The Effectiveness of Social Advertising

Shan Huang[1] and Song Lin[2]

[1]University of Washington, Seattle

[2]Hong Kong University of Science and Technology

Abstract

Although social advertising has grown to be one of the major online advertising channels in recent years, its effectiveness is not been fully understood. In this study, we use data from a large-scale field experiment on a major social media platform (WeChat Moments) to investigate how the display of social cues (friends' likes) in an advertisement affects users' responses. In the experiment, we randomly manipulate the presence of social cues in ads shown to 37 million users. We distinguish two types of consumer response: publicly observable responses that reveal whether a user has liked an ad, and private responses whereby a user clicks on an ad. We find that on average, displaying the first social cue significantly enhances the liking and clickthrough rates. However, while showing additional social cues can further increase users' tendency to like, it does not increase the clickthrough rate any further. This empirical pattern is consistent with the coexistence of informational social influence and normative social influence in social advertising. We find evidence that informational influence has a greater impact on the clickthrough rate, whereas normative social influence has a more prominent effect on the liking rate. Our results provide rich implications for advertisers and social media platforms in designing social advertising policies.

Keywords: *social networks; social advertising; social influence; field experiment*

Contact. shanh18@uw.edu.

1. INTRODUCTION

Over the last decade, the growth of online social media platforms such as Facebook, Instagram, Twitter, Snapchat, and WeChat has not only changed how consumers interact with each other, but also provided new channels for advertisers to communicate with their targeted customers. As a result, consumers are increasingly discovering products and updating brand preferences based on the information they gain from social media, and marketers have come to realize the value of advertising on social media platforms. According to *eMarketer*, in the United States, advertising spending on social media is projected to account for 39% of digital ad spending in 2020. However, despite its rapid growth, the effectiveness of social advertising is still not fully understood.

Social advertising, unlike other advertising channels, is embedded in social networks and thus uses social influence to communicate with consumers. By far the most common practice is to embed visible social cues in an ad, from which consumers can learn about how others in their network react to the ad. The objective of this study is to understand the effectiveness of these social cues. Importantly, we identify two dimensions of the effectiveness. The first dimension is how consumers react to social cues by *clicking* on an ad. This measurable ad response is what differentiates online media from traditional broadcast media. On a social media platform, this response is typically private; it is not disclosed to a user's network. The second dimension is how consumers respond to social cues in terms of revealing their interests, attitudes, or feelings about an ad. Many social media platforms allow users to express their *liking* of an ad, which can be revealed to others in their network and induce additional ad responses. This type of ad response is often unobservable in other online media and broadcast media, and thus constitutes a unique feature of social media.

The differences between these two types of response have important implications for advertising policies. Firms often differ in their advertising objectives. Some may be primarily interested in driving product sales, and thus their foremost priority is to persuade consumers to click on their ads. Other firms may simply aim to capture consumers' attention or build their brand image or relationships, with less focus on clicks and more on public attitude. Understanding the effectiveness of social advertising in relation to these different dimensions can not only inform firms about how to design their advertising campaigns and allocate their advertising budgets, but also can provide insights into how social media platforms can use social cues in their ads and determine the metrics for measuring advertising effectiveness.

To investigate the effectiveness of social advertising, we conducted a large-scale randomized field experiment on one of the world's largest social networking platforms, WeChat. We randomly allocated ads shown on WeChat Moments, on which WeChat users share their status, personal stories, and thoughts with their friends. The presence of friends' likes (social cues) on an ad for a user was randomized over more than 57 million ad-user pairs. This intervention lasted for 21 days and covered 99 ads. With the experimental randomization, the impression-level data on social cues and the users' ad responses, the individual-level data on the users' demographics and historical behavior, and the data on the ad characteristics, we are able to evaluate the causal effects of social advertising, and explore the underlying mechanisms.

