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

Shan Huang[1], Sinan Aral[2], Jeffrey(Yu) Hu[3], and Erik Brynjolfsson[4]

[1]MIT

[2]MIT

[3]Georgia Tech

[4]MIT

Abstract

Almost all of the empirical evidence of a lift from social advertising focuses on a single product at a time. As a result, we know little about how social advertising effectiveness varies across product categories or how product characteristics impact social advertising effectiveness. We therefore collaborated with WeChat to conduct a randomized field experiment measuring social ad effectiveness across 71 products in 25 categories among a random sample of more than 37 million users of WeChat Moments Ads. We found that some product categories, like food, clothes, and cars, experienced significantly stronger social advertising effectiveness than other categories like financial services and electrical appliances. More generally, we found that status goods, which rely on normative social influence, displayed strong social advertising effectiveness, while social ads for experience goods, which rely on informational social influence, did not perform any better or worse than their theoretical counterpart search goods. The status and expertise of the user displayed in the ad also moderated these effects differently across different products. Understanding the heterogeneous effects of social advertising across products will help marketers differentiate their social advertising strategies and lead researchers to a more general theory of social influence in product adoption.

Contact. shanh@mit.edu.

1. INTRODUCTION

Spending on social advertising is increasing dramatically. The explosive growth of Facebook's ad revenue alone highlights the high expectations advertisers place on social advertising. 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, Iyengar et al. 2011, Aral 2011, Berger and Milkman 2012).

But, while recent work has demonstrated that social ads achieve significant lift from the social proof in peer endorsements (e.g. Bakshy et al. 2012, Aral and Walker 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. Although 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), no 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 (Aral 2011). As a result, we know little about how social influence and social advertising effectiveness vary across product categories.[1]

The goal of this research is to identify the heterogeneous effects of social advertising across products and to investigate how product characteristics impact social advertising effectiveness. Are social ads more effective for electronic products or fashion accessories? Are we more likely to be swayed by the opinions of our friends when shopping for status goods, like a luxury car, or when we are seeking trusted information about a product? We simply don't know the answers to these

[1]Social advertising is a broad term and it's applications range widely, from social 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 foundational and 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 efforts place the images and names of people who have liked a brand in ads for that brand shown to those people's friends. Google's 'friendorsements' do the same thing in search results for products, placing the names, images and product ratings of friends in product search results. In this paper, we define social advertising as the placement of social cues or endorsements in ads shown to friends of those who have engaged with a brand or product. These social ads rely on the power of social influence in product adoption and the value of social cues in social media engagement to encourage lift in ad effectiveness.

questions and it is difficult to generalize a theory of social influence in product adoption from one product to the next while parameter estimates of influence in consumer adoption decisions remain highly idiosyncratic (Friedman and Friedman 1979, Bearden and J.Etzel 1982, Kulviwat et al. 2009, Stephen and Galak 2012).

Although social influence is of central importance in marketing and the social science more broadly (e.g., Deutsch and Gerard 1955, Burnkrant and Cousineau 1975, Sacerdote 2001, Cialdini and Goldstein 2004, 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 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 influence across 71 products in 25 product categories and to examine the effects of product characteristics on the effectiveness of social ads (Aral 2016). The experiment was conducted on a random sample of more than 37 million users of WeChat Moments Ads, a type of social advertisement displayed in WeChat users' newsfeeds. WeChat is a world leading mobile social networking platform with a billion monthly active users. Our experiment involves user-ad level randomization of the number of social cues shown on WeChat Moments Ads. Social influence in our experiment is measured as the degree to which social cues (i.e. likes), representing friends' endorsements of products, affect users' responses to social advertising (i.e. click-throughs). By randomly assigning the number of social cues displayed on (otherwise identical) ads in a real-world context and controlling for many other factors, we were able to obtain unbiased estimates of the impact of social influence on ad engagement across many different products simultaneously.

Our study moves the research frontier from identifying the presence of social influence and its heterogeneous effects across individuals and relationships to investigating how product characteristics moderate social influence. Our results, which provide the first large-scale experimental evidence of the heterogeneous effects of social advertising across products, demonstrate that not all products are created equal in terms of their diffusion through social networks. We found that some product categories, like food, clothing, and cars, experienced significantly stronger social advertising effectiveness than other categories like financial services and electrical appliances. More generally, we found that status goods, which rely on normative social influence, displayed strong social advertising effectiveness, but social ads for experience goods, which rely on informational social influence, did not perform any better than their theoretical counterpart search goods. The status and expertise of the user displayed in the ad moderated these effects differently across different products as well. Understanding the heterogeneous effects of social influence across products will help marketers better target their social advertising strategies and advance research toward a more general theory of social influence.

2. THEORY

2.1. Social Advertising Effectiveness Across Products

Multiple theories inform our understanding of when social spillovers are more or less likely. Several of these theories make clear predictions about when social influence will be salient for consumer decisions. For example, social spillovers 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 information in peers' product adoption decisions to infer their value (Lin et al. 2015, Zhang and Liu 2012). On the other hand, Keynes argued that human needs fall into two categories: "those needs which are absolute in the sense that we feel them whatever the situation of our fellow human beings... and those which are relative in the sense that we feel them only if their satisfaction lifts us above, makes us feel superior to, our fellows"(Keynes 2010, p.3). In this way, consumption and status are likely related for some products but not for others (Veblen 2009, Bernheim 1994, O'Cass and McEwen 2004). We therefore examined and compared the effects of social influence and social advertising for experience (or search) goods and status (or non-status) goods.

People have different motives for changing their attitudes or behaviors and social influence can be informational or normative (Deutsch and Gerard 1955, Bearden et al. 1989, Turner 1991, Iyengar et al. 2011). Informational social influence exists when one accepts information obtained from another as evidence about reality, where peers are the mediators of facts. Normative social influence, on the other hand, exists when individuals desire to conform to the expectations of another person or group. Informational and normative social influence operate through three distinct processes, which are distinguished by their motivational antecedents (Kelman 1958, Burnkrant and Cousineau 1975, Bearden and J.Etzel 1982, Turner 1991, Iyengar et al. 2015).

Informational social influence mainly works through the process of internalization, by which individuals desire to make informed decisions and perceive information offered by peers as instrumental to attaining their goals. Normative social influence operates through the processes of identification and compliance. Compliance is driven by individuals' desire to conform to another person or group or to realize a reward or avoid a punishment mediated by that person or group. Rewards or punishments are not only material but also status driven. People are more likely to comply with others' behavior to realize rewards or avoid punishment if they believe their behaviors are visible or known to those others. Therefore, compliance is likely only if adoption decisions or consumption patterns are visible and identifiable by peers. Identification occurs when the behavior is associated with a satisfying self-defining relationship to the other and is beneficial to an individual's sense of self. Identification can take two forms: attempts to resemble others, or attempts to attach to those one likes or agrees with. Product characteristics can affect informational and normative social influence by impacting the three distinct operational processes: internalization, compliance, and identification (See Table 1).

Table 1: Product Types and Social Influence

Product Types	Affected Social Influence	Affected Processes
Experience/Search Goods	Informational	Internalization
Status/Non-Status Goods	Normative	Identification and Compliance

2.2. Experience versus 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 should impact social influence mainly by affecting the processes of internalization and informational influence, which are reflected in customers' desire to utilize the information inferred from others' preferences or behaviors to characterize a product. In theory, informational social influence should affect experience goods more than search goods through the process of internalization. The quality of search goods can be evaluated before purchase after a (costly) search, while the quality of experience goods can only be evaluated by experiencing them or being exposed to the experiences of others. Friends' endorsements, reflecting their personal experience with products, provide additional information for ad viewers when deciding whether to engage with a social advertisement. Faced with a lack of product information for experience goods, individuals will tend to rely more on the experience of their trusted peers to evaluate experience goods than search goods, which are easier to evaluate with simple information about the product found online. Experience goods are therefore more likely to be susceptible to social influence than search goods, which are more easily evaluated using non-social information about the product's characteristics. As a result, we hypothesize that:

Hypothesis 1: *Friends' endorsements, in the form of social cues, have a greater effect on ad engagement for experience goods than for search goods.*

2.3. Status versus Non-Status Goods

While the needs that some products fulfill are individually experienced, the needs that others fulfill are relative. We may therefore expect that 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. Social status has been theorized as a powerful driver of consumption choices (Veblen 2009, Bagwell and Bernheim 1996, Corneo and Jeanne 1997, Wang and Griskevicius 2014). O'Cass and McEwen (2004) 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). Normative social influence should effect status goods more than non-status goods through the processes of identification and compliance. Identification is reflected in the need for psychological association with a person or group that an individual likes or wants to resemble. Consumers are motivated to identify themselves with higher-status individuals to improve their own social standing. Compliance stems from the motivation to comply with the product decisions of others in exchange for social returns. Consuming status goods in common with higher status individuals can establish common ground for communication and thus stronger relationships (Kuksov and Xie 2012). For example, golf vacations have been regarded as an effective networking tool in forming relationships with wealthy people. For these reasons, we hypothesize:

Hypothesis 2: *Friends' endorsements, in the form of social cues, have a greater effect on ad engagement for status goods than non-status goods.*

2.4. The Moderating Effects of User Status and Expertise

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 social interactions online has begun to uncover 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. In three seminal experiments, Aral and Walker (2012) identified individual-level variation in causal social influence on social media; Aral and Walker (2014) identified relationship-level variation in social influence online; and Taylor et al. (2015) documented the causal effect of different individual-level identities on the persuasion of social media posts. These studies show how individual-level characteristics of the source of online messages can cause viewers to change their opinions about the content of those messages and to subsequently change their behavior in response to those messages. In particular, Taylor et al. (2015) suggest that heterogeneity in responses to message sources intimate an opportunity to create "personalized social advertising" programs that "optimally select whom to show in a social advertisement and to whom, paving the way for personalized social advertising," although they don't examine social advertising in their work.

