewer than two ads during our 21-day experiments.

We dropped 17 ads with invalid data and analyzed experimental results on 82 ads for 71 distinct products across 25 product categories.[10] Among the 71 products, 48 were experience goods, 23 were search goods, 22 were status goods, and 49 were nonstatus goods (see Table 1). We also excluded data with an incorrect number of displayed likes.[11] This ensured the integrity of our manipulation: no like was displayed to users in the control group, and one like or the organic

Table 1. Product Types Across Products

	Experience goods	Search goods
Status goods	17 products	5 products
Nonstatus goods	31 products	18 products

Note. Product types are measured at the product level.

Table 2. Manipulation Checks

			# displayed likes			
# likes	# user–ad pairs	# users	Mean	SD	Max	Min
0	1,860,622	1,775,820	0.000	0.000	0	0
1	1,873,401	1,787,240	1.000	0.000	1	1
2	1,837,093	1,755,895	1.672	1.742	100	1

Note. In the table, "0" represents the control group, "1" represents Treatment Group 1, and "2" represents Treatment Group 2.

number of likes was correctly displayed to the users in the two treatment groups (see Table 2).

The number of social cues we can display on ads is limited by the number of organic likes posted by friends of the ad viewer. We did not generate fake likes but manipulated real likes. Some ads had no organic likes, and we were therefore unable to display any real social cues on these ads as part of our manipulation. We excluded the user–ad pairs with zero organic likes and filtered the data on the condition that there was at least one organic like, in *both* control and treatment groups, to guarantee that exactly one social cue could be displayed on ads in Treatment Group 1, at least one social cue could be shown on ads in Treatment Group 2, and that users were equally distributed across control and treatment groups after the filtering.

This process created a sample of 82 ads for 71 distinct products across 25 product categories, 5,571,116 user–ad pairs, and 4,884,070 distinct users across three treatment groups: 1,860,622 user–ad pairs in the control group, 1,873,401 user–ad pairs in Treatment Group 1, and 1,837,093 user–ad pairs in Treatment Group 2. There are no economically meaningful mean differences between these three groups in terms of their age, gender, city (i.e., a first-, second-, or third-class city), network degree (i.e., the number of WeChat friends), or level of WeChat Moments activity (i.e., the number of log-in days in November 2015, the month before the experiment; see Tables 2 and 3; Ding and VanderWeele 2016).

5.2. Average Effects of Social Advertising

In the main analysis, we estimated social ad effectiveness as the effect of showing one social cue to an ad on users' ad engagement during their first ad impressions. We first report average treatment effects in social advertising, which we define as displaying social cues to advertisements in WeChat Moments (users' news feeds). We estimated marginal social ad effectiveness across all user–ad pairs as the relative risk of users' average response rates (click-throughs) across control and treatment groups during users' first impressions on ads. Control group ad units were displayed without any social cues, whereas treatment group units were displayed with one social cue (i.e., a randomly chosen friend's like). We found that the social influence enabled by social cues significantly improved ad effectiveness. Displaying a social cue (a like) made users 33.75% more likely to click an ad on average ($p < 0.01$).

Our estimate of average social ad effectiveness is larger than those found in some prior studies. But, several differences between our experiment and prior work can explain the differences in magnitudes. First, our experiment compared users' engagement across ads without any social cues to ads with one or multiple social cues. In contrast, prior work, for example, Bakshy et al. (2012), compared groups with one, two, and three social cues or a group displaying the total number of endorsements (e.g., 100 people like this) against a group displaying one social cue. As prior work did not include a baseline group without any displayed social cues, their effect sizes were understandably smaller, for example, ranging from 3.8% to 10.5% in Bakshy et al. (2012). We believe the difference between one social cue and no social cues conveys dramatically more social proof and therefore creates a further step in the effect of social cues on social advertising effectiveness.

5.3. Heterogeneous Effects of Social Advertising Across Products

We are interested in characterizing the heterogeneity in social advertising effectiveness across products. Figure 3 displays the marginal effect of socializing ads for each the 71 distinct products and confirms that social influence lifts the click-through rates for most products. Thirty-nine out of 71 products exhibit statistically significantly positive lift from social advertising (see the dots in Figure 3), whereas 32 products experience no statistically significant lift (see the crosses in Figure 3), and none perform worse when social cues are added to the ads. Displaying a friend's like in an ad causes up to a 270% increase in the click-through rate for a social advertisement (see the highest dot on the right in Figure 3). There is significant heterogeneity in social advertising effectiveness across products. The highest product-level social ad effectiveness (max influence = 3.70, $p < 0.01$) is 2.64 times as large as the

Table 3. Mean Comparisons Between Control and Treatment Groups

	|#0 – #1|			|#0 – #2|			|#1 – #2|		
	t-statistic 1	*t*-statistic 2	SD	*t*-statistic 1	*t*-statistic 2	SD	*t*-statistic 1	*t*-statistic 2	SD
Age	19.627	–28.502	0.0058	13.932	–34.009	0.0058	18.363	–29.620	0.0058
Gender	2.065	–6.617	0.0005	4.540	–4.075	0.0005	6.817	–1.812	0.0005
City	13.027	–15.047	0.0006	11.743	–16.201	0.0006	12.773	–15.226	0.0006
Log(*Network degree*)	39.356	–32.491	0.0003	51.760	–19.791	0.0003	48.447	–23.270	0.0003
Login days	138.956	–139.316	0.0011	138.765	–139.230	0.0011	139.313	–139.419	0.0011

Notes. In the column headings, "0" represents the control group. "1" represents the treatment group 1. "2" represents the treatment group 2. We compared means between three groups and ran equivalence test H0: $|\mu_a - \mu_b| > \Delta, \Delta = 0.5\%\bar{X}_b$, which is two-one-sided t-tests. The H0 of t-test 1 is $\mu_a - \mu_b < -\Delta$. The H0 of t-test 2 is $\mu_a - \mu_b > \Delta$ (Lakens et al. 2018). The values of City can be 1,2 or 3, which indicate the first, second, or third class of cities. All the one-sided t-tests in this table show statistically significant results and are therefore rejected at the 5% level of significance. When these one-sided tests are statistically rejected, we can conclude that $|\mu_a - \mu_b| \leq \Delta$. Please note that our dataset with a large *N* allows the tests to detect tiny differences across groups with high statistical power.

average product-level effectiveness (average influence = 1.40, $p < 0.01$), and 3.06 times as large as the lowest positive product-level influence (lowest influence = 1.21, $p < 0.01$; see the lowest dot on the left in Figure 3).

Next, we aggregated the products into 25 categories and identified influence at the product category level (see Figure 4). Nineteen categories exhibit significantly positive lift from social advertising, whereas 6 experience no statistically significant lift from social ads. Food products exhibit the greatest lift, with social ads in the food category causing an 84% relative increase in the click-through rates on ads in that category (average marginal influence = 1.84, $p < 0.01$). Showing a friend's like on an ad for food is 1.64 times more effective than doing so for mobile games (average marginal influence = 1.12, $p < 0.01$), 1.57 times more effective than doing so for electrical appliances (average marginal influence = 1.17, $p < 0.01$), and 1.55 times more effective than doing so for financial services (average marginal influence = 1.18, $p < 0.01$). Baby food, clothes, and cars are other categories in which social influence creates a large increase in advertising effectiveness, whereas TV shows, e-commerce platforms, and credit cards do not exhibit statistically significant lift from social ads. In summary, we found a highly heterogeneous lift distribution from social ads across products and product categories. The results