We find that displaying friends' likes in ads significantly increases users' likelihood of both clicking and liking an ad. Compared with displaying no likes, displaying one like in an advertisement causes, on average, a 0.98% increase in the liking rate and a 0.96% increase in the clickthrough rate. This result provides evidence that social influence affects social advertising, and that the overall effect is positive. We then examine how additional social cues influence the users' responses. Interestingly, displaying more friends' likes on average further enhances the users' liking propensity, but does not increase the clickthrough rate further. Hence, our findings suggest that social cues can drive different user response dynamics. On the one hand, user likes can create a snowball effect whereby one like generates another, which further generates more likes. On the other hand, more likes do not necessarily induce more clicks, as our findings indicate that only the first social cue significantly affects users' private responses.

The diverging effects of social cues on clicks and likes identified in our field experiment provide a rare opportunity to investigate two distinct mechanisms of social influence, namely *informational* influence and *normative* influence. Deutsch and Gerard (1955) were among the first to distinguish the two types of social influence. They define informational social influence as "an influence to accept information obtained from another as evidence about reality." Campbell and Fairey (1989) further interpret informational social influence as based on the desire to be accurate, stating that "other's responses are used as a source of information about reality, and people conform because they believe that the others may be correct." In contrast, normative social influence reflects the urge to conform to the expectations of another (Deutsch and Gerard 1955), which is based on the desire to maximize social outcomes (Campbell and Fairey 1989). It is well known that the effect of normative social influence is reduced when individual actions are not observable by others (e.g., Deutsch and Gerard 1955, Asch 1956, Mouton et al. 1956, Argyle 1957, Levy 1960, Insko et al. 1983, Insko et al. 1985, Campbell and Fairey 1989). In our social advertising setting, although clicks and likes are probably subject to informational social influence, because likes are publicly observable but clicks are not, normative social influence is likely to have a greater impact on likes. Thus, users are likely to conform to the actions or expectations of others in their network to obtain social approval or to build relationships with them. This normative force presumes that the conformity inducing action is publicly observable to others. In contrast, users' clicks are not observable to the public and thus they are primarily driven by personal interests, in which the informational value of social cues plays a more prominent role.

To test this proposition, we identify scenarios in which either normative social influence or informational social influence is more salient and compare the impacts of displaying social cues on likes versus clicks. First, when consumers are more familiar with a brand and trust it, such as a well-known or reputable brand, social cues have less informational value. Thus, we find that exposing more social cues leads to a further increase in the liking rate, but not in the clickthrough rate for these well-known brands. In fact, the effects of increasing social cues on clickthrough rate are even slightly reduced for them. Conversely, when consumers are less familiar with a brand, such as a lesser-known brand, exposure to more social cues can increase not only the liking rate but also the clickthrough rate. Second, when consumers are more willing to engage with friends on social media newsfeeds, such as by endorsing (expressing likes) or commenting on friends' posts, they are more likely to be subject to normative social influence. Liking and commenting on friends' posts on social media have become meaningful ways of showing conformity with the opinions and behavior of friends. We find that increasing exposure to social cues has a much stronger positive impact on the liking rates for these users than for those who are less socially engaged, although the impact on the clickthrough rate remains similar across these two user segments. Together, these results are consistent with the prediction that normative social influence has a stronger effect on public responses (i.e., the liking rate) than on private responses (i.e., clickthrough rate), whereas informational social influence plays a more critical role in determining private responses than public responses.

The coexistence of informational and normative social influence, and the different roles they play in shaping private and public responses, shed important light on the management of social advertising. For firms interested in brand advertising, social advertising can be useful in driving public responses through the mechanism of normative social influence. These firms can further enhance the effects of social cues on liking by targeting users who are more susceptible to normative social influence (e.g., socially engaged users). In contrast, for firms interested in performance advertising, the effectiveness of social media advertising campaigns is more subtle when they are evaluated based on the clickthrough rate. Although displaying one social cue can significantly increase the clickthrough rate, the impact of more social cues on clicks is dependent on the relative weights of the informational and normative social influences. Although normative influence can stimulate the growth of social cues, it accidentally introduces noise when users try to make inferences about the product quality based on the social signals. For example, when normative social influence takes the dominant role in the liking response for a well-known or reputable brand, users may find the social cues (likes) less informative. Hence, policies that aim to stimulate social cues (public responses) by strengthening the normative social influence may negatively affect the informational social influence manifested in private responses such as clicks. Thus, an important message for social media platforms seeking to monetize from advertising is that they should carefully design their social advertising policies (e.g., social cues disclosure policies and targeting strategies) and manage social cues by taking into account the objectives of the advertisers (i.e., brand or performance advertising) and distinguishing between the two types of advertising metrics (i.e., the extent of public versus private responses) and their underlying mechanisms (i.e., normative versus informational social influence).