We therefore explore the moderating effects of users' characteristics on social advertising effectiveness. We theorize that the status and product expertise 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). A product endorsement is likely to be more influential if it comes from a high-status individual than if it comes from a low-status individual.

However, in this paper, we are not only concerned with individual-level variation in the influence of endorsements, but also with product-level variation. It is likely that status and expertise are not equally meaningful for all products, but instead that they are more meaningful for some products than for others. For example, status-based consumption is driven by relative comparisons to peers or friends. Individuals seek to identify themselves and communicate with those who have similar or higher social status (O'Cass and Frost 2002). If status considerations play a role in influence processes, friends' status relative to others will likely affect the degree of influence of a social endorsement. Furthermore, an endorser's status is likely to be more meaningful if they are endorsing a status good than if they are endorsing a non-status good. We therefore propose:

Hypothesis 3: *Friends' endorsements, in the form of social cues, have a greater effect on ad engagement for status goods than for non-status goods when the user whose cue is used in the ad is of higher status than the ad viewer.*

Expertise is another critical enabler of source persuasiveness. A meta-analysis of source effects in persuasion demonstrated that 16% of the variation in persuasion accounted for by source effects is explained by whether the source is perceived as an expert (Lipsey and Wilson 1993). It is therefore reasonable to assume that the expertise of a message source will impact the persuasiveness of their endorsements. Peers with high expertise tend to be regarded as credible information sources (Taylor et al. 2015). So those who are more knowledgeable, for example about a product, are likely to be more influential in the product evaluation process. Consumers are not likely to be concerned with the absolute product expertise of their friends but instead with their relative product expertise. Knowing that a friend is more knowledgeable than oneself about a particular product will likely increase the persuasive power of the friend's endorsement of that product. The experience of an expert will be more persuasive than the experience of a novice when those experiences are expressed in the form of endorsements in social advertising. But, relative expertise is not likely to be equally relevant for all products. Search goods are easily researched by seeking readily available factual information. Experience goods, on the other hand, are only really evaluated through hands on experience or the knowledge of someone else's hands on experience. We, therefore, expect relative expertise differences to be meaningful for evaluations of experience goods, where the experience of peers is relevant to the product evaluation process. Search goods, by comparison, are likely to be evaluated by researching factual information about the product online. We therefore propose:

Hypothesis 4: *Friends' endorsements, in the form of social cues, 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 expertise about the product being advertised than the ad viewer.*

3. 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, using instrumental variables (Bramoull et al. 2009, Tucker 2008), dynamic matched-sample estimation (Aral et al. 2009) and structural models (Ghose and Han 2010). 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).

The experiment was conducted on WeChat Moments Ads. WeChat Moments, like Facebook's news feed, 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 feed. Users can express their attitudes towards 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 page, landing page, and product photos (see Figure 1 for an example of a WeChat Moments Ad). Unlike Facebook, only first-degree friends' likes or comments are visible to WeChat users. As a result, social influence in our experiment has a consistent meaning that reflects the effect of a social cue (i.e. a like), representing first-degree friends' endorsements of an ad, on users' engagement with the ad (i.e. click-throughs). This reduces statistical interference from the confounding effects of correlated preferences among friends of friends.

During the experiment, as users were served new ads, they were randomly assigned to three groups: one without any social cue (the Control Group), 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) (See Figure 2). In this way, our treatments varied the existence of social cues and the

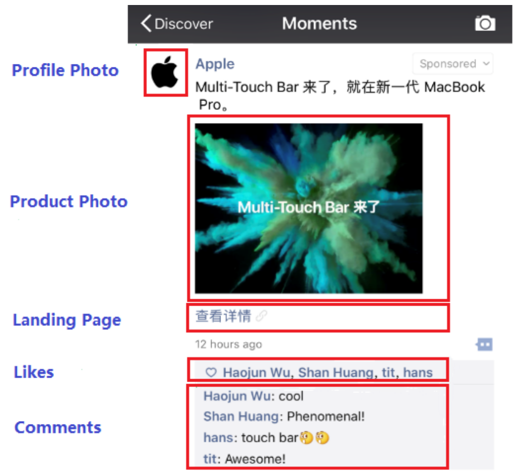

Figure 1: An Example WeChat Moments Ad
*Note.*This figure displays an example of WeChat Moments Ads including the brand profile, product photo, brand landing page, as well as social likes and comments.

number of social cues separately, allowing us to independently estimate the effects of exposure to social cues as well as the number of social cues displayed on ad engagement.

There are two types of social cues in Moments ads: friends' likes and comments. Since 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 during the experiment. Every ad stayed in a user's newsfeed for a maximum of 48 hours. After 48 hours, the old ad disappeared, and a new ad was served. In this way, users only saw one ad at a time in WeChat Moments during the experiment, which also reduced our susceptibility to statistical interference. Randomization occurred each time a user received a new ad. Users could be assigned to a different treatment group for each ad they saw, and the randomization took place at the user-ad level. In Treatment Group 1, when there was more than one organic like, the displayed like was the first available organic like.

Our experimental design avoids many known sources of bias in influence identification and networked experiments. First, it eliminates bias created by homophily by randomly assigning the social cues such that observed and unobserved attributes of users are equally distributed across

Figure 2: 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) respectively.

different groups. Second, the 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 only different in friends' names or profile pictures, eliminating the heterogeneity of unmeasurable 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. Since 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. Finally, our design avoids statistical interference and guarantees the stable unit treatment value assumption (SUTVA) is met (Rubin 1990). Users are randomized into different conditions for each ad, reducing the likelihood that they have a different experience and therefore talk to their friends about. Also, since users are only in one treatment condition at a time, they are not assigned to and therefore unlikely to be affected by simultaneous assignment to different treatment conditions.

During data collection, we recorded the number of organic social cues (i.e. organic likes), the number of social cues displayed on the ads due to treatment (i.e. displayed likes), the exact friends who were shown in the ads and to which viewers, ad viewers' responses to the ads (i.e., whether they click and their response time) and users' and their friends' demographic (i.e. age, gender and city) and behavioral characteristics on WeChat (i.e., number of login days to WeChat Moments and number of WeChat friends). For our main analysis, we only considered users' responses during their first impressions of a new ad and measured social influence as the effect of displaying a 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 showed 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 category that a product belongs to. We adopted the product categorization used in the WeChat ads department, which is a standard one used in the advertising industry.

We used multiple raters to classify the 71 experimental products on the two theoretically motivated product dimensionsin our study: experience vs. search goods and status vs. non-status goods, based on definitions that we provided them (See the Appendix for these definitions). Four independent judges, all undergraduate economics students at a prestigious Chinese university, separately classified the 71 products. Their inter-rater agreement ranged from 0.87 to 0.99 as measured by the Intraclass Correlation (ICC) and between 0.62 to 0.83 by Fleiss' Kappa. The four judges resolved all disagreements by consensus.

4. ANALYSIS

4.1. Model Specification and Estimation: Product Types

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 4.1.[2] Each observation represents a

[2] 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.

user-ad pair. The model simultaneously estimates the impact of two dimensions of product types: search/experience goods and status/non-status goods as shown here:

$$\log\left(\frac{Pr(Y_{ij}=1)}{1-Pr(Y_{ij}=1)}\right) = \beta_0 + \beta_1 S_{ij} + \beta_2 P_j + \gamma_1 (S_{ij} * P_j) + \theta_n C_{ij} + \psi_k (S_{ij} * C_{ij}) \tag{4.1}$$

where Y_{ij} is a dummy variable indicating whether user i clicked ad j during their first ad impression. S_{ij} is a dummy variable that indicates whether a user-ad pair (i, j) is in a particular social cue treatment group for ad j. P_j is a vector of product-type dummies, each of which indicates whether the product in ad j is an experience good or a search good and whether it is a status good or a non-status good. The coefficient β_1 captures the marginal effect of social cues on ad engagement. The coefficient β_2 captures the effect of product types on the tendency of users to enage with an ad. The coefficient γ_1 captures the impact of product types on the effectiveness of social ads.