Figure 3. (Color online) Social Advertising Effectiveness Across Products

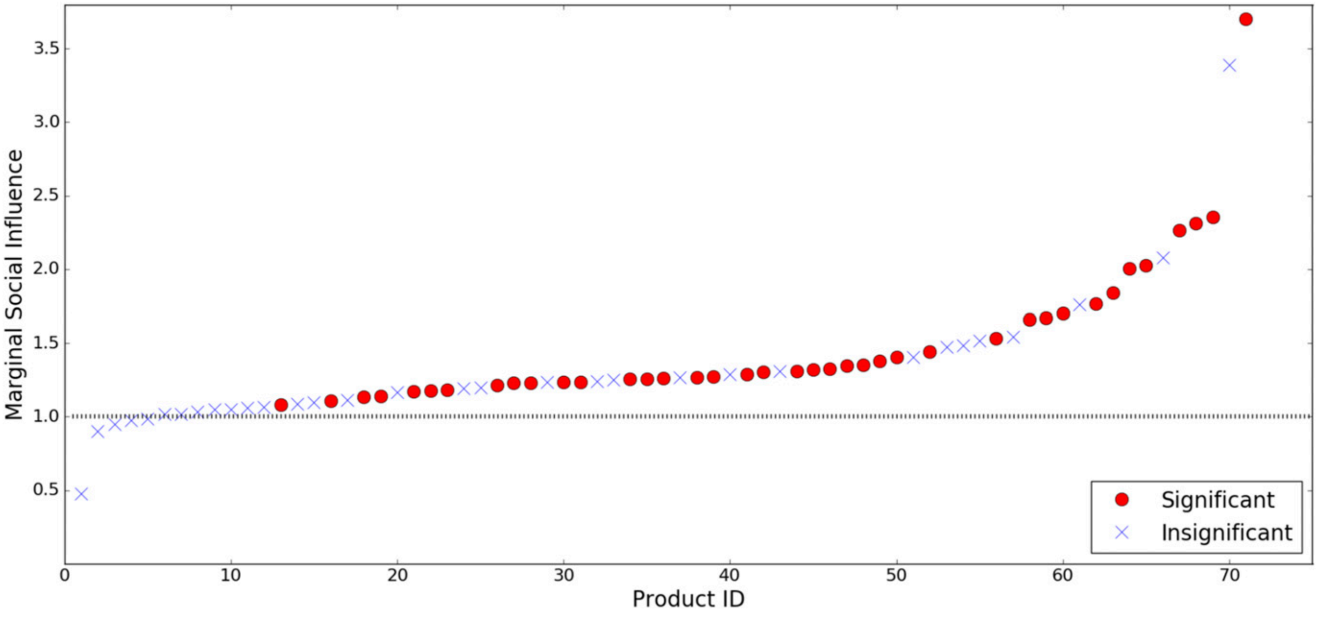

Note. We ordered the products in ascending order according to their marginal social ad effectiveness, representing the relative increase in click-through rates caused by displaying one like in the ad for that product.

Figure 4. (Color online) Social Advertising Effectiveness by Product Category

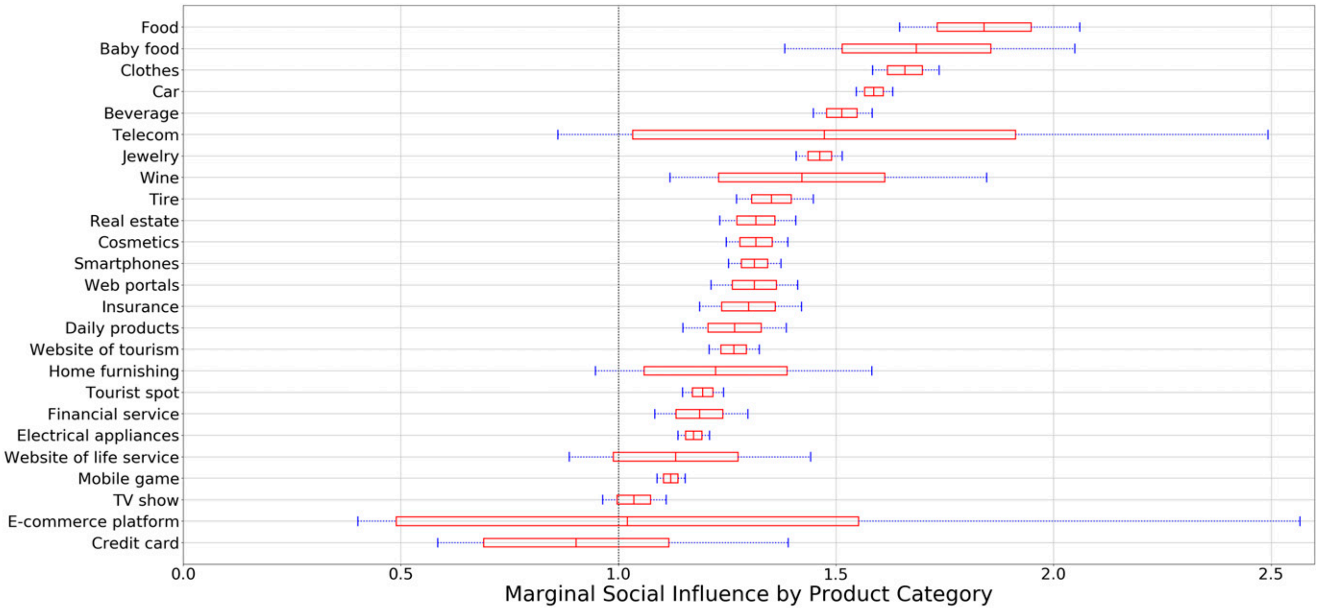

Note. Marginal social ad effectiveness in each product category is shown with standard errors (boxes) and 95% bootstrapped confidence intervals (whiskers).

in Figures 3 and 4 utilize bootstrapped confidence intervals, which are robust to correlated observations. For example, ad viewers' behaviors may be correlated for the same ads, which were delivered during the same time and share the same design (Bakshy and Eckles 2013).

5.4. How Product Types Moderate Social Advertising Effectiveness

The evidence in the previous section establishes significant heterogeneity in the effectiveness of social advertising across products and product categories. The natural next step in our investigation then is to study why this heterogeneity exists. In Section 2, we argued that the theoretically motivated dimensions of status/nonstatus goods and experience/search goods may moderate the importance of social cues for consumer decisions. In this section, we provide experimental evidence for the moderating effects by estimating the impact of these two product dimensions on the effectiveness of social advertising. We evaluate data from the control group with no social cues and the treatment group with one social cue (a like) to estimate the relative marginal effect of socializing ads for search/experience goods and for status/nonstatus goods.

Estimates of the impact of these product types on social advertising effectiveness are displayed in Table 4, whereas Figure 5 compares the e^{γ_1} estimates for search/experience goods and status/nonstatus goods and displays the standard errors (boxes) and 95% confidence intervals (whiskers) of these estimates. The e^{γ_1} estimates represent relative risk of marginal social advertising effectiveness across these different product types. The odds approximate the probability ($\Pr(Y_{ij}=1)$), when the probability ($\Pr(Y_{ij}=1)$) is near zero. The click-through rates in different experimental groups for different types of products are well below 10% in our experiment. We therefore use odds ratio to approximate relative risks.

The results corroborate the average marginal effectiveness of social advertising found in the previous section. Displaying a social cue lifts the click-through rate for social ads by 16.36% ($e^{0.1515}, p < 0.05$) in this analysis, down from the 33.75% in earlier analyses with fewer covariates. We also found that social advertising is 14.82% ($e^{0.1382}$) more effective for status

Table 4. Social Advertising Effectiveness Across Search/Experience Goods and Status/Nonstatus Goods

	1	2
	Ad clicks	Ad clicks
Social cue	0.1907***	0.1515**
	(0.0240)	(0.0669)
SC × Experience goods	0.0339	0.0604
	(0.0606)	(0.0684)
SC × Status goods	0.1807**	0.1382**
	(0.0721)	(0.0659)
Controls		Yes
Log-likelihood	–534,818	–479,761
Observations	3,734,023	3,734,023

Notes. SC, Social cue. An observation is a user–ad pair. Robust standard errors, clustered at the ad level, are reported in parentheses.
$^{**}p < 0.05$; $^{***}p < 0.01$.

Figure 5. (Color online) Social Ad Effectiveness Across Search (Experience) Goods and Status (Nonstatus) Goods

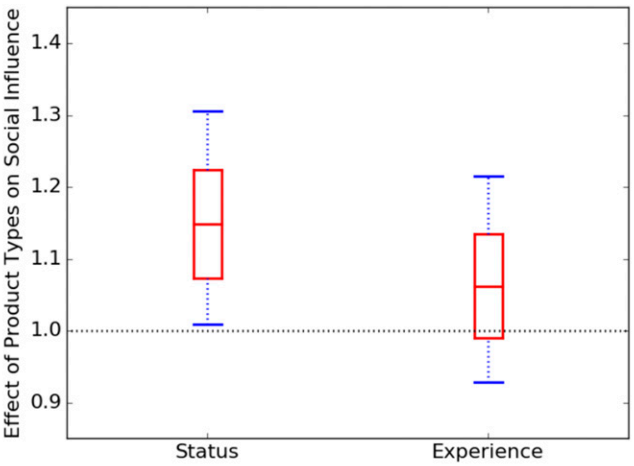

Note. The effects of product types on marginal social ad effectiveness are shown with standard errors (boxes) and 95% confidence intervals (whiskers).

goods than for nonstatus goods ($p < 0.05$), whereas there is no statistically significant difference in the performance of social advertising between experience and search goods ($p > 0.1$). When peer endorsements are presented in advertisements for status goods, they cause an increase in ad effectiveness. In contrast, the results also suggest that informational social influence is a weak driver of social advertising effectiveness. When peer endorsements are presented in advertisements for experience goods, the ads perform about the same as when social cues are added to ads for search goods. Consumers seem to be more affected by peer influence through their consideration of social status than their desire to seek an endorsement of a product experience. It could be, however, that although there is no difference in social advertising effectiveness across experience and search goods on average, differences emerge when we consider the relative status and involvement of the friend shown in the ad. We explore this possibility in the next section.