Related Literature

This study contributes to the nascent but growing body of literature on social advertising. In one of the earliest empirical studies on the effectiveness of social advertising, Bakshy et al. (2012) conduct two field experiments on the Facebook News Feed. They compare the effectiveness of an ad when it displays one, two, and three social cues (with friends' names identified), as well as when it displays the total number of endorsements (without friends' names), and find evidence that showing social cues can enhance ad performance. Our experiment includes a baseline group without any displayed social cues and treatment groups that display one or organic number of social cues. More importantly, we provide further evidence of the effect of social influence in social advertising in a different context from the Facebook network. WeChat Moments shows 100% of users' contents, including their posts, likes, and comments, but only to their *first-degree* friends. In contrast, Facebook users, on average, receive less than 10% of the organic feeds. Moreover, the users' public behavior can be observed by strangers and acquaintances, such as their second-degree friends on Facebook. This important difference of the product strategies and network settings implies that WeChat Moments exerts a stronger *normative* social influence on users than Facebook, leading to different impacts of social advertising. Different from Bakshy et al. (2012), we also explore the underlying behavioral mechanisms, find that private (clicks) and public (likes) responses to social ads follow different dynamics of social influence, and highlight the importance of distinguishing between informational and normative social influences when investigating the effectiveness of social advertising.

Social influence receives considerable interest from social science researchers (e.g., Asch 1955, Deutsch and Gerard 1955, Burnkrant and Cousineau 1975, Sacerdote 2001, Cialdini and Goldstein 2004, Christakis and Fowler 2013). As previously mentioned, informational influence and normative influence are two distinct types of social influence. With respect to informational social influence, individuals often face uncertainty when making decisions and information from others can help inform their decisions. In some cases, the information can be obtained by observing the decisions of others. It is then possible that individuals take the same actions as others after observing their choices, which is known as herding or observational learning (e.g., Banerjee 1992, Bikhchandani et al. 1992, Zhang 2010). In other cases, individuals may learn about the opinions or preferences of others, which directly influence their decisions. This is often referred to as word-of-mouth (e.g., Arndt 1967, Godes and Mayzlin 2004, Chevalier and Mayzlin 2006). Both mechanisms are studied extensively in the literature, and our study provides additional evidence that although informational social influence can be effective in the context of social advertising, it may be confounded by normative social influence.

With respect to normative social influence, individuals may follow others' actions or conform to others' expectations because doing so can directly enhance their utility. Social psychologists and behavioral scientists have long investigated why individuals choose to yield to the influence of others to explain social norms, conformity, and compliance (see Cialdini and Trost 1998 for a comprehensive review). However, distinguishing normative social influence from informational social influence is empirically challenging, because both forces can lead to conformity in actions. Thus, most studies rely on laboratory experiments to isolate these two influences. In this regard, the data from our field experiment offer a unique opportunity to investigate the difference between the two types of influence, and thus add valuable field evidence to this line of research. The concept of normative social influence typically rests on the premise that individuals desire to gain the approval of or build relationships with their group, friends, or associates (Cialdini and Goldstein 2004). To effectively fulfill these goals, an individual's actions must be identified by others (Deutsch and Gerard 1955). Thus, the data on both private (clicks) and public (likes) responses allow us to shed light on the relative effects of normative and informative influences on social advertising.

The remainder of this paper is organized as follows. In Section 2, we describe the research setting, experimental design, and data. Section 3 presents the empirical results on the impact of social cues on public versus private responses. We explore the underlying mechanism of the empirical pattern in Section 4, and discuss its managerial implications in Section 5. Section 6 concludes the paper.