We also include a comprehensive set of controls in the analyses to account for their effects on both clicking and social influence. First, since different ads target different users, it is necessary to control for the variables used for ad targeting. Previous studies also suggest that individual demographic characteristics significantly affect the magnitude of social influence (Aral and Walker 2012). Since our experiment was conducted at the very early experimental stages of WeChat Moments Ads, the targeting conditions were based simply on users' age, gender, and city, further reducing the confounding effects of ad targeting and isolating the effects of adding social cues to ads. Second, we controlled for characteristics of the affiliated friends whose social cues were displayed in the ads. Targeting conditions of ads also partly decide the demographic conditions of affiliated friends. Both influence and susceptibility are key factors that drive social influence (Aral and Walker 2012). As a result, it is necessary to take affiliated friends' demographic characteristics into consideration. Third, brand characteristics have 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. Fourth, since our experimental period covered the Christmas and New Year holidays, we included week dummies to control for time effects. The coefficients $\theta_n C_{ij}$ capture the variation in ad engagement explained by the vectors of control variables for the characteristics of individuals, affiliated friends, brands, products and the ad delivery week respectively. We follow Aral and Walker (2012) in controlling for the selection bias by specifying a comprehensive set of controls for $\psi_k(S_{ij} * C_{ij})$ which interact the treatment group of a user-ad pair (whether they are treated with social ads) with the control variables that describe that pair. Finally, users' adoption outcomes may be correlated for the same ads. To account for this, we specifed clustered standard errors at the ad level.

4.2. Measuring Users' Status and Expertise

Closeness centrality, which in our case is the average shortest path length between a node and all the other nodes in the entire WeChat social network, has been used as a reliable measure of social status in networked environments (e.g., Tucker 2008, Ransbotham et al. 2012). The more central a node is, the closer it is to all other nodes. If a user plays a central role in a social network, he or she is typically of high social status compared to others (Freeman 1978, Stephenson and Zelen 1989). Although this metric is computationally intensive to collect, because the social networks in our calculations are very large and cover almost all WeChat users and their contacts, our measure of closeness centrality as social status is more precise for the same reason. We divide the status of users and their friends into high, medium and low status using 25% and 75% cutoffs. Relative social status is then measured as the friends' status minus the ad viewers' status.

Product expertise is measured based on the accumulated number of product-related articles that an individual has read on WeChat historically. The WeChat team used a machine learning algorithm to predict users' interests in different fields, such as finance, technology and fashion etc. The input of the model is the number and type of articles that a user read on WeChat and the output is a vector of scores that measure the user's interests in various fields. If a product category matches a field, we use the users' score in that field to represent their expertise in that product category. This measure is then transformed into high, medium and low levels of expertise using the

25% and 75% cutoffs. Relative product expertise is measured as the friends' expertise minus the users' expertise.

4.3. Modeling the Moderating Effects of User Status and Expertise

We first specify a logistic regression model for experience and search goods and for status and non-status goods separately as in Equation 4.2. The model estimates the impacts of the relative social status and product expertise between the user shown in the ad and the ad viewer on users' tendency to engage with the ads for different types of products as follows:

$$\begin{aligned}\log\left(\frac{Pr(Y_{ij}=1)}{1-Pr(Y_{ij}=1)}\right) &= \beta_0 + \beta_1 S_{ij} + \beta_2 St_{ij} + \beta_2 Ex_{ij} + \gamma_1 (S_{ij} * St_{ij}) \\ &+ \gamma_2 (S_{ij} * Ex_{ij}) + \theta_n C_{ij} + \psi_k (S_{ij} * C_{ij})\end{aligned} \tag{4.2}$$

where St_{ij} and Ex_{ij} indicate the social status of the affiliated friend ij relative to user i (i.e. the status of friend ij minus the status of user i), and the expertise of the affiliated friend ij relative to user i (i.e., the friend ij's expertise minus user i's expertise, in that product category). The coefficients γ_1 and γ_2 on the interaction terms capture the impact of relative social status and expertise on social advertising effectiveness.[3]

We then estimate a logistic regression model in Equation 4.3 to test whether the γ_1 coefficient in Equation 4.2, the impact of the relative social status of the friend featured in the ad compared to the status of the ad viewer on ad performance, is significantly greater for status goods than for

[3] We also controlled for whether the product in the ad is an experience good for the sample of status and non-status goods and for whether the product in the ad is a status good for the sample of experience and search goods.

non-status goods as follows:

$$\log(\frac{Pr(Y_{ij}=1)}{1-Pr(Y_{ij}=1)}) = \beta_0 + \beta_1 S_{ij} + \beta_2 St_{ij} + \beta_3 Ex_{ij} + \beta_4 StG_j + \gamma_1 (S_{ij} * St_{ij})$$
$$+\gamma_2(S_{ij} * Ex_{ij}) + \gamma_3(S_{ij} * StG_j) + \gamma_4(St_{ij} * StG_j) + \gamma_5(Ex_{ij} * StG_j)$$
$$+\pi_1(S_{ij} * St_{ij} * StG_j) + \pi_2(S_{ij} * Ex_{ij} * StG_j)$$
$$+\theta_n C_{ij} + \psi_k(S_{ij} * C_{ij}) \tag{4.3}$$

where StG_j is a dummy variable that indicates whether the product in ad j is a status good or a non-status good. The coefficients π_1 and π_2 on the three-way interaction terms capture the difference in the impact of relative social status and expertise on social advertising effectiveness across status and non-status goods.

We run a similar regression in Equation 4.4 to test whether γ_2 in Equation 4.2, the impact of the relative product expertise of the friend featured in the ad compared to that of the ad viewer on social advertising effectiveness, is greater for experience goods or search goods, as follows:

$$logit(Pr(Y_{ij}=1) = \beta_0 + \beta_1 S_{ij} + \beta_2 St_{ij} + \beta_3 Ex_{ij} + \beta_4 ExG_j + \gamma_1 (S_{ij} * St_{ij})$$
$$+\gamma_2(S_{ij} * Ex_{ij}) + \gamma_3(S_{ij} * ExG_j) + \gamma_4(St_{ij} * ExG_j) + \gamma_5(Ex_{ij} * ExG_j)$$
$$+\pi_1(S_{ij} * St_{ij} * ExG_j) + \pi_2(S_{ij} * Ex_{ij} * ExG_j)$$
$$+\theta_n C_{ij} + \psi_k(S_{ij} * C_{ij}) \tag{4.4}$$

where ExG_j is a dummy variable indicating whether the product in ad j is an experience good or a search good. The coefficients ϕ_1 and ϕ_2 on the three-way interaction terms capture the difference in the impact of relative social status and product expertise on social advertising effectiveness across experience and search goods.

5. RESULTS

5.1. Descriptive Statistics

The experiment was conducted over a 21-day period in 2015-2016, during which 57,510,157 user-ad pairs, 37,951,299 distinct users, and 99 ads participated in the experiment. 19,198,084 user-ad pairs were randomly assigned to the control group with no social cues. 19,174,955 user-ad pairs were randomly assigned to the treatment group displaying a maximum of one like. 19,137,118 user-ad pairs were randomly assigned to the treatment group displaying the organic number of likes. We dropped 17 ads with invalid data and finally analyzed 82 ads for 71 distinct products across 25 product categories.[4] Among the 71 products, there are 48 experience goods, 23 search goods, 22 status goods and 49 non-status goods.

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. Some ads have no likes and we are therefore unable to display any real social cues on these ads as part of our manipulation. We therefore filtered the data on the condition that there was at least one organic like to guarantee that at least one social cue could be displayed on ads in Treatment Groups 1 and 2. This created a sample of 5,571,116 user-ad pairs and 4,884,070 distinct users across three treatment groups.

Assignment to treatment groups was random, with no economically meaningful mean differences between treatment groups in terms of their age, gender, city, network degree (i.e. number of WeChat friends) and level of WeChat Moments activity (i.e. log-in days) ($p > 0.1$) (See Table 3).[5] This evidence taken together confirms the integrity of the randomization procedure. No likes were displayed to users in the control group, and one like or the organic number of likes were correctly displayed to the users in the two treatment groups (See Table 2 for manipulation checks).

[4] We dropped 10 old ads, which were left over from the pre-experiment 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, perhaps indicating termination of under-performing ad campaigns.

[5] We tested for the Hypothesis: $\frac{\Delta Mean}{Mean} <$ max 0.5%, since hypothesis test: $\mu_0 = \mu_1$ would always be rejected in a very large dataset (Lin and Lucas 2013).

Table 2: Manipulation Checks

#Like	#user-ad pairs	#Displayed likes			
		Mean	Std	Max	Min
0	1,860,622	0.000	0.000	0	0
1	1,873,401	1.000	0.000	1	1
2	1,837,093	1.672	1.742	100	1

Note. "0" represents the control group. "1" represents treatment group 1. "2" represents treatment group 2.