5.5. How User Status and Product Involvement Moderate Social Advertising Effectiveness Across Products

In this section, we explore how the social status and product involvement of the friends shown in ads, relative to the ad viewers, moderate social advertising effectiveness, and how these moderating effects vary across products. We find the relative social status and product involvement of friends shown in ads are critical to social advertising effectiveness.

Models 1 and 2 in Table 5 and Figure 6 display the moderating effects of the status and product involvement of the friend shown in a social ad, relative to the viewer of the ad, on social advertising effectiveness for status and nonstatus goods. Friends' relative social status significantly improves social ad effectiveness for both status and nonstatus goods. Friends exert 4.67% ($e^{0.0456}$, $p < 0.01$) more influence on viewers' ad engagement for status goods when the status difference between the friends increases by one standard deviation (SD). Friends exert 2.64% ($e^{0.0261}$, $p < 0.05$) more influence on viewers' ad engagement for nonstatus goods when the difference in status between the friend and the viewer increases by one standard deviation. Relative product involvement significantly moderates social ad effectiveness for both status ($p < 0.01$) and nonstatus goods ($p < 0.01$). For status goods, friends exert 4.14% ($e^{0.0406}$, $p < 0.01$) more influence on viewers' ad engagement when their relative product involvement increases by one standard deviation, compared with the ad viewer.

Figure 6. (Color online) Effects of Users' Status and Product Involvement on Social Advertising Effectiveness Across Products

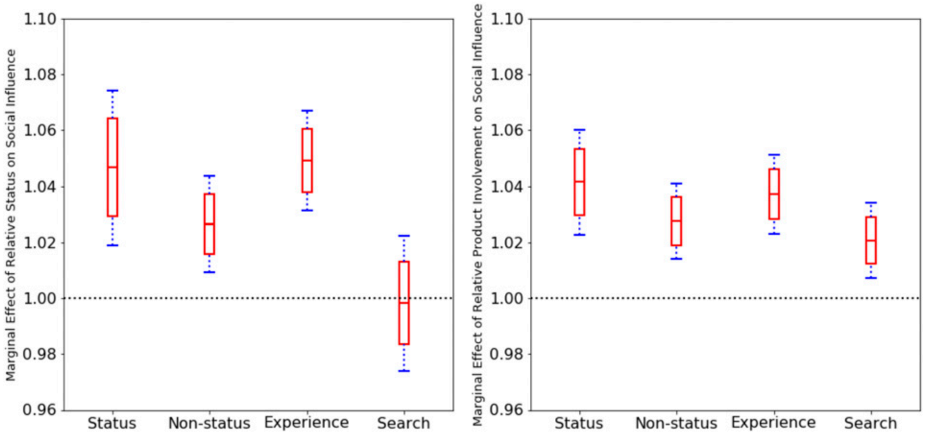

Note. The effects of the status and product involvement of the friend shown in a social ad, relative to the viewer of the ad, on marginal social ad effectiveness for different types of goods are shown with standard errors (boxes) and 95% confidence intervals (whiskers).

Table 5. Effects of Users' Status and Product Involvement on Social Advertising Effectiveness Across Products

	1	2	3	4	5	6
	Status good ad clicks	Nonstatus good ad clicks	Experience good ad clicks	Search good ad clicks	Status vs. nonstatus ad clicks	Experience vs. search ad clicks
SC × *FRS*	0.0456***	0.0261**	0.0479***	−0.0018	0.0261**	−0.0018
	(0.0168)	(0.0104)	(0.0110)	(0.0147)	(0.0104)	(0.0145)
SC × *FRI*	0.0406***	0.0271***	0.0365***	0.0204**	0.0271***	0.0204**
	(0.0115)	(0.0082)	(0.0085)	(0.0081)	(0.0082)	(0.0080)
SC × *FRS* × *SG*					0.0195	
					(0.0196)	
SC × *FRI* × *SG*					0.0135	
					(0.0140)	
SC × *FRS* × *EG*						0.0497***
						(0.0182)
SC × *FRI* × *EG*						0.0161
						(0.0117)
Log-likelihood	−249,990	−227,862	−368,998	−109,018	−477,852	−478,016
Observations	2,387,250	1,346,773	3,215,964	518,059	3,734,023	3,734,023

Notes. SC, Social cue; FRS, friend's relative status; FRI, friend's relative involvement; *SG*, status good; EG, experience good. An observation is a user–ad pair. Robust standard errors, clustered at the ad level, are reported in parentheses.
**$p < 0.05$; ***$p < 0.01$.

For nonstatus goods, friends exert 2.75% ($e^{0.0271}$, $p < 0.01$) more influence on ad viewers when the difference in the product involvement between them increases by one standard deviation.

Models 3 and 4 in Table 5 display the impact of friends' relative status and product involvement on social ad effectiveness for search and experience goods. Friends' product involvement significantly moderates the effect of social cues in ads for experience goods ($p < 0.01$) and search goods ($p < 0.05$). For experience goods, the lift from social cues increases by 3.72% ($e^{0.0365}$, $p < 0.01$) when the product involvement gap for the product between the friend and the ad viewer increases by one standard deviation. For search goods, the lift from social cues increases by 2.06% ($e^{0.0204}$, $p < 0.05$) when the product involvement gap for the product between the friend and the ad viewer increases by one standard deviation.

Friends' relative status significantly moderates the lift from social cues for experience goods ($p < 0.01$) but not for search goods ($p > 0.1$). For experience goods, friends exert 4.91% ($e^{0.0479}$, $p < 0.01$) more influence on viewers' ad engagement when the difference in their status increases by one standard deviation. This suggests consumers value the endorsement of their higher-status friends more than their lower-status friends for experience goods but not for search goods. High-status friends play a significant role not only when consumers evaluate and communicate social status but also when they face product uncertainties and thus desire more information about products. Consumers simply seem to trust high-status individuals more in these situations. Whereas friends' relative status significantly moderated the effectiveness of social ads for experience goods but not search goods, we found the effects of relative status and involvement on social advertising effectiveness did not statistically significantly vary across the other types of products (see Model 5 and 6).

5.6. Robustness Checks

We conducted multiple tests to ensure the robustness of our findings to alternative measurement approaches and model specifications.

First, we added the dyadic similarities between friends and ad viewers as additional control variables in identifying the moderating effects of status and involvement on social ad effectiveness. It is also valuable to examine how the similarities between friends and ad viewers, in their age, gender, and location, moderate the effects of social cues on ad engagement. Dyadic similarity was shown in prior work (e.g., Aral and Walker 2014) to be a significant moderator of social influence. We therefore used three dummy variables to indicate whether friends and ad viewers share the same age range, same gender, and same city class. We found that after adding the three variables representing the dyadic similarities between friends and ad viewers, the results on the moderating effects of relative status and involvement did not change significantly. We also observed that showing friends with a different gender and showing those from the same class of city are significantly more influential in affecting users to engage with social ads, for experience and status goods (see Table B.1).

Second, we included tie strength between friends and ad viewers as an additional control variable in our analysis to increase the internal validity of our model. Tie strength has been shown to be a critical factor that significantly impacts social influence (Xiang et al. 2010, Bakshy et al. 2012, Aral and Walker 2014, Berger 2014). Tie strength indicates the intensity of the relationships as well as the similarity (homophily) between friends, both of which can impact peer influence. At the same time, tie strength between friends and ad viewers can also correlate with friends' social status and product involvement relative to the ad viewers.

Tie strength was measured by a dummy variable indicating whether the friend whose like was shown in ads had sent (received) messages to (from) the ad viewer during the month before the experiment. Table B.2 shows that the results remain consistent after controlling for tie strength. Moreover, we find that tie strength greatly and significantly moderates social ad effectiveness for all types of the products. Close friends exert 28.97% ($e^{0.2544}$, $p < 0.01$) more influence for status goods, 15.63% ($e^{0.1452}$, $p < 0.01$) more influence for nonstatus goods, 24.06% ($e^{0.2156}$, $p < 0.01$) more influence for experience goods, and 15.93% ($e^{0.1478}$, $p < 0.01$) more influence for search goods on viewers' ad engagement. The moderating effect of tie strength in social ad effectiveness is significantly larger for status than nonstatus goods. These results not only indicate the robustness of our analysis but also confirm the importance of the role of tie strength in moderating social influence and social advertising effectiveness for various types of products (Aral and Walker 2014).