2. THE FIELD EXPERIMENT

We use data from a field experiment conducted on WeChat Moments ads (see Figure 1). Owned by the technology conglomerate Tencent, WeChat is one of the world's largest mobile messaging applications, with over one billion monthly active users spending, on average, more than 90 minutes a day on the app. An important function of the app is WeChat Moments, which, like Facebook's newsfeeds, supports the posting of images and texts and the sharing of music, articles, and short videos. WeChat Moments introduced ads in spring 2015. Users can click, endorse (like), and comment on (but cannot share) ads in WeChat Moments (see Figure 1). The field experiment started in December 2015. As our experiment was conducted at the very early experimental stages of WeChat Moments ads, the targeting conditions were based solely on the users' age, gender, and city of residence.

The design of WeChat Moments' ads allows us to identify social influence by comparing the effectiveness of the social ads between the groups with and without social cues shown in ads. It is well known that identifying social influence is empirically challenging because of the potential confounding effects of homophily and particular external factors (Manski 1993, Aral et al. 2009). That is, individuals may share actions with their peers in the network, not because others have influenced them, but because they are just similar in terms of preferences. Accordingly, some researchers use field experiments to identify causal effects of social influence in ad engagements (Bakshy et al. 2012, Huang et al. 2020) and product adoptions(Aral and Walker 2011, 2014). In the next subsection, we provide the details of the randomized experiment.

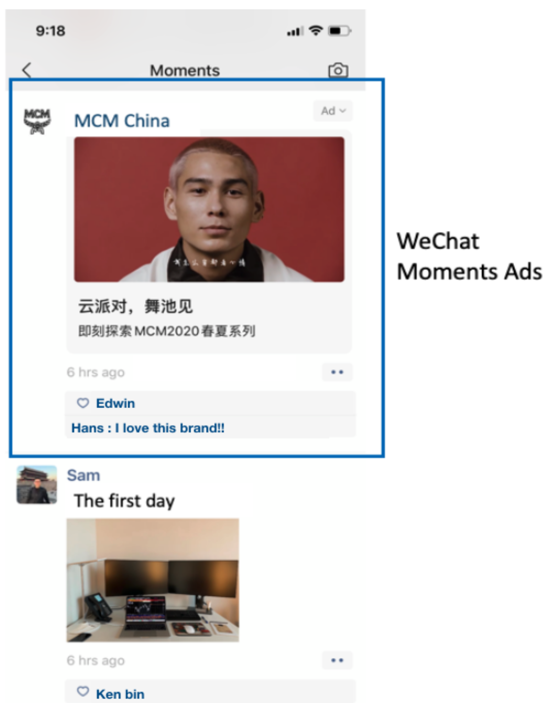

Figure 1: Example of WeChat Moments Ads

Note: This figure provides an example of WeChat Moments and Moments ads. WeChat Moments supports the posting of images and text and the sharing of music, articles, and short videos. Similar to Facebook ads, WeChat Moments ads appear on the timeline of Moments. Users can click, endorse (like), and comment on these ads.

2.1. Experimental Design

In our experiment, we randomized the presence of social cues shown in ads. When the users received an opportunity[1] to see a new ad, they were randomly assigned into three experimental groups, each with an 8% probability, or outside the experiment, with the remaining 76% probability. In the control group, the users did not see any social cues. In the two treatment groups, the users either saw a maximum of one like shown in ads (Treatment Group 1), or saw the organic number of likes (Treatment Group 2). Figure 2 shows an example of the experimental treatments. Note that this randomization was at the ad-user level, and occurred whenever the users received a new ad. Therefore, a user could be in a different experimental group or outside the experiment for different ads. It is unlikely that the users suspected they were in an experiment and communicated with others about their treatment. By randomly assigning social cues, our experiment eliminated the bias due to homophily and external confounding factors, such that the users were equally distributed and exposed to the external confounding factors across the control and treatment groups. Homophily suggests that similar people tend to associate together and become friends, and thus the more friends endorse an ad, the more likely a user is to respond to it (Mcpherson et al. 2001). The randomization of the presence of social cues in our experiment breaks the correlation between the number of social cues displayed in ads and the number of friend endorsers affiliated with ad viewers. Users saw only one ad at a time and therefore were unlikely to have been affected by the simultaneous assignment of different treatment conditions for different ads. Each ad remained in the users' newsfeed for a maximum of 48 hours. After 48 hours, the ad was removed and a new ad was probably received.