Table 3: Mean Comparisons Between Treatment Groups

	$(\|\#0-\#1\| < \#0 * X\%)$	$(\|\#0-\#2\| < \#0 * X\%)$	$(\|\#1-\#2\| < \#1 * X\%)$
	t-statistic (X%)	t-statistic (X%)	t-statistic (X%)
Age	-0.371 (0.10%)	-13.882 (0.50%)	0.835 (0.10%)
Gender, 1=Male	1.410 (0.10%)	-0.629 (0.10%)	1.637 (0.10%)
City, 1=$Class_1$	0.618 (0.10%)	1.556 (0.10%)	0.269 (0.10%)
City, 1=$Class_2$	0.092 (0.10%)	0.371 (0.10%)	-0.843 (0.10%)
Degree	-1.269 (0.50%)	-0.778 (1.80%)	-1.179 (1.50%)
Login Days	-27.647 (0.10%)	-27.567 (0.10%)	-27.820 (0.10%)

Note. "0" represents the control group. "1" represents the treatment group 1. "2" represents the treatment group 2. We compared mean between three groups and tested the hypothesis : mean(a) - mean(b) < X% * mean(b). We reported the t-statistics. X% are in parentheses. All the tests are insignificant, indicating that mean(a) - mean(b) <X%*mean(b) for all the covariates between the three groups.

5.2. Heterogeneous Effects of Social Advertising Across Products

We are interested in characterizing the heterogeneity in social advertising effectiveness across products with data-driven exploratory analyses and theoretically motivated hypothesis tests. We therefore estimate the average treatment effect of a social cue on ad engagement and then investigate how those effects vary across different products (e.g. a BMW 325 automobile or a Kitchen-Aid blender), product categories (e.g. fashion products or electronics products) and theoretically motivated product types (e.g. search/experience goods or status/non-status goods). In this section we report on the heterogeneity in social ad effectiveness across products and product categories. In the next section we test hypotheses about how product types moderate the effectiveness of social advertising.

We first report average treatment effects in social advertising, which we define as adding social cues to advertisements in WeChat users' newsfeeds. We estimated marginal social influence across all user-ad pairs as the relative risk of users' average response rates (clicks) across control and treatment groups during users' first impressions on ads. Control group ad units were displayed without any social cues while 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) makes users 33.75% relative more likely to click an ad on average ($p < 0.01$).

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. 39 out of 71 products exhibit significantly positive lift from social advertising (see the dots in Figure 3), while 32 products experience no significant lift (See the crosses in Figure 3) and none perform worse when social cues are added to the ads.[6] 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). We also observed that there is significant heterogeneity in social advertising effectiveness across products. The highest product-level social influence (max influence = 3.70, p < 0.01) is 2.64 times as large as the average product-level influence (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). To report this heterogeneity, we grouped the product categories into four quantile groups according to their associated social influence. The average influence in four quantiles are all significantly positive and different from one another ($p < 0.01$) (See Figure 4). This demonstrates that significant heterogeneity exists in social advertising effectiveness across products.

Next, we aggregated the products into 25 categories and identified influence at the product category level (See Figure 5). 19 categories exhibit significantly positive lift from social advertising, while 6 experience no 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,

[6]There may not be enough statistical power to detect significant effects for every product.

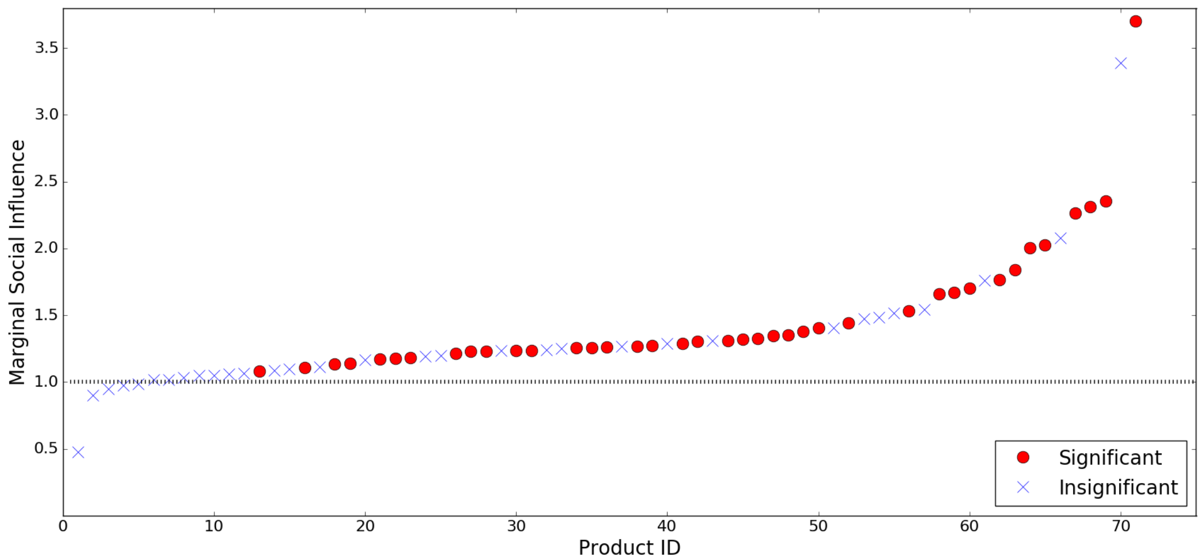

Figure 3: Social Advertising Effectiveness Across Products

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

$p < 0.01$), 1.57 times more effective than for electrical appliances (average marginal influence = 1.17, $p < 0.01$), and 1.55 times more effective than 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, while TV shows, E-commerce platforms and credit cards do not exhibit significant lift from social ads. In summary, we found a distribution of lifts from social ads, which is highly heterogenous across products and product categories.

5.3. How Product Types Moderate Social Advertising Effectiveness

The evidence in the previous section establishes that there is significant heterogeneity in the effectiveness of social advertising across products and product categories. The natural next step in our investigation then is to test theoretically motivated hypotheses about why this heterogeneity exists. In Section 2, we argued that the theoretically motivated dimensions of status/non-status goods and experience/search goods should moderate the importance of social cues for consumer decisions.

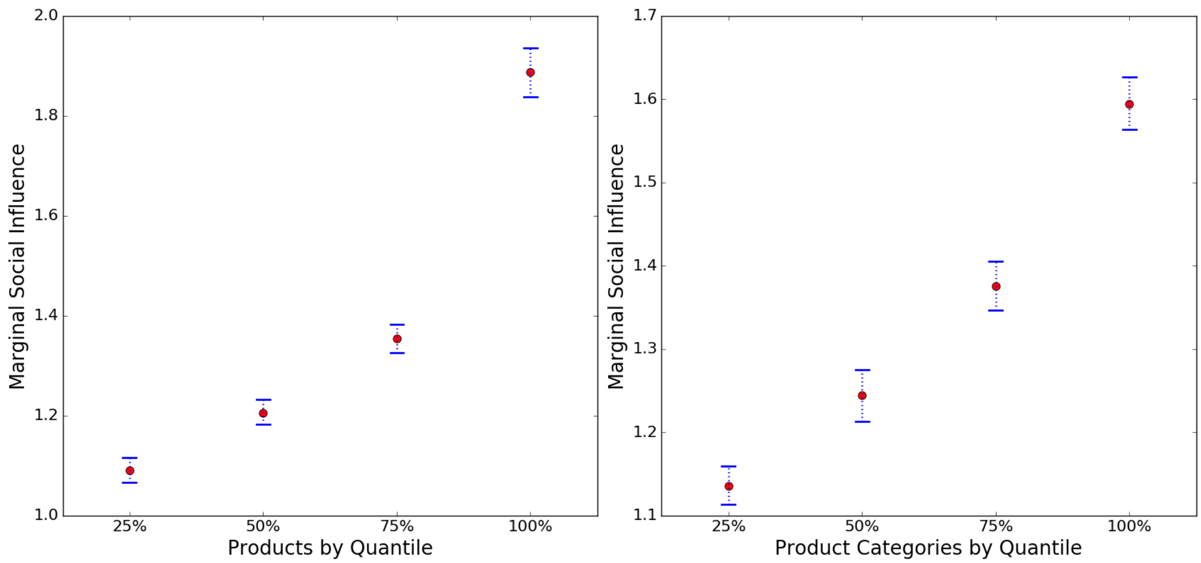

Figure 4: Social Advertising Effectiveness Across Products and Product Categories by Quantile
Note. Marginal social influence in each quantile of products is shown with SEs (boxes) and 95% bootstrapped confidence intervals (whiskers).

In this section we provide experimental evidence testing these two hypotheses 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 status/non-status goods. We do this using the same econometric framework we used to establish the heterogeneity in social advertising effectiveness across products and product categories.