Third, the alternative measures of social influence, such as how social cues impact ad viewers' responses during the entire span of their ad impressions (mean impressions = 3.03, SD = 2.52), are also important. Although our main analysis measures the more precise and unconfounded marginal effect of showing social cues to the first impression of a social ad, responses during the entire span of an ad's impressions evaluate users' decisions over a longer time horizon. We therefore adopted a new dependent variable in our robustness checks by counting any click on a given ad as long as users clicked the ad at any point during the time the ad was shown in their WeChat news feed. Note that the ads stayed in the WeChat news feed for no more than 48 hours during the experiment.

Results of analyses with the new dependent variable are consistent with our main findings. Displaying a social cue in ads was 11.67% ($e^{0.1104}$) more effective for status goods than nonstatus goods ($p < 0.1$), but there was no statistically significant difference in the performance of social ads across experience and search goods ($p > 0.1$) using all ad impressions (see Tables 4 and B.3). It seems that users were less affected by the opinions of their friends with higher status but were more influenced by the social cues of friends with greater involvement, when we considered their responses during the entire span of ad impressions (compare the first and second rows between Tables 5 and B.4). These results may imply that the influence of friends' relative status in ads is stronger on a first impression than after a long span of consideration and reinforcement, but that ad viewers are more persuaded by friends' product involvement during a longer span of consideration than on a first impression. Putting these nuances aside, the general pattern of results and the importance of peers' status and involvement are confirmed by this robustness analysis.

Fourth, it is valuable for marketers to understand the influence of the organic number of social cues to grasp the full magnitude of social ad effectiveness. We would benefit from knowing whether influence is stronger for more socially liked advertisements and products. To examine social influence created by the organic number of social cues an ad received (mean = 1.672, SD = 1.742; see Table 2), we replicated our estimation comparing the control group with no social cues to Treatment Group 2 with the organic number of social cues.[12] The results are presented in Tables B.5 and B.6. We found that the impact of product types (status/nonstatus and experience/search goods) on social influence does not change significantly when we estimate social ad effectiveness using the organic number of likes.

Displaying organic social cues was 20.42% ($e^{0.1858}$) more effective for status goods than for nonstatus goods ($p < 0.01$), but there, again, was no statistically significant difference in social ad effectiveness across experience and search goods ($p > 0.1$). Status goods experienced an even larger social influence effect when showing organic social cues in ads (20.42% > 14.82%) than when we showed only one social cue. This may be due to the additional influence of multiple sources of social proof. Consistent with our previous results, users were significantly more influenced by the social cues of friends with higher status and greater product involvement in most of the cases. The effect of the relative social status of organic social cues on social advertising effectiveness was significantly greater for status goods than for nonstatus goods and for experience goods than search goods ($p < 0.01$). We observe that friends' social status became more influential when more friends were displayed on social ads for status goods, but not for nonstatus goods (see Tables 5 and B.6). Further, we replicated our analysis comparing Treatment Group 1 with one social cue and Treatment Group 2 with the organic number of social cues. The results are qualitatively consistent (see the online appendix for more details).

Fifth, 8.64% of the users in the sample for our main analysis viewed more than one ad and were therefore associated with multiple observations. Ignoring correlations among behaviors of the same user across different ads may underestimate the standard errors. However, we do not have sufficient within-subject variability to specify fixed or random effects models at the individual level. As a robustness check, we therefore dropped the 8.64% of users who viewed multiple ads and replicated our analysis on the remaining sample of users (91.36%) who were exposed only to a single ad during our experiment.[13] The results from our robustness analysis in Tables B.7 and B.8 are very consistent with those in our main analysis.

Sixth, because of ad targeting, users' characteristics are not balanced across different ads. Instead of controlling for the variables (age, gender, and city) used for targeting in our main analyses, we directly controlled for the 44 targeting conditions used by the 82 ads involved in our experiment.[14] Table B.9 shows that controlling for the exact targeting conditions and the variables used in targeting leads to almost identical results, confirming the robustness of our results.

6. Discussion

We conducted a very large-scale randomized field experiment to causally identify the heterogeneous effects of social ads and social influence across products and product characteristics. This section summarizes our study's key findings and discusses limitations and implications for research and marketing practice.

6.1. Key Findings

The collective evidence, reported across multiple specifications and operationalizations of key variables, points to four broad patterns in our results. First, we find significant heterogeneity in social advertising effectiveness across 71 products in 25 categories. The highest product-level social ad effectiveness is 3.06 times as large as the lowest positive product-level social ad effectiveness. Food, clothes, and cars are the three best-performing categories. Displaying a friend's like on an ad for food is 1.64 times more effective than doing so for mobile games, 1.57 times more effective than doing so for electrical appliances, and 1.55 times more effective than doing so for financial services.

Second, several theories explain why some products perform better than others. Theories of status production online and the role of status in consumption decisions help explain why status goods perform better than nonstatus goods in the context of social advertising. Our results show that social advertising is 14.82% more effective for status goods than for nonstatus goods. We look to our peers' opinions especially when status is a factor of consumption. This may be because status symbols enable ad viewers to display their status to others and identify themselves with friends by engaging with products their friends endorsed. Greater social risks associated with status goods may also increase the influence of friends and social ad effectiveness. On the other hand, social ads are not significantly more effective for products with greater quality uncertainty, such as experience goods, than their theoretical counterparts, search goods. Our results suggest that status symbols in a product may be a more important factor than quality uncertainty for the effectiveness of social ads.

Third, we find the relative characteristics of ad viewers and the friends shown in ads significantly moderate social advertising effectiveness across products. Relative status significantly moderated social advertising effectiveness for status goods, nonstatus goods, and experience goods. Friends exerted 4.67% more influence for status goods, 2.64% more influence for nonstatus goods, and 4.91% more influence for experience goods on ad engagement when the difference in status between the friend and the viewer increased by one standard deviation. Showing friends with greater social status in ads led to larger social influence and significantly increased social ad effectiveness. Relative product involvement also significantly affected social advertising effectiveness for all types of goods. Friends exerted 4.14% more influence for status goods, 2.75% more influence for nonstatus goods, 3.72% more influence for experience goods, and 2.06% more influence for search goods on ad engagement when the difference in product involvement between the friend and the viewer increased by one standard deviation. These results shed light on the conditions under which informational social influence may operate, namely, that when the peer shown in the ad is more involved with the product than the ad viewer, informational social influence is more pronounced, and social ads are more effective. Although we observe that the influence of friends of higher status is larger for status goods than nonstatus goods (4.67% > 2.64%) and the impact of friends with more product involvement is greater for experience goods than search goods (3.72% > 2.06%), these differences are not statistically significant.

We also identified how dyadic relationships between friends and the ad viewers moderate the effects of social ads across product types. When ad viewers and the friends shown in ads were from the same city or of different genders, social cues exerted significantly larger effects for status goods and experience goods, but not for nonstatus goods and search goods. Tie strength between the friends and the viewers significantly impacted the effects of social ads for all types of products. The moderating effects of tie strength in social ad effectiveness are significantly greater for status goods than for nonstatus goods.

Fourth, we provide some of the first experimental evidence of social ad effectiveness comparing ads *without* any social cue to ads with one or multiple social cues. Displaying a social cue (a friend's like) made users 33.75% more likely to click an ad on average ($p < 0.01$), and caused up to a 270% increase in the click-through rate for a social advertisement. This indicates that social ads outperform ads without social cues by a significant margin.

6.2. Limitations

First, although we can evaluate how social ad effectiveness varies across the characteristics of the friends that are shown in the ads, investigations of the causal effect of friends' characteristics should ideally experimentally vary which friends are shown in the ads to estimate the causal effect of showing certain friends with certain characteristics. We hope that future research will address this important challenge and move research forward in the area of personalized social advertising. In thinking about such an investigation, we urge our colleagues to seek settings in which there is statistical support for the distributions of characteristics being evaluated to avoid selection bias in the types of friends that can be randomized. For example, if one wants to estimate the effect of the education level of the friend shown in the ad, the ad viewers considered in the experiment must have friends in all strata of education for an ego-network-level randomization to estimate such effects without bias. Absent such a setting, this design will not be able to distinguish true social influence effects from selection effects driven by endogenous link formation due to, for example, homophily. It is also worth noting, as with most applied interventions, that people (in this case WeChat users) may acclimate to behavioral nudges, making them less effective over time. We encourage investigations of such attenuation effects in personalized social advertising programs.