It is important to note that we manipulated only the *display* of social cues, not their

[1] Our experiment was conducted at the very early stage of WeChat Moments Ads. During our experiment, users were targeted solely based on their age, gender, and city for different ads.

occurrence, because the social cues needed to be organic and fake social cues are prohibited on WeChat. Comparing the control group (in which the users did not see any social cue) and Treatment Group 1 (in which the users were shown a maximum of one social cue) enables us to identify the marginal social influence of social ads or, in other words, the effect of one social cue on the user response. The variation in the number of social cues in Treatment Group 2 (in which the users were shown the organic number of social cues) further allows us to estimate the social influence exerted by more than one friend's likes. There are two types of social cues on WeChat Moments ads: likes and comments. Because comments can have both positive and negative sentiments, we focused on likes only and hid (controlled for) all of the comments on the ad interface throughout the experiment. Likes from different users were shown in identical format on Moments except for the user name, thus eliminating the heterogeneous effects attributable to the format in which the social cues were displayed. Almost all of the ads included in the experiment were new and distinctive, and the users were unlikely to have been exposed to the ads through any external sources outside WeChat.

Figure 2: Experimental Treatments

Note: This figure illustrates the experimental design. When the users received a new ad, they were randomly assigned into three groups: without any social cue (control group), with maximum one like displayed (Treatment Group 1) and with organic likes shown in the ad (Treatment Group 2). This randomization occurred whenever the users received a new ad during the experiment.

2.2. Data Collection

We collected four sets of data for this study. First, we recorded impression-level data on the number of social cues shown in the ads and the number of social cues hidden from the ad interface due to the experimental treatment. The data on the numbers of organic peer endorsements (displayed and hidden social cues) for ads allow us to examine the spontaneous behaviors among friends (homophily) and identify how the increasing numbers of social cues affect the levels of social influence (the effects of social cues). Second, the dependent variables for this study were the users' public and private responses to social ads. We collected impression-level data on the users' binary responses to an ad (i.e., whether they publicly liked and/or privately clicked on an ad), and their response times. We counted any click on a given ad as long as the users clicked the profile page, link to the landing page, or product photos. Third, to explore the heterogeneous effects of social cues and construct the control variables, we collected data on the characteristics of the ad viewers and their affiliated friends whose likes were shown in the ads. Specifically, we recorded their demographic information (e.g., age, gender, city) and historical behavior on WeChat.[2] Fourth, we also collected the names of the products and brands associated with the ads used in the experiments.

2.3. Descriptive Statistics

The experiment was conducted over a 21-day period starting on December 22, 2015, and involved 57,605,029 user-ad pairs, 37,985,501 distinct users, and 99 ads. A total of 19,198,166 user-ad pairs were randomly assigned to the control group with no displayed social cue. Another 19,201,745 user-ad pairs were randomly assigned to the treatment group with a

[2] *Affiliated friends* were those who had liked an ad and created social cues before the users saw the ad. Some of these social cues were hidden due to the experimental manipulations in the control group and Treatment Group 1.

maximum of one displayed like. The other 19,205,118 user-ad pairs were randomly assigned to the treatment group with the organic number of displayed likes. Our sample was representative of the population of social-ad users on WeChat (see Figure A1 in the Appendix). The users were randomly assigned to the experimental groups, with no statistically significant differences between the three experimental groups in terms of the users' age, gender, city, network degree (i.e., number of WeChat friends), and level of WeChat Moments activity (i.e., log-in days in Nov 2015) (see Table 1). This confirms the validity of the randomization procedure used in our experiment. On average, each user was exposed to fewer than two ads during the 21-day experiment.[3]

We dropped the user-ad pairs (3.46% of the sample) associated with 10 old ads[4] and with invalid data (i.e., incomplete user information). We also excluded a very small amount of data (0.16% of the sample) due to technical errors that caused an incorrect number of likes to be displayed on some user-ad pairs during the experiment. For example, we excluded the user-ad pairs assigned to Treatment Group 1 with more than one like or no likes shown in ads, and the user-ad pairs assigned to Treatment Group 2, in which the number of displayed likes did not match the number of organic likes. This ensured the integrity of our manipulation: no like was displayed to users in the control group, exactly one like was displayed to users in the treatment group 1, and the organic number of likes were correctly displayed to the users in the treatment group 2.