Estimates of the impact of these product types on social advertising effectiveness are displayed in Table 4, while the forest plot in Figure 6 compares the e^{γ_1} estimates for search/experience goods and status/non-status goods and displays the standard errors (boxes) and 95% confidence intervals (whiskers) of these estimates. The e^{γ_1} estimates represent relative marginal social advertising effectiveness across these different product types.

The results corroborate the average marginal effectiveness of social advertising found in the previous section. Adding a social cue lifts the click-through rate for social ads by 20.19% in this

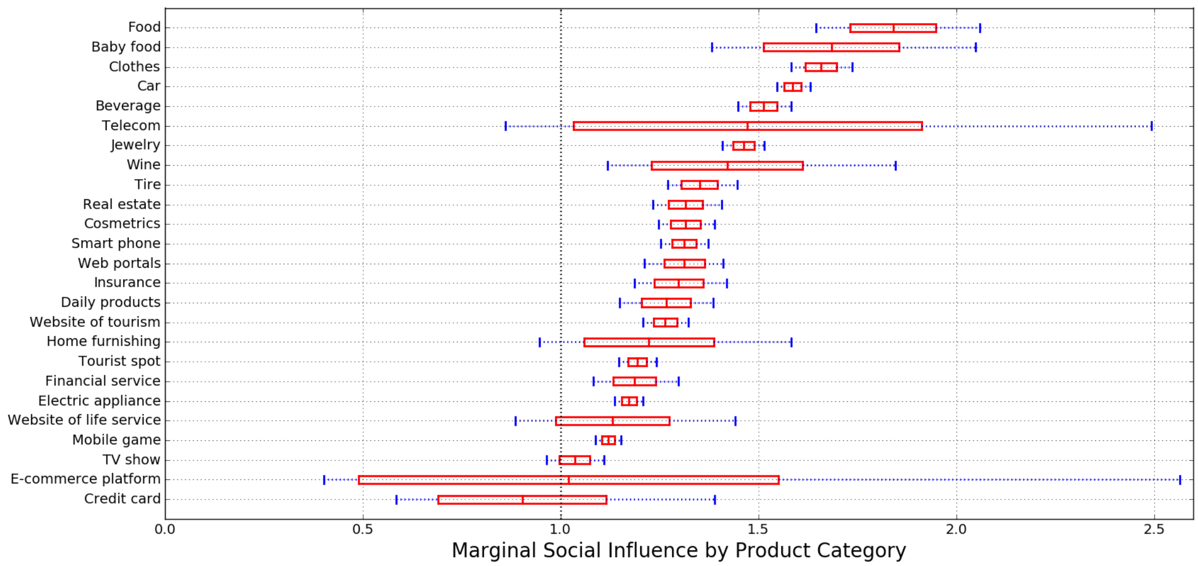

Figure 5: Social Advertising Effectiveness by Product Category

Note. Marginal social influence in each product category is shown with SEs (boxes) and 95% bootstrapped confidence intervals (whiskers).

analysis, down from the 33.75% in earlier analyses with fewer covariates.[7] We also found that social advertising is 17.14% more effective for status goods than non-status goods ($p < 0.05$), while there is no statistically significant difference in the performance of social advertising between experience and search goods. These results suggest that normative social influence, based on the Keynesian notion of relative utility, is a strong driver of social advertising effectiveness. 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. Hypothesis 2 is supported, while Hypothesis

[7]The *odds* approximates the probability ($Pr(Y_{ij} = 1)$), when the probability ($Pr(Y_{ij} = 1)$) is near zero: $odds_{ij} = \frac{Pr(Y_{ij}=1)}{1-Pr(Y_{ij}=1)} \approx Pr(Y_{ij} = 1), [Pr(Y_{ij} = 1) \approx 0]$ Since the click-through rate for ads in general and for our specific dataset are both close to zero, odds ratios closely approximate the *relative effects* of the variables on $Pr(Y_{ij} = 1)$.

Table 4: Social Advertising Effectiveness Across Search/Experience Goods and Status/Non-Status Goods

	1 Clicking	2 Clicking
Social Cue (SC)	1.2101*** (0.0290)	1.2019** (0.0913)
Experience Goods	0.5929 (0.2834)	0.6280 (0.2751)
Status Goods	0.5170 (0.2316)	0.6682 (0.2920)
SC*Experience Goods	1.0345 (0.0627)	1.0186 (0.0581)
SC*Status Goods	1.1981** (0.0863)	1.1714** (0.0867)
Controls		YES
Log-Likelihood	-534,820	-525,480
Observations	3,734,023	3,734,023

Note. An observation is a user-ad pair. The reported coefficients are the odd ratios. Robust standard errors, clustered at the ad-level, are reported in parentheses.
$*p < 0.10; **p < 0.05; ***p < 0.01$.

1 is not. It could be, however, that while there is no difference in social advertising effectiveness across search and experience goods on average, differences emerge when we consider the status and expertise of the user shown in the ad. We explore this possibility in the next section.

5.4. How User Status and Expertise Moderate Social Advertising Effectiveness Across Products

Models 1, 2 and 3 in Table 5 display and compare the moderating effects of the status and product expertise of the friend shown in a social ad, relative to the viewer of the ad, on social advertising effectiveness for status and non-status goods. Social status and product expertise significantly improve social ad effectiveness for both status ($p < 0.01$) and non-status goods ($p < 0.1$). Peers exert 7.02% more influence on viewers' ad engagement for status goods when the status difference between the friends increases by one class (e.g. low medium or high) ($p < 0.01$). For non-status goods, peers only exert 2.77% more influence on viewers' ad engagement when the difference in

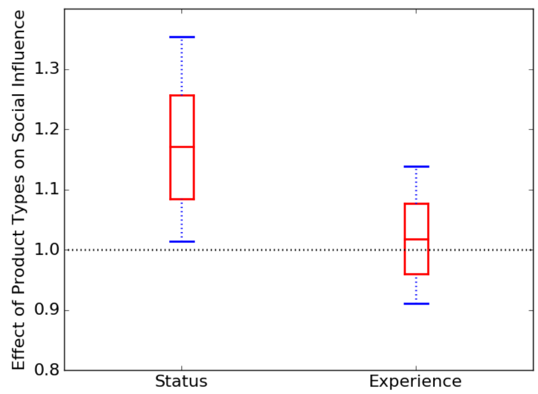

Figure 6: Relative Social Ads Effectiveness Across Search (Experience) Goods and Status (Non-Status) Goods
Note. The effects of product types on marginal social influence are shown with SEs (boxes) and 95% confidence intervals (whiskers).

status between the friend and the viewer increases by one class ($p < 0.1$). The effect of the relative social status of the peer compared to the ad viewer on social ad effectiveness is itself significantly greater for status-goods than non-status goods ($p < 0.1$). These results confirm that the relative status of peers is critical to social advertising effectiveness and that it is more important for status goods than for non-status goods. H3 is supported.

We also observe that relative product expertise significantly moderates social ads effectiveness for both status ($p < 0.05$) and non-status goods ($p < 0.01$), although these moderating effects are not significantly different from each other ($p > 0.1$). For status goods, friends' exert 3.97% more influence on viewers' ad engagement when their relative product expertise increases by one class, compared to the ad viewer ($p < 0.05$). For non-status goods, friends' exert 3.09% more influence on ad viewers when the difference in the product expertise between them increases by one class ($p < 0.01$). These results imply that consumers are persuaded by the social influence of friends with greater product expertise when evaluating products, while the distinction between status and non-status goods does not impact the extent of this persuasiveness.

Models 4, 5 and 6 in Table 5 compare the impact of friends' status and product expertise on

Table 5: Effects of Users' Status and Expertise on Social Advertising Effectiveness Across Products

	1 Status Goods Clicking	2 Non-Status Goods Clicking	3 Status vs. Non-Status Clicking	4 Experience Goods Clicking	5 Search Goods Clicking	6 Experience vs. Search Clicking
SC*FRS	1.0702*** (0.0161)	1.0277* (0.0164)	1.0277* (0.0164)	1.0608*** (0.0170)	1.0046 (0.0161)	1.0046 (0.0171)
SC*FRE	1.0397** (0.0156)	1.0309*** (0.0113)	1.0309*** (0.0113)	1.0411*** (0.0104)	1.0140 (0.0101)	1.0140 (0.0091)
SC*FRS*SG			1.0413* (0.0229)			
SC*FRE*SG			1.0085 (0.0182)			
SC*FRS*EG						1.0559** (0.0243)
SC*FRE*EG						1.0268* (0.0144)
Log-Likelihood	-262,810	-252,860	-515,640	-405,110	-112,760	-517,870
Observations	2,387,250	1,346,73	3,734,023	3,215,964	518,059	3,734,023

Note. SC = "Social Cue," FRS = "Friend's Relative Status," FRE = "Friend's Relative Expertise," SG = "Status Good," EG = "Experience Good." An observation is a user-ad pair. The reported coefficients are the odd ratios. Robust standard errors, clustered at the ad-level, are reported in parentheses.
$*p < 0.10; **p < 0.05; ***p < 0.01$.

social ads effectiveness for search and experience goods. Friends' product expertise significantly moderates the effect of social cues in ads for experience goods ($p < 0.01$) but not for search goods ($p > 0.1$). For experience goods, the lift from social cues increases by 4.11% ($p < 0.01$), when the expertise gap for the product between the friend and the ad viewer increases by one class. The effect of the relative product expertise of the peer compared to the ad viewer on social ads effectiveness is itself significantly greater for experience goods than for search goods (p < 0.1). H4 is supported. Together with the null result of the average affect of the difference between search and experience goods on social ads effectiveness (See Section 5.3), these results suggest that the distinction between search and experience goods only plays a role in social ads effectiveness when the expertise gap between the friend and the ad viewer is large. This suggests that informational social influence is important as a behavioral mechanism in social advertising only when the friend has more information or expertise than the ad viewer.