Second, as with all studies in this area, the platform we studied skews heavily toward a particular community, in our case, Chinese users, and a specific product, early-stage WeChat Moments ads. Although most studies of social influence online have this limitation, we feel it is important to circumscribe our conclusions accordingly. Although the early-stage WeChat Moments ads enabled us to control our experiment effectively, it is also possible that the novelty effects associated with this early-stage product may have led to larger effect sizes in baseline click-through rates. Users may have been more curious about Moments ads and friends' ad endorsements when they were first introduced. That said, to the best of our knowledge, no existing theory or prior empirical work suggests that this novelty will bias our estimates of the heterogeneous effects of social ads, which is the focus of our study. Although our results generalize to Chinese consumers and may generalize more broadly to all consumers, we must bound our generalizations to the community we studied. It could be, for example, that status processes and reactions to advertising vary across cultures (Van den Bulte and Stremersch 2004). Although status may be an important part of social influence processes in Chinese advertising, it may or may not be so in the United States or Europe. Further work is needed to understand whether such cross-cultural variation exists in social advertising.

Third, we studied ad clicking, which represents customer engagement with an ad. Although theories and recent empirical evidence suggest that ad click-throughs increase conversions (e.g., Jones et al. 2017), click-through rates are only a proxy metric for ad effectiveness. Future research that identifies social advertising effectiveness relying on lower funnel advertising outcomes, such as online and offline purchases, may further advance the literature on social advertising. It is also possible that individuals feel negatively about ads with zero endorsement, because no friends "like" them. This negative feeling may be especially true for unknown products. Although we controlled for the effects of big brands, future research that explores how the potential negative effects of zero endorsement vary across products may provide a deeper understanding of social advertising effectiveness. Because of our agreement with WeChat, we can only report the effect sizes using relative risks rather than absolute risks, which obscures absolute lift and return on investment (ROI). Future work on how social ads impact ROI may provide additional insights on social advertising effectiveness.

Finally, future research could explore other theoretically motivated dimensions of the phenomenon. Other characteristics of products, experiences, and behaviors may also moderate social influence. For example, products with strong network externalities may drive social influence because consumers will obtain greater value from peers using the same product. The uniqueness and novelty of products may also arouse greater curiosity about the products that friends liked and therefore increase social influence, especially among curious individuals. For example, ad viewers may be curious about a product whose endorsers' social status mismatches what the product reflects, and status goods can be subject to greater curiosity effects than nonstatus goods. Many other individual and relative characteristics of ad viewers and friends shown in ads may also be relevant. More empirical and theoretical work is needed to produce a more complete theory of social influence and behavioral contagion. We hope our work will move us one step closer to these goals.

6.3. Implications for Research and Practice

Our results directly support the use of social advertising to promote various kinds of products on digital platforms. We causally identified the effects of social ads for 71 products. This provides valuable evidence that marketers can use when they invest in social ads and platform managers can use to price and promote social advertisements. The striking heterogeneity of social advertising effectiveness that we found suggests that there are significant opportunities for brands to improve their returns from social advertising. Simply allocating social advertising dollars to the better-performing categories that we identify will generate higher returns. Furthermore, there are additional opportunities for improving the performance of other categories or making nuanced budget allocations that consider the cost of social advertising across categories as well as its performance. The results also can help platform managers optimally price social ads for different products by considering the heterogeneity in the returns to personalized social advertising strategies across products.

Our study paves the way toward more sophisticated personalized social advertising programs, in which marketers could potentially maximize returns by deciding who, specifically, to show in advertisements, to whom, and for which products. Strategically displaying social cues in ads can significantly increase their effectiveness, and this effect varies predictably across products. We document that well-connected individuals are influential, and targeting those with greater network degree could well improve the effectiveness of social ads. The costs of collecting the data of users' network degree (the number of friends) and strategically showing these high-degree friends on ads are usually not high for social advertising platforms. Our study also suggests that individuals who have historically been more involved with product-related content can exert greater influence on ad viewers. This result indicates that even *indirect* involvement with products, such as reading product-related content, can be a significant indicator of influence. Our findings are particularly valuable for identifying and targeting influential individuals to maximize social ad effectiveness, because social advertising platforms, such as WeChat and Facebook, may lack the data on direct product adoption and usage. Furthermore, we found that displaying friends of higher social status, greater product involvement, and stronger ties in ads will increase social ad effectiveness for almost all types of products. Previous research mostly focuses on one product or average effects across products. Our results are therefore more applicable for practitioners to design personalized social advertising strategies for a particular product.

Our findings not only shed light on the heterogeneity of social advertising effectiveness, but also have broad implications for how marketers should promote product diffusion in social networks and how they can differentiate their strategies in network marketing campaigns. For instance, we distinguish status goods, non-status goods, experience goods, and search goods from each other and find that status goods will gain more from influence-based marketing campaigns, whereas there appears to be surprisingly little evidence of a difference between experience goods and search goods. Benefits can be attained by initially targeting users who are likely to influence their friends in local social networks, such as those of higher social status and greater product involvement, and then displaying these user-endorsed ads to peers, especially peers who are susceptible to influence, such as those connected by stronger ties to product adopters.

Finally, we hope the experimental, across-product empirical evidence in our analyses will help researchers build more complete and generalizable theories of social influence and social contagion in product adoption and aid marketers in optimizing their social advertising and network marketing programs in general.

Acknowledgments

The authors gratefully acknowledge generous support from the WeChat team, who implemented this large-scale field experiment for our research. A special acknowledgment goes to Xiaolong Zhang, Bin Zhang, Haojun Wu, Yuxuan Zhang, and Zhiyuan Lin. For valuable comments, the authors are also grateful to Dean Eckles, Catherine Tucker, Ray Reagans, Juanjuan Zhang, Ezra Zuckerman, Song Lin, Michael Zhao, and seminar participants at the MIT Initiative on the Digital Economy (IDE), the Conference on Digital Experimentation (CODE), the Workshop on Information Systems Economics (WISE), the International Conference on Information Systems (ICIS), Massachusetts Institute of Technology, University of Pennsylvania, University of Washington, Northwestern University, and Carnegie Mellon University.

Appendix A. Definitions of Product Types

Product types	Description
Experience\search goods	Experience goods must be experienced to be truly evaluated. Search goods can be evaluated with only published information and do not necessarily need to be experienced to be evaluated. Examples of experience goods include clothes, food, and video games. Examples of search goods include laptops, cell phones, and credit card services.
Status\nonstatus goods	Consumers have the tendency to purchase goods and services for gaining and displaying social status or prestige. A consumer may seek to purchase or consume goods and services, which exhibit or serve as status symbols, for the status they confer, regardless of the consumer's objective income or social class level. Examples of status goods include status-conferring clothing, cars, wines, restaurants, and hotels. Examples of nonstatus goods include toothpaste, beverages, and website services.

Appendix B. Robustness Checks

Table B.1. Effects of Users' Status and Product Involvement on Social Advertising Effectiveness Across Products, Controlling for Dyadic Similarities

	1	2	3	4	5	6
	Status good clicking	Nonstatus good clicking	Experience good clicking	Search good clicking	Status vs. nonstatus clicking	Experience vs. search clicking
SC × FRS	0.0450***	0.0259**	0.0478***	−0.0022	0.0259**	−0.0022
	(0.0169)	(0.0104)	(0.0109)	(0.0146)	(0.0103)	(0.0144)
SC × FRI	0.0396***	0.0266***	0.0367***	0.0213**	0.0266***	0.0213**
	(0.0113)	(0.0082)	(0.0086)	(0.0085)	(0.0082)	(0.0083)
SC × AGESIMI	−0.0091	0.0025	−0.0060	−0.0072	0.0025	−0.0072
	(0.0245)	(0.0234)	(0.0227)	(0.0192)	(0.0234)	(0.0189)
SC × GENDERSIMI	−0.0581**	−0.0202	−0.0522***	−0.0272	−0.0202	−0.0272
	(0.0242)	(0.0149)	(0.0180)	(0.0178)	(0.0149)	(0.0175)
SC × CITYSIMI	0.0842**	0.0195	0.0682**	−0.0072	0.0195	−0.0072
	(0.0388)	(0.0220)	(0.0266)	(0.0302)	(0.0219)	(0.0298)
SC × FRS × SG					0.0191	
					(0.0196)	
SC × FRI × SG					0.0130	
					(0.0138)	
SC × AGESIMI × SG					−0.0116	
					(0.0336)	
SC × GENDERSIMI × SG					−0.0379	
					(0.0281)	
SC × CITYSIMI × SG					0.0647	
					(0.0441)	
SC × FRS × EG						0.0500***
						(0.0181)
SC × FRI × EG						0.0154
						(0.0120)
SC × AGESIMI × EG						0.0012
						(0.0295)
SC × GENDERSIMI × EG						−0.0250
						(0.0251)
SC × CITYSIMI × EG						0.0754*
						(0.0399)
Log-likelihood	−250,034	−227,862	−368,999	−109,019	−477,934	−478,020
Observations	2,387,250	1,346,773	3,215,964	518,059	3,734,023	3,734,023

Notes. SC, Social cue; FRS, friend's relative status; FRI, friend's relative involvement; SG, status good; EG, experience good; AGESIMI, dyadic similarity in age; GENDERSIMI, dyadic similarity in gender; CITYSIMI, dyadic similarity in city. An observation is a user–ad pair. Robust standard errors, clustered at the ad level, are reported in parentheses.