As previously mentioned, we did not generate fake likes but manipulated the display of real likes. As some ads had no organic likes[5], we were unable to display any real social cues

[3]In the early stages of WeChat Moments Ads, each ad only targeted a small group of users.

[4]The old ads were left over from the pre-experiment period. Users had already been exposed to these ads before the experiment started. The randomization in our experiment was at the ad-user level. In other words, the users for an ad were randomized into three experimental groups or outside the experiment. As a result, excluding any ads (with all the corresponding user-ad pairs) from the sample will not affect the integrity of the randomization in our experiment.

[5]Ninety percent of the user-ad pairs were not associated with any organic likes. In other words, no first-degree

on these ads as part of our manipulation. We excluded the user-ad pairs with zero organic likes in *both* the control and treatment groups, to guarantee that only one social cue could be displayed in the ads in Treatment Group 1, at least one social cue could be shown on ads in Treatment Group 2, and the users were equally distributed across the control and treatment groups after the filtering. This process resulted in a final sample of 89 ads, 5,571,225 user-ad pairs, and 4,884,154 distinct users across the three groups: 1,860,654 user-ad pairs in the control group, 1,873,434 user-ad pairs in Treatment Group 1, and 1,837,137 user-ad pairs in Treatment Group 2.

Table 1: Mean Comparisons between the Control and Treatment Groups

	$\lvert \#0 - \#1 \rvert$		$\lvert \#0 - \#2 \rvert$		$\lvert \#1 - \#2 \rvert$	
	t-statistic	SD	t-statistic	SD	t-statistic	SD
Age	0.0902	0.00197	0.332	0.00197	0.468	0.00197
Gender	1.927	0.000163	1.756	0.000163	-1.1707	0.000163
City	-1.382	0.000216	-0.281	0.000216	1.101	0.000216
#Login Days	0.580	0.000111	0.966	0.000111	0.386	0.000111
Log(Network Degree)	0.508	0.000674	1.146	0.000674	0.639	0.000674

Note. This table reports the results of the t-tests of the mean differences between the control group, Treatment Group 1, and Treatment Group 2. The variables reported are Age, Gender (female =0, male = 1), City (= 1,2, or 3, indicating the first, second, or third class of cities), #Login Days (number of days a user was logged into WeChat Moments), and Network Degree (number of WeChat friends). “0” represents the control group, “1” represents Treatment Group 1, and “2” represents Treatment Group 2.

3. EFFECTS OF SOCIAL CUES ON PUBLIC VERSUS PRIVATE RESPONSES

Our large-scale field experiment provides a unique dataset for investigating how social influence operates in social advertising through its impacts on the public (i.e., liking) and private (i.e., clicking) ad responses. In this section, we document the main effects and illustrate how displaying social cues can have different effects on the liking and clickthrough rates. We note that users’ responses to an ad mostly occurred at their first ad impression, and thus we focus

friends, who were also targeted by the ad, liked the ad before the users saw the ad at their fist ad impressions. WeChat Moments only displayed the social cues of first-degree friends.

on the first ad impressions to measure the ad responses in the main analysis. However, it is likely that some users logged onto WeChat multiple times, and thus may have responded to the ads after the first time they saw them. In Subsection 3.3, we extend our analysis to consider responses in the overall impressions of ads, and the conclusions are qualitatively the same.

3.1. Effects of One Social Cue

As a preliminary step, we first examine the average treatment effects of displaying one social cue on liking and clicking across all user-ad pairs by comparing the control group and Treatment Group 1. This illustrates our approach to identifying social influence. Recall that no social cues were displayed in the ads for the user-ad pairs in the control group, while exactly one social cue was shown in the ads in Treatment Group 1. It turns out that compared with the control group, the users who were shown one like were 0.98% ($p < 0.01$) more likely to like an ad and