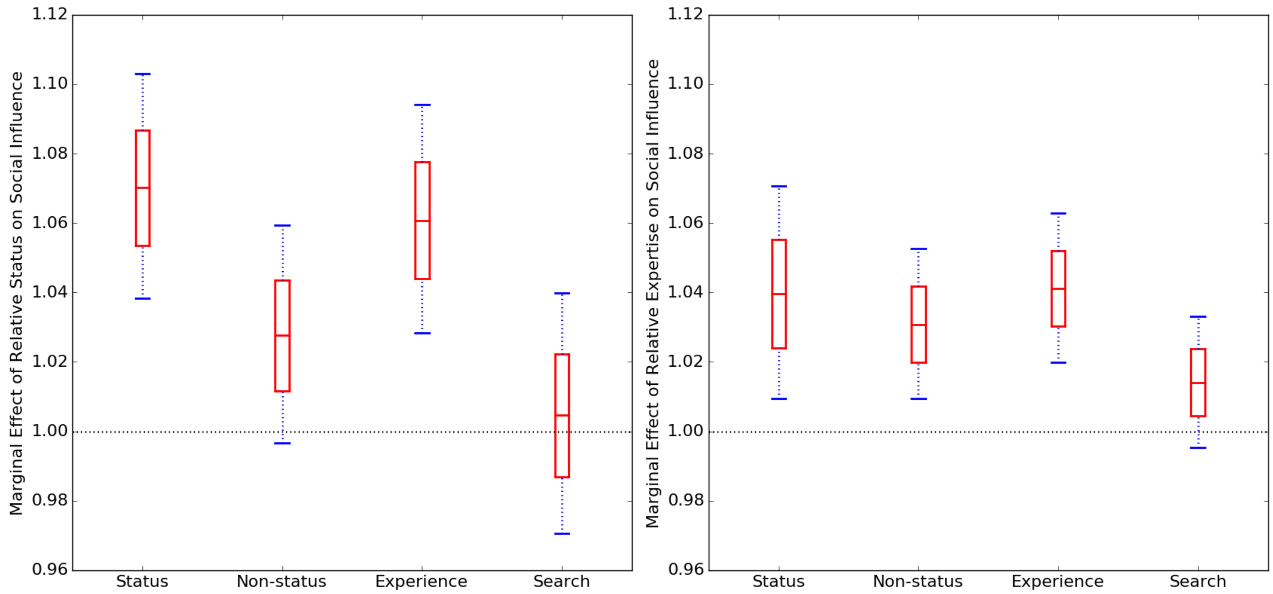

Figure 7: Effects of Users' Status and Expertise on Social Advertising Effectiveness Across Products
Note. The effects of the status and product expertise of the friend shown in a social ad, relative to the viewer of the ad, on marginal social influence for different types of goods are shown with SEs (boxes) and 95% confidence intervals (whiskers).

We also observe that friends' status significantly moderates the lift from social cues for experience goods ($p < 0.01$) but not for search goods ($p > 0.1$). The difference in the moderating effect of status on social advertising effectiveness is significant between experience and search goods (p < 0.01). For experience goods, friends' exert 6.08% more influence on viewers' ad engagement (p < 0.01), when the difference in their status increases by one class. This indicates that while consumers seem to value the product expertise of their friends for both search and experience goods, they value status more when evaluating experience goods than when evaluating 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.

5.5. Robustness Checks

We conducted multiple tests to ensure the robustness of our findings to alternative measurement approaches and model specifications. First, we measured the goodness-of-fit of the logistic regression models described in sections 5.3 and 5.4. We report the log-likelihood of our models in Tables 4 and 5, and the log-likelihood ratio chi-squared tests are significant ($p < 0.01$), indicating that the observed relationships are unlikely to have been found due to chance.

Second, in the main analysis, we had estimated social influence as the effect of adding one social cue to an ad on users' ad engagement during their first ad impression. This is one of the most relevant measures of social influence in social advertising because it is not confounded by, for instance, variation in the number of organic likes a product or brand receives. However, alternative measures of social influence, such as influence from the organic number of social cues, are also worthy of study because we'd benefit from knowing if influence is stronger for more socially liked products and how the strength of the impact of social cues on ad engagement varies with the number of social cues that appear in ads. We cannot measure the casusal impact of the number of social cues on ad engagement without statistical support for a variety of numbers of social cues available for experimental manipulation. However, we can see how ad engagment varies with the number of social cues. To examine social influence created by the organic number social cues an ad received (Mean = 1.672, Std = 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.[8] The results are presented in Tables B1 and B2. We find that the impact of product types (status/non-status and experience/search goods) on social influence does not change significantly when we estimate social ads effectiveness using the organic number of likes.

Displaying organic social cues was 24.53% more effective for status goods than non-status goods

[8]We measured the expertise gap across multiple friends as the maximum product expertise 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 expertise gap by the average expertise 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. The results are available upon request.

($p < 0.05$), but there, again, was no statistically significant difference in social ads 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 (24.53% > 17.1%) than when we showed only one social cue. This may be due to the additional influence of multiple sources of social proof or to the selection effect of choosing products which many friends have endorsed. These results broadly corroborate the finding that showing social cues in ads is more effective for status goods than non-status goods, but does not increase ad engagement for experience goods more than for search goods.

We also found that users were more influenced by the social cues of friends with higher-status and greater product expertise when they were evaluating both status goods and non-status goods ($p < 0.05$). The effect of relative social status on social advertising effectiveness was significantly greater for status goods than non-status goods ($p < 0.01$). This evidence confirms that status processes are a strong behavioral mechanism driving social advertising effectiveness and that they are especially strong when customers are evaluating status goods compared to non-status goods.

Users were persuaded by friends with greater product expertise only when they were evaluating experience goods ($p < 0.01$) although this effect was not significantly different than for search goods ($p = 0.104 > 0.1$). This result is slightly different than the one in Section 5.4 - users' tendency to conform to the peer with more product expertise was significantly stronger in evaluating experience goods than search goods when we only considered one social cue ($p = 0.060 < 0.01$). This could be because multiple cues improve social ads effectiveness across all products, making them all more effective or that the greater normative social influence created by organic social cues may interact with users' perceptions of product uncertainties (Huang 2016). These results corroborate that users' desire to seek advice from friends with greater expertise is stronger for experience goods than for search goods.

Third, it is important for marketers to understand how social cues impact ad viewers' responses during the entire span of their ad impressions (Mean Impressions = 3.03, Std = 2.52) to grasp the full magnitude of social ad effectiveness. While our main analysis measures the more precise and unconfounded marginal effect of adding social cues at the very first sight of a social ad, responses during the entire span of ad impressions evaluates users' decisions over a longer time horizon. We adopted a new dependent variable in our robustness check by counting any click on a given ad as long as users click the ad at any point during the time the ad was shown in their WeChat newsfeed.[9] The results with the new dependent variable are consistent with our main findings with a focus on users' first ad impressions. Displaying a social cue in ads was 14.33% more effective for status goods than non-status goods ($p < 0.05$), 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 Table 4 and Table B3). Users were less motivated to conform to the opinons friends with higher status but were more influenced by the social cues of friends with more expertise, when we considered their responses during the entire span of ad impressions (See the first and second rows in Table 5 and Table B4). These results imply that the influence of friends' 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 expertise during a longer span of consideration than on a first impression. But, putting these nuances aside, the general pattern of results and the importance of peers' status and expertise is confirmed by this robustness analysis.

Finally, to assess whether our estimates were affected by the regression models we specified, we replicated our analysis using probit models. The results estimated by probit models were consistent with the results estimated by the logit models (See Table B5 and B6 in Appendix), indicating that our findings are robust to these different specifications.

6. DISCUSSION

Stepping back from the detailed parameter estimates and robustness analyses, the collective evidence, reported across multiple specifications and operationalizations of key variables, points to three broad patterns in our results.

[9]The ads stayed in the WeChat news feed for no more than 48 hours during the experiment.