*$p < 0.10$; **$p < 0.05$; ***$p < 0.01$.

Table B.2. Effects of Users' Status and Product Involvement on Social Advertising Effectiveness Across Products, Controlling for Tie Strength

	1	2	3	4	5	6
	Status good clicking	Nonstatus good clicking	Experience good clicking	Search good clicking	Status vs. nonstatus clicking	Experience vs. search clicking
SC × FRS	0.0444***	0.0258**	0.0473***	−0.0023	0.0258**	−0.0023
	(0.0168)	(0.0104)	(0.0109)	(0.0145)	(0.0104)	(0.0143)
SC × FRI	0.0417***	0.0281***	0.0387***	0.0226***	0.0281***	0.0226***
	(0.0113)	(0.0084)	(0.0087)	(0.0086)	(0.0084)	(0.0085)
SC × TS	0.2544***	0.1452***	0.2156***	0.1478***	0.1452***	0.1478***
	(0.0279)	(0.0223)	(0.0271)	(0.0350)	(0.0222)	(0.0346)
SC × FRS × SG					0.0186	
					(0.0195)	
SC × FRI × SG					0.0136	
					(0.0139)	
SC × TS × SG					0.1092***	
					(0.0353)	
SC × FRS × EG						0.0496***
						(0.018)
SC × FRI × EG						0.0161
						(0.0122)
SC × TS × EG						0.0678
						(0.0439)
Log-likelihood	−249,888	−227,817	−368,856	−108,996	−477,744	−477,854
Observations	2,387,250	1,346,773	3,215,964	518,059	3,734,023	3,734,023

Notes. SC, Social cue; FRS, friend's relative status; FRI, friend's relative involvement; SG, status good; EG, experience good; TS, tie strength. An observation is a user–ad pair. Robust standard errors, clustered at the ad level, are reported in parentheses.
$^{**}p < 0.05$; $^{***}p < 0.01$.

Table B.3. Social Advertising Effectiveness Across Search/Experience Goods and Status/Nonstatus Goods with Multiple Impressions

	1	2
	Clicking	Clicking
Social cue	0.1794***	0.1342**
	(0.0183)	(0.0667)
SC × experience goods	0.0387	0.0549
	(0.0537)	(0.0596)
SC × status goods	0.1550**	0.1104*
	(0.0640)	(0.0594)
Controls		Yes
Log-likelihood	−777,493	−701,778
Observations	3,734,023	3,734,023

Notes. SC, Social cue. An observation is a user–ad pair. Robust standard errors, clustered at the ad level, are reported in parentheses.
$^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table B.4. Effects of Users' Status and Product Involvement on Social Advertising Effectiveness Across Products with Multiple Impressions

	1	2	3	4	5	6
	Status good clicking	Nonstatus good clicking	Experience good clicking	Search good clicking	Status vs. nonstatus clicking	Experience vs. search clicking
SC × FRS	0.0404***	0.0234***	0.0408***	0.0041	0.0234***	0.0041
	(0.0149)	(0.0072)	(0.0099)	(0.0115)	(0.0072)	(0.0113)
SC × FRI	0.0454***	0.0313***	0.0431***	0.0213***	0.0313***	0.0213***
	(0.0084)	(0.0085)	(0.0077)	(0.0072)	(0.0085)	(0.0071)
SC × FRS × SG					0.0170	
					(0.0163)	
SC × FRI × SG					0.0141	
					(0.0119)	
SC × FRS × EG						0.0367**
						(0.0150)
SC × FRI × EG						0.0218**
						(0.0104)
Log-likelihood	−381,426	−317,967	−549,674	−149,926	−699,462	−699,603
Observations	2,387,250	1,346,773	3,215,964	518,059	3,734,023	3,734,023

Notes. SC, Social cue; FRS, friend's relative status; FRI, friend's relative involvement; SG, status good; EG, experience good. An observation is a user–ad pair. Robust standard errors, clustered at the ad level, are reported in parentheses.
**$p < 0.05$; ***$p < 0.01$.

Table B.5. Social Advertising Effectiveness Across Search/Experience Goods and Status/Nonstatus Goods in the Control Group vs. Treatment Group 2

	1	2
	Clicking	Clicking
Social cue	0.1979***	0.2104***
	(0.0366)	(0.0772)
SC × experience goods	0.0537	0.0750
	(0.0765)	(0.0800)
SC × status goods	0.2308***	0.1858**
	(0.0843)	(0.0725)
Controls		Yes
Log-likelihood	−540,521	−486,148
Observations	3,697,715	3,697,715

Notes. SC, Social cue. An observation is a user–ad pair. Robust standard errors, clustered at the ad level, are reported in parentheses.
**$p < 0.05$; ***$p < 0.01$.

Table B.6. Effects of Users' Status and Product Involvement on Social Advertising Effectiveness Across Products in the Control Group vs. Treatment Group 2

	1	2	3	4	5	6
	Status good clicking	Nonstatus good clicking	Experience good clicking	Search good clicking	Status vs. nonstatus clicking	Experience vs. search clicking
SC × FRS	0.0577***	0.0244***	0.0492***	0.0178**	0.0244***	0.0178**
	(0.0138)	(0.0069)	(0.0097)	(0.0087)	(0.0068)	(0.0086)
SC × FRI	0.0520***	0.0270***	0.0456***	0.0213	0.0270***	0.0213
	(0.0134)	(0.0097)	(0.0102)	(0.0147)	(0.0096)	(0.0145)

Table B.6. (Continued)

	1	2	3	4	5	6
	Status good clicking	Nonstatus good clicking	Experience good clicking	Search good clicking	Status vs. nonstatus clicking	Experience vs. search clicking
SC × FRS × SG					0.0333**	
					(0.0152)	
SC × FRI × SG					0.0250	
					(0.0163)	
SC × FRS × EG						0.0314**
						(0.0129)
SC × FRI × EG						0.0243
						(0.0177)
Log-likelihood	−255,625	−228,961	−375,429	−109,376	−484,606	−484,810
Observations	2,358,227	1,339,488	3,182,600	515,115	3,697,715	3,697,715

Notes. SC, Social cue; FRS, friend's relative status; FRI, friend's relative involvement; SG, status good; EG, experience good. An observation is a user–ad pair. Robust standard errors, clustered at the ad level, are reported in parentheses.
$^{**}p < 0.05$; $^{***}p < 0.01$.

Table B.7. Social Advertising Effectiveness Across Search/Experience Goods and Status/Nonstatus Goods in the Control Group vs. Treatment Group 1 Among Users Exposed to One Ad

	1	2
	Clicking	Clicking
Social cue	0.2086***	0.1619*
	(0.0321)	(0.0875)
Social cue	0.1963***	0.1492*
	(0.0293)	(0.0812)
SC × experience goods	0.0530	0.0905
	(0.0722)	(0.0786)
SC × status goods	0.1955**	0.1456*
	(0.0814)	(0.0761)
Controls		Yes
Log-likelihood	−435,171	−391,686
Observations	3,113,824	3,113,824

Notes. SC, Social cue. An observation is a user–ad pair. Robust standard errors, clustered at the ad level, are reported in parentheses.
$^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table B.8. Effects of Users' Status and Product Involvement on Social Advertising Effectiveness Across Products in the Control Group vs. Treatment Group 1 Among Users Exposed to One Ad

	1	2	3	4	5	6
	Status good clicking	Nonstatus good clicking	Experience good clicking	Search good clicking	Status vs. nonstatus clicking	Experience vs. search clicking
SC × FRS	0.0403**	0.0270***	0.0446***	−0.0040	0.0270***	−0.0040
	(0.0176)	(0.0092)	(0.0112)	(0.0128)	(0.0092)	(0.0127)
SC × FRI	0.0404***	0.0328***	0.0381***	0.0254**	0.0328***	0.0254**
	(0.0140)	(0.0111)	(0.0104)	(0.0109)	(0.0111)	(0.0107)

Table B.8. (Continued)

	1	2	3	4	5	6
	Status good clicking	Nonstatus good clicking	Experience good clicking	Search good clicking	Status vs. nonstatus clicking	Experience vs. search clicking
SC × FRS × SG					0.0132	
					(0.0197)	
SC × FRI × SG					0.0076	
					(0.0177)	
SC × FRS × EG						0.0486***
						(0.0169)
SC × FRI × EG						0.0127
						(0.0149)
Log-likelihood	−209,774	−180,465	−307,430	−82,957	−390,240	−390,388
Observations	1,999,136	1,114,688	2,709,267	404,557	3,113,824	3,113,824

Notes. SC, Social cue; FRS, friend's relative status; FRI, friend's relative involvement; SG, status good; EG, experience good. An observation is a user–ad pair. Robust standard errors, clustered at the ad level, are reported in parentheses.
$^{**}p < 0.05$; $^{***}p < 0.01$.