First, there is clear and significant heterogeneity in social advertising effectiveness across products. Whether we examined vertical categories or aggregated quintiles, there was tremendous variation in social ads' performance. Food, clothes and cars were the best performing categories, while TV shows, e-commerce platforms and credit cards performed the worst. Although nearly all products experienced lift from social ads, the best performing products were over three times as effective as the worst performing products. This heterogeneity suggests significant opportunities for brands to optimize their returns from social advertising. Simply allocating social advertising dollars to higher performing categories will generate significant returns, not to mention the myriad opportunities for improving the performance of poorly performing categories or making nuanced budget allocations which consider the cost of social advertising across categories as well as its performance.

Second, theory helps explain why some product categories perform better than others, contributing to the development of a more general theory of social contagion in product adoption. Previous research examined social contagion in product adoption by studying a single product at a time. Our findings show that generalizing to all products from studies of single products creates significant bias in our understanding of the effects of social influence on ad engagement and product adoption. Several theories explain why some products perform better than others. For example, theories of status production online and the role of status in consumption decisions help explain why status goods perform better than non-status goods in the context of social advertising. We look to our peers' opinions more when status is a factor of consumption. This may be because normative social influence is stronger for status goods, which triggers stronger identification and compliance processes when ad viewers evaluate ads. Although experience goods did not perform better than search goods on average, as we had hypothesized based on informational social influence, they did perform better when the friend shown in the ad was of higher status or had more expertise than the ad viewer. This highlights the third broad pattern in the collective results.

Third, the relative characteristics of the ad viewer and the friend shown in the ad significantly moderate social advertising effectiveness. This 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. We found that social ads that showed friends who were of higher status than the ad viewer were significantly more effective. Relative status moderated social advertising effectiveness for status goods significantly more than for non-status goods and for experience goods significantly more than for search goods. It seems that normative social influence is more important for status goods and experience goods than for non-status goods and search goods respectively. We also found that social ads that showed friends who had more expertise than the ad viewer were significantly more effective as well. Relative expertise significantly moderated social advertising effectiveness for status goods, non-status goods and experience goods, but not for search goods. These results shed light on the conditions under which informational social influence may operate, namely that the peer must be perceived to be more knowledgeable than the ad viewer for informational social influence to be impactful.

These broad results 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 programs. Although we conducted the first large-scale, experimental investigation of social advertising effectiveness across products and although the results (which are robust to multiple modeling specifications) are both theoretically interesting and practically relevant for researchers and marketing practitioners, our work is not without its limitations.

First, while we can evaluate how social ads effectiveness varies across the characteristics of the friends that are shown in the ads, we cannot make causal claims about these estimates. Rigorous 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.

Second, as with all studies in this area, the platform we studied skews heavily toward a particular community, in our case Chinese users. Although most studies of social influence online have this characteristic, they rarely highlight this limitation. In contrast, we feel it is important to circumscribe our conclusions accordingly. While 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. While status may be an important part of social influence processes in advertising in China, it may or may not be in the United States. Further work is needed to understand whether such cross cultural variation exists in social advertising.

Third, we have only scratched the surface of exploring variation in social influence processes across products. Many more dimensions of products and behaviors may be relevant to the strength of social influence and the subsequent effectiveness of social ads. While distinctions between status/non-status goods and experience/search goods are important, theoretically motivated dimensions of the phenomenon, other characteristics of products, experiences and behaviors may also moderate social influence. Furthermore, many individual and relative characteristics of ad viewers and the friends shown in ads may also be relevant. We strongly feel that more empirical and theoretical work is needed to produce a more complete theory of social influence and behavioral contagion. Despite these limitations, we hope our work will move us one step closer these goals.

7. CONCLUSION

Consumers make billions of decisions about what products to buy every day and prior work shows that they rely heavily on the opinions of others in making those decisions. However, the nature of this social influence, and in particular whether, how and why it varies across products, has not been well understood. To address these gaps in our knowledge, we designed and conducted a randomized field experiment on WeChat involving 37 million users to measure social influence and social advertising effectiveness across 71 products in 25 categories.

As a result, we were able to causally identify the large and very heterogeneous effects of social influence across products and product characteristics. Social influence was over three times higher for some products than for others. Furthermore, the differences in social influence could be partially predicted by the nature of the products themselves. Our findings not only shed light on the heterogeneous role of social influence across products, but also have broad implications for how marketers should promote product diffusion in social networks and how they can differentiate and price their social advertising strategies. For instance, strategically displaying social cues in users' first ad impressions will significantly increase the effectiveness of social ads, and this effect varies predictably across products. Further benefits can be attained by initially targeting users who are likely to endorse an ad even in the absence of influence and then displaying these user-endorsed ads to peers, especially peers who are susceptible to influence. The results can also help platform managers optimally price social ads for different products by precisely measuring the heterogeneity in the returns to social advertising across products.

Our work also contributes to the rapidly growing literature on social influence and behavioral contagion. Building more complete theories of social influence and behavioral contagion requires theoretically motivated hypothesis tests that uncover the nuances of when, where and how influence processes operate. By examining theoretically motivated dimensions of products such as status/non-status goods and experience/search goods, and by testing the moderating effects of the theoretically motivated characteristics of peers shown in social ads, our work moves us toward more nuanced and complete theories of social influence and behavioral contagion. While the types of product and user heterogeneity we studied only represent first steps toward explaining how social influence and social advertising effectiveness vary across products, they are far from the final word. Fortunately, large-scale field experiments, like the one we describe in this study, provide a powerful tool for revealing the causal relationships among product characteristics, social influence and user behavior.

References

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

Aral S (2016) Networked Experiments. *The Oxford Handbook of the Economics of Networks* (Oxford University Press).

Aral S, Muchnik L, Sundararajan A (2009) Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences* 106(51):21544–21549.

Aral S, Walker D (2011) Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Science* 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 Science* 60(6):1352–1370.

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

Bakshy E, Eckles D, Yan R, Rosenn I (2012) Social influence in social advertising: Evidence from field experiments. *Proceedings of the 13th ACM Conference on Electronic Commerce*, volume 1, 146–161.

Bakshy E, Karrer B, Adamic La (2009) Social influence and the diffusion of user created content. *Electronic Commerce*, 325–334.

Bapna R, Umyarov A (2015) Do your online friends make you pay? A randomized field experiment on peer influence in online social networks. *Management Science* 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. *Journal of Marketing Research* 51(3):270–285.

Bearden WO, JEtzel M (1982) Reference group influence on product and brand purchase decisions. *Journal of Consumer Research* 9(2):183–194.

Bearden WO, Netemeyer RG, Teel JE, Bearden W, Netemeyer RG, Teel JE (1989) Measurement of consumer susceptibility to interpersonal influence. *Journal of Consumer Research* 15(4):473–481.

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

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

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

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

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.

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

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

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

Chaiken S (1980) Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology* 39(5):752–766.

Chaiken S (1987) The heuristic model of persuasion. *In Social influence: the ontario symposium* 5:3–39.

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

Cialdini RB, Goldstein NJ (2004) Social influence: Compliance and conformity. *Annual Review of Psychology* 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. *Journal of Marketing* .

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

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

Freeman LC (1978) Centrality in social networks conceptual clarification. Social networks. *Social Networks* 1(3):215–239.

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

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

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

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

Hass R (1981) Effects of source characteristics on cognitive responses and persuasion. *Congitive Responses in Persuasion*, 141–172 (Psychology Press).

Huang S (2016) Social influence in public and private behaviors. *Proceedings of Thirty Seventh International Conference on Information Systems(ICIS)*, 1–16 (Dublin).

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

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

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 US presidential election. *PloS one* 12(4):e0173851.

Kelman H (1958) Compliance, identification, and internalization three processes of attitude change. *Journal of Conflict Resolution* 2(1):51–60.

Keynes JM (2010) Keynes JM. Economic possibilities for our grandchildren. *In Essays in persuasion*, 321–332 (Palgrave Macmillan UK).

Kuksov D, Xie Y (2012) Competition in a status goods market. *Journal of Marketing Research* 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. *Journal of Business Research* 62(7):706–712.

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. *Information Systems Research* 25(3):468–488.

Lin M, Lucas HC (2013) Too big to fail : Large samples and the p-value problem. *Information Systems Research* 24(4):906–917.

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

Lipsey MW, Wilson DB (1993) The efficacy of psychological, educational, and behavioral treatment. Confirmation from meta-analysis. *The American psychologist* 48(12):1181–1209.

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

Manski CF (1993) Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies* 60(3):531.

Muchnik L, Aral S, Taylor SJ (2013) Social influence bias: a randomized experiment. *Science (New York, N.Y.)* 341(6146):647–51.

Nelson P (1970) Information and consumer behavior. *Journal of Political Economy* 78(2):311–329.

O'Cass A, Frost H (2002) Status brands: Examining the effects of non-product-related brand associations on status and conspicuous consumption. *Journal of Product & Brand Management* 11(2):67–88.

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

Pesendorfer W (1995) Design innovation and fashion cycles. *The American Economic Review* 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. *Journal of Interactive Marketing* 27(4):281–298.

Ransbotham S, Kane G, Lurie NH (2012) Network characteristics and the value of collaborative user generated content. *Marketing Science* 31(3):387–405.

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

Sacerdote B (2001) Peer effects with random assignment: Results for Dartmouth roommates. *The Quarterly Journal of Economics* 116:681–704.

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

Sahni NS (2016) Advertising spillovers: evidence from online field experiments and implications for returns on advertising. *Journal of Marketing Research* 53(4):459–478.

Schmalensee R (1978) A model of advertising and product quality. *Journal of Political Economy* 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. *Journal of Marketing Research* 49(5):624–639.

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

Stephenson K, Zelen M (1989) Rethinking centrality: Methods and Examples. *Social Networks* 11(1):1–37.

Taylor S, Bakshy E, Aral S (2013) Selection effects in online sharing: Consequences for peer adoption. *Proceedings of the fourteenth ACM conference on Electronic commerce* 1(212):821–836.

Taylor SJ, Bakshy E, Eckles D, Aral S (2015) Heterogeneous social influence from personalized social cues.

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

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

Tucker CE (2016) Social advertising: How advertising that explicitly promotes social influence can backfire.

Turner JC (1991) *Social influence.* (Thomson Brooks/Cole Publishing Co).

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

Veblen T (2009) *The theory of the leisure class* (Oxford University Press).

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

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

Zhang J, Liu P (2012) Rational herding in microloan markets. *Management Science* 58(5):892–912.

Appendix A DEFINITIONS OF PRODUCT TYPES

Table A1: Product Type

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\Non-Status 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 consumers objective income or social class level. Examples of status goods include status conferring clothing, cars, wines, restaurants and hotels. Examples of non-status goods include toothpaste, beverages and website services.

Appendix B ROBUSTNESS CHECKS

Table B1: Social Advertising Effectiveness Across Search/Experience Goods and Status/Non-Status Goods: Control Group vs. Treatment Group 2

	1 Clicking	2 Clicking
Social Cue (SC)	1.2188*** (0.0451)	1.3114*** (0.1167)
Experience Goods	0.5929 (0.2834)	0.635 (0.2781)
Status Goods	0.517 (0.2316)	0.6631 (0.2891)
SC*Experience	1.0552 (0.0812)	1.0253 (0.0677)
SC*Status	1.2596*** (0.1058)	1.2453** (0.1096)
Controls		YES
Log-Likelihood	-540,520	-531,280
Observations	3,697,715	3,697,715

Note. An observation is a user-ad pair. The reported coefficients are the odds ratios. Robust standard errors, clustered at the ad-level, are reported in parentheses.
$*p < 0.10; **p < 0.05; ***p < 0.01$.

Table B2: Effects of Users' Status and Expertise on Social Advertising Effectiveness Across Products: Control Group vs. Treatment Group 2

	1 Status Goods Clicking	2 Non-Status Goods Clicking	3 Status vs. Non-Status Clicking	4 Experience Goods Clicking	5 Search Goods Clicking	6 Experience vs. Search Clicking
SC*FRS	1.0787*** (0.0108)	1.0222** (0.0112)	1.0222** (0.0112)	1.0575*** (0.0137)	1.0228 (0.0153)	1.0228 (0.0153)
SC*FRE	1.0526*** (0.0158)	1.0254** (0.0103)	1.0254** (0.0103)	1.049*** (0.0115)	1.0152 (0.0173)	1.0152 (0.0173)
SC*FRS*SG			1.0553*** (0.0158)			
SC*FRE*SG			1.0264 (0.0185)			
SC*FRS*EG						1.034* (0.0196)
SC*FRE*EG						1.0332 (0.0207)
Log-Likelihood	-268,270	-253,340	-521,610	-411,020	-112,880	-523,900
Observations	2,358,227	1,339,488	3,697,715	3,182,600	515,115	3,697,715

Note. SC = "Social Cue," FRS = "Friend's Relative Status," FRE = "Friend's Relative Expertise," SG = "Status Good," EG = "Experience Good." An observation is a user-ad pair. The reported coefficients are the odd ratios. Robust standard errors, clustered at the ad-level, are reported in parentheses.
$*p < 0.10; **p < 0.05; ***p < 0.01$.

Table B3: Social Advertising Effectiveness Across Search/Experience Goods and Status/Non-Status Goods: Multiple Impressions

	1 Clicking	2 Clicking
Social Cue (SC)	1.1965*** (0.0219)	1.1743** (0.0883)
Experience Goods	-0.4977 (0.2643)	-0.4487 (0.2606)
Status Goods	-0.5845 (0.2293)	-0.3078 (0.3056)
SC*Experience Goods	0.0388 (0.0558)	0.0122 (0.0491)
SC*Status Goods	1.1677** (0.0747)	1.1433** (0.0706)
Controls		YES
Log-Likelihood	-777,490	-766,060
Observations	3,734,023	3,734,023

Note. An observation is a user-ad pair. The reported coefficients are the odds ratios. Robust standard errors, clustered at the ad-level, are reported in parentheses.
$*p < 0.10; **p < 0.05; ***p < 0.01$.

Table B4: Effects of Users' Status and Expertise on Social Advertising Effectiveness Across Products: Multiple Impressions

	1 Status Goods Clicking	2 Non-Status Goods Clicking	3 Status vs. Non-Status Clicking	4 Experience Goods Clicking	5 Search Goods Clicking	6 Experience vs. Search Clicking
SC*FRS	1.0662*** (0.0142)	1.0233** (0.0093)	1.0233** (0.0093)	1.0588*** (0.0120)	-0.0052 (0.0102)	-0.0052 (0.0100)
SC*FRE	1.0538*** (0.0114)	1.0343*** (0.0102)	1.0343*** (0.0102)	1.0505*** (0.0090)	1.0180* (0.0098)	1.0180* (0.0096)
SC*FRS*SG			1.0420** (0.0166)			
SC*FRE*SG			0.0187 (0.0148)			
SC*FRS*EG						1.0643*** (0.0161)
SC*FRE*EG						1.0320** (0.0131)
Log-Likelihood	-399,380	-353,110	-752,490	-601,390	-154,820	-756,220
Observations	2,387,250	1,346,773	3,734,023	3,215,964	518,059	3,734,023

Note. SC = "Social Cue," FRS = "Friend's Relative Status," FRE = "Friend's Relative Expertise," SG = "Status Good," EG = "Experience Good." An observation is a user-ad pair. The reported coefficients are the odds ratios. Robust standard errors, clustered at the ad-level, are reported in parentheses.
$*p < 0.10; **p < 0.05; ***p < 0.01$.

Table B5: Social Advertising Effectiveness Across Search/Experience Goods and Status/Non-Status Goods: Probit Model

	1 Clicking	2 Clicking
Social Cue (SC)	0.0063*** (0.0010)	0.0044** (0.0020)
Experience Goods	-0.0199 (0.0193)	-0.0166 (0.0166)
Status Goods	-0.0215 (0.0158)	-0.0116 (0.0138)
SC*Experience	0.0007 (0.0013)	0.0004 (0.0013)
SC*Status	0.0048*** (0.0017)	0.0042** (0.0020)
Controls		YES
Log likelihood	-534,650	-525,310
Observations	3,734,023	3,734,023

Note. An observation is a user-ad pair. The reported coefficients are the marginal effects. Robust standard errors, clustered at the ad-level, are reported in parentheses.
$*p < 0.10; **p < 0.05; ***p < 0.01$.

Table B6: Effects of Users' Status and Expertise on Social Advertising Effectiveness Across Products: Probit Model

	1 Status Goods Clicking	2 Non-Status Goods Clicking	3 Status vs. Non-Status Clicking	4 Experience Goods Clicking	5 Search Goods Clicking	6 Experience vs. Search Clicking
SC*FRS	0.0014*** (0.0003)	0.0011** (0.0005)	0.0007* (0.0004)	0.0014*** (0.0003)	0.0002 (0.0009)	0.0001 (0.0005)
SC*FRE	0.0007** (0.0003)	0.0011*** (0.0003)	0.0008*** (0.0003)	0.0009*** (0.0002)	0.0004 (0.0004)	0.0002 (0.0003)
SC*FRS*SG			0.0010* (0.0006)			
SC*FRE*SG			0.0001 (0.0005)			
SC*FRS*EG						0.0015** (0.0006)
SC*FRE*EG						0.0007* (0.0004)
Log likelihood	-262,570	-252,800	-515,370	-405,240	-112,770	-518,010
Observations	2,387,250	1,346,773	3,734,023	3,215,964	518,059	3,734,023

Note. SC = "Social Cue," FRS = "Friend's Relative Status," FRE = "Friend's Relative Expertise," SG = "Status Good," EG = "Experience Good." Robust standard errors, clustered at the ad-level, are reported in parentheses. $*p < 0.10$;$**p < 0.05$;$***p < 0.01$.