Table B.9. Social Advertising Effectiveness Across Search/Experience Goods and Status/Nonstatus Goods, Controlling for Targeting Conditions

	1	2
	Clicking	Clicking
Social cue	0.1907***	0.4560***
	(0.0240)	(0.0485)
SC × Experience goods	0.0339	0.0692
	(0.0606)	(0.0485)
SC × Status goods	0.1807**	0.1522***
	(0.0721)	(0.0572)
Controls		Yes
Log-likelihood	−534,820	−514,860
Observations	3,734,023	3,734,023

Notes. SC, Social cue. An observation is a user–ad pair. Robust standard errors, clustered at the ad level, are reported in parentheses.
$^{**}p < 0.05$; $^{***}p < 0.01$.

Endnotes

[1] Conversions include purchases, sign-ups, or store location lookups.

[2] Unlike Facebook ads, users cannot directly share Moments ads with their friends on WeChat.

[3] WeChat shows the names of all the friend endorsers under the ad images. If there are many names to show, the name list expands to multiple lines. Facebook shows more kinds of social cues on its social ads than WeChat does. For example, the social cues could be the names of endorsers, the count of endorsers, or the combination of the names and the count (e.g., David, Tom, and 285 others). The endorsers shown or counted by Facebook ads are not limited to ad viewers' first-degree friends.

[4] WeChat Moments ads were introduced in early 2015. Our experiment was conducted in December of 2015.

[5] A "like" button was available on the ads, and ad viewers could still "like" the ads, even in the control group.

[6] We did not randomly select friends but randomly assigned the presence and number of friends' likes shown in ads. If multiple friends' likes were presented, they were shown chronologically.

[7] WeChat is known as China's "app for everything" or a "super app" (https://en.wikipedia.org/wiki/WeChat). Users, on average, spend more than an hour each time they use WeChat. This makes WeChat one the most important information channels for people in China. We therefore believe our data proxy well for the amount and type of information that users acquire on a daily basis.

[8] We did not estimate hazard models for the analysis because adoption time does not implicate the degree of social influence in our context. Adoption times depend more on users' WeChat use habits, like their level of engagement on the platform. As we study peer effects on users' first reactions to ads, the variation of (relative) adoption time is small and less meaningful.

[9] We use age dummies to indicate whether users' ages are in [0,20], (20,25], (25,30], (30,35], (35,40], (40,45], (45,50], or (50,+∞), and city dummies to indicate whether users' most common log-in cities are first class, second class, or other.

[10] We dropped 10 old ads that were left over from the preexperimental period, and another 7 ads whose click-through rate in the control group was 0. Users were already exposed to the old ads before the experiment started, and the sample sizes for the 17 dropped ads were very small.

[11] We dropped a small amount of data (0.16% of the sample) because of the technical errors that caused an incorrect number of likes to be displayed on some user–ad pairs during the experiment.

For example, in our analysis, we excluded the user–ad pairs assigned to Treatment Group 1 with more than one like or no likes (if the number of organic likes is positive) shown in ads, and the user–ad pairs assigned to Treatment Group 2 in which the number of displayed likes did not match the number of organic likes. Specifically, 77 user–ad pairs were dropped in the control group, 26,183 were dropped in Treatment Group 1 and 66,852 were dropped in Treatment Group 2. The resulting imbalance is not economically meaningful; see Table 3 (Ding and VanderWeele 2016).

[12] We measured the involvement gap across multiple friends as the maximum product involvement gap between the ad viewer and any of the friends used in the social cues and the status of the group of friends with organic likes as the average network centrality of these friends. We then checked the robustness of these results by measuring the involvement gap by the average involvement gap between the ad viewer and any of the friends used in the social cues and the status gap as the maximum network centrality of the group of friends used in the social cues. The results are consistent across all operationalizations and are available upon request.

[13] We dropped 16.61% of the user–ad pairs associated with the 8.64% of users.

[14] An example of a typical targeting condition could be that users' age ranges from 15 to 30, their most common log-in cities are first or second class, and gender is female or male. We replaced the variables for age, gender, and city with the dummies of 44 targeting conditions in the model.

References

Aral S (2011) Commentary-identifying social influence: A comment on opinion leadership and social contagion in new product diffusion. *Marketing Sci.* 30(2):217–223.

Aral S, Walker D (2011) Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Sci.* 57(9):1623–1639.

Aral S, Walker D (2012) Identifying influential and susceptible members of social networks. *Science* 337(July):337–341.

Aral S, Walker D (2014) Tie strength, embeddedness, and social influence: A large-scale networked experiment. *Management Sci.* 60(6):1352–1370.

Aral S, Muchnik L, Sundararajan A (2009) Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proc. Natl. Acad. Sci. USA* 106(51):21544–21549.

Bagwell LS, Bernheim BD (1996) Veblen effects in a theory of conspicuous consumption. *Amer. Econom. Rev.* 86(3):349–373.

Bakshy E, Eckles D (2013) Uncertainty in online experiments with dependent data: An evaluation of bootstrap methods. *Proc. 19th ACM SIGKDD Internat. Conf. Knowledge Discovery Data Mining* (Association for Computing Machinery, New York), 1303–1311.

Bakshy E, Karrer B, Adamic LA (2009) Social influence and the diffusion of user created content. *Proc. 10th ACM Conf. Electronic Commerce* (Association for Computing Machinery, New York), 325–334.

Bakshy E, Eckles D, Yan R, Rosenn I (2012) Social influence in social advertising: Evidence from field experiments. *Proc. 13th ACM Conf. Electronic Commerce*, vol. 1 (Association for Computing Machinery, New York), 146–161.

Bapna R, Umyarov A (2015) Do your online friends make you pay? A randomized field experiment on peer influence in online social networks. *Management Sci.* 61(8):1902–1920.

Bart Y, Stephen AT, Sarvary M (2014) Which products are best suited to mobile advertising? A field study of mobile display advertising effects on consumer attitudes and intentions. *J. Marketing Res.* 51(3):270–285.

Bearden WO, Etzel MJ (1982) Reference group influence on product and brand purchase decisions. *J. Consumer Res.* 9(2): 183–194.

Berger J (2014) Word of mouth and interpersonal communication: A review and directions for future research. *J. Consumer. Psych.* 24(4):586–607.

Berger J, Milkman KL (2012) What makes online content viral? *J. Marketing Res.* 49(2):192–205.

Berger J, Schwartz EM (2011) What drives immediate and ongoing word of mouth? *J. Marketing Res.* 48(5):869–880.

Bernheim BD (1994) A theory of conformity. *J. Political. Econom.* 102(5):841–877.

Bond RM, Settle JE, Fariss CJ, Jones JJ, Fowler JH (2017) Social endorsement cues and political participation. *Political Comm.* 34(2): 261–281.

Bond RM, Fariss CJ, Jones JJ, Kramer ADI, Marlow C, Settle JE, Fowler JH (2012) A 61-million-person experiment in social influence and political mobilization. *Nature* 489(7415):295–298.

Bramoullé Y, Djebbari H, Fortin B (2009) Identification of peer effects through social networks. *J. Econometrics* 150(1):41–55.

Brynjolfsson E, Hu Y, Simester D (2011) Goodbye Pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Sci.* 57(8):1373–1386.

Burnkrant RE, Cousineau A (1975) Informational and normative social influence in buyer behavior. *J. Consumer Res.* 2(3):206–215.

Cai H, Chen Y (2009) Observational learning: Evidence from a randomized natural field experiment. *Amer. Econom. Rev.* 99(3):864–882.

Cameron RC, Miller DL (2015) A practitioner's guide to cluster-robust inference. *J. Human Resources* 50(2):317–372.

Chaiken S (1980) Heuristic vs. systematic information processing and the use of source vs. message cues in persuasion. *J. Personality Soc. Psych.* 39(5):752–766.

Chaiken S (1987) The heuristic model of persuasion. Zanna MP, Olson JM, Herman CP, eds. *Social Influence: The Ontario Symposium*, vol. 5 (Lawrence Erlbaum Associates, Mahwah, NJ), 3–39.

Christakis N, Fowler J (2013) Social contagion theory: Examining dynamic social networks and human behavior. *Statist. Medicine* 32(4):1–32.

Cialdini RB, Goldstein NJ (2004) Social influence: Compliance and conformity. *Annual Rev. Psych.* 55:591–621.

Colicev A, Malshe A, Pauwels K, O'Connor P (2017) Improving consumer mind-set metrics and shareholder value through social media: The different roles of owned and earned. *J. Marketing* 82(1):37–56.

Corneo G, Jeanne O (1997) Conspicuous consumption, snobbism and conformism. *J. Public Econom.* 66(1):55–71.

Deutsch M, Gerard HB (1955) A study of normative and informational social influences upon individual judgment. *J. Abnormal Psych.* 51(3):629–636.

Ding P, VanderWeele TJ (2016) Sensitivity analysis without assumptions. *Epidemiology* 27(3):368–377.

Friedman HH, Friedman L (1979) Endorser effectiveness by product type. *J. Advertising Res.* 19(5):63–71.

Godes D, Mayzlin D (2009) Firm-created word-of-mouth communication: Evidence from a field test. *Marketing Sci.* 28(4):721–739.

Goldenberg J, Libai B, Muller E (2001) Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing Lett.* 12(3):211–223.

Goldenberg J, Han S, Lehmann DR, Hong JW (2009) The role of hubs in the adoption process. *J. Marketing* 73(2):1–13.

Hanssens DM, Weitz BA (1980) The effectiveness of industrial print advertisements across product categories. *J. Marketing Res.* 17(3):294–306.

Hass R (1981) Effects of source characteristics on cognitive responses and persuasion. Petty R, Ostrom TM, Brock TC, eds. *Congitive Responses in Persuasion* (Psychology Press, London), 141–172.

Hu Y, Van den Bulte C (2014) Nonmonotonic status effects in new product adoption. *Marketing Sci.* 33(4):509–533.

Huang S (2016) Social influence in public and private behaviors. *Proc. 37th Internat. Conf. Inform. Systems* (Association for Information Systems), 1–16.

Iyengar R, Van den Bulte C, Lee JY (2015) Social contagion in new product trial and repeat. *Marketing Sci.* 34(3):408–429.

Iyengar R, Van den Bulte C, Valente TW (2011) Opinion leadership and social contagion in new product diffusion. *Marketing Sci.* 30(2):195–212.

Jackson MO (2008) *Social and Economic Networks*, vol. 3 (Princeton University Press, Princeton, NJ).

Jacoby J, Kaplan LB (1972) The components of perceived risk. *Proc. 3rd Annual Conf. Assoc. Consumer Res.* (Association for Consumer Research, Chicago), 382–393.

Johnson G, Lewis RA, Nubbemeyer E (2017) The online display ad effectiveness funnel & carryover: Lessons from 432 field experiments. Preprint, submitted October 1, http://dx.doi.org/10.2139/ssrn.2701578.

Jones JJ, Bond RM, Bakshy E, Eckles D, Fowler JH (2017) Social influence and political mobilization: Further evidence from a randomized experiment in the 2012 U.S. presidential election. *PLoS One* 12(4):e0173851.

Kuksov D, Xie Y (2012) Competition in a status goods market. *J. Marketing Res.* 49(5):609–623.

Kulviwat S, Bruner GC, Al-Shuridah O (2009) The role of social influence on adoption of high tech innovations: The moderating effect of public/private consumption. *J. Bus. Res.* 62(7):706–712.

Lakens D, Scheel AM, Isager PM (2018) Equivalence testing for psychological research: A tutorial. *Adv. Methods Practices Psych. Sci.* 1(2):259–269.

Levina N, Arriaga M (2014) Distinction and status production on user-generated content platforms: Using Bourdieu's theory of cultural production to understand social dynamics in online fields. *Inform. Systems Res.* 25(3):468–488.

Li X (2018) Impact of average rating on social media endorsement: The moderating role of rating dispersion and discount threshold. *Inform. Systems Res.* 29(3):739–754.

Lin S, Zhang J, Hauser JR (2015) Learning from experience, simply. *Marketing Sci.* 34(1):1–19.

Lovett MJ, Peres R, Shachar R (2013) On brands and word of mouth. *J. Marketing Res.* 50(4):427–444.

Manski CF (1993) Identification of endogenous social effects: The reflection problem. *Rev. Econom. Stud.* 60(3):531–542.

Muchnik L, Aral S, Taylor SJ (2013) Social influence bias: a randomized experiment. *Science* 341(6146):647–651.

Nelson P (1970) Information and consumer behavior. *J. Political Econom.* 78(2):311–329.

O'Cass A, McEwen H (2004) Exploring consumer status and conspicuous consumption. *J. Consumer Behav.* 4(1):25–39.

Pesendorfer W (1995) Design innovation and fashion cycles. *Amer. Econom. Rev.* 85(4):771–792.

Peters K, Chen Y, Kaplan AM, Ognibeni B, Pauwels K (2013) Social media metrics—A framework and guidelines for managing social media. *J. Interactive Marketing* 27(4):281–298.

Rubin DB (1990) Comment: Neyman (1923) and causal inference in experiments and observational studies. *Statist. Sci.* 5(4):472–480.

Sacerdote B (2001) Peer effects with random assignment: Results for Dartmouth roommates. *Quart. J. Econom.* 116(2):681–704.

Sahni NS (2015) Effect of temporal spacing between advertising exposures: Evidence from online field experiments. *Quant. Marketing Econom.* 13(3):203–247.

Sahni NS (2016) Advertising spillovers: Evidence from online field experiments and implications for returns on advertising. *J. Marketing Res.* 53(4):459–478.

Salganik MJ, Dodds PS, Watts DJ (2006) Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311(5762):854–856.

Schmalensee R (1978) A model of advertising and product quality. *J. Political Econom.* 86(3):485–503.

Stephen AT, Galak J (2012) The effects of traditional and social earned media on sales: A study of a microlending marketplace. *J. Marketing Res.* 49(5):624–639.

Stephen AT, Toubia O (2010) Deriving value from social commerce networks. *J. Marketing Res.* 47(2):215–228.

Taylor S, Bakshy E, Aral S (2013) Selection effects in online sharing: Consequences for peer adoption. *Proc. 14th ACM Conf. Electronic Commerce* (Association for Computing Machinery, New York), 821–836.

Taylor SJ, Bakshy E, Eckles D, Aral S (2015) Heterogeneous social influence from personalized social cues. MIT Sloan Working Paper, Massachusetts Institute of Technology, Cambridge.

Trusov M, Bucklin RE, Pauwels K (2009) Effects of word-of-mouth vs. traditional marketing: Findings from an Internet social networking site. *J. Marketing* 73(5):90–102.

Tucker C (2008) Identifying formal and informal influence in technology adoption with network externalities. *Management Sci.* 54(12):2024–2038.

Tucker CE (2016) Social advertising: How advertising that explicitly promotes social influence can backfire. Preprint, submitted June 1, http://dx.doi.org/10.2139/ssrn.1975897.

Turner JC (1991) *Social Influence* (Thomson Brooks/Cole Publishing, Pacific Grove, CA).

Van Alstyne M, Brynjolfsson E (2005) Global village or cyber-Balkans? Modeling and measuring the integration of electronic communities. *Management Sci.* 51(6):851–868.

Van den Bulte C (2000) New product diffusion acceleration: Measurement and analysis. *Marketing Sci.* 19(4):366–380.

Van den Bulte C, Lilien GL (2001) Medical innovation revisited: Social contagion vs. marketing effort. *Amer. J. Sociol.* 106(5):1409–1435.

Van den Bulte C, Stremersch S (2004) Social contagion and income heterogeneity in new product diffusion: A meta-analytic test. *Marketing Sci.* 23(4):530–544.

Van den Bulte C, Wuyts S (2007) *Social Networks and Marketing* (Marketing Science Institute, Cambridge, MA).

Veblen T (1899) *The Theory of the Leisure Classes: An Economic Study of Institutions* (Macmillan, New York).

Wang Y, Griskevicius V (2014) Conspicuous consumption, relationships, and rivals: Women's luxury products as signals to other women. *J. Consumer Res.* 40(February):834–854.

Xiang R, Neville J, Rogati M (2010) Modeling relationship strength in online social networks. *Proc. 19th Internat. Conf. World Wide Web* (Association for Computing Machinery, New York), 981–990.

Zhang J (2010) The sound of silence: Observational learning in the U.S. kidney market. *Marketing Sci.* 29(2):315–335.

Zhu F, Zhang X (2010) Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics. *J. Marketing* 74(2):133–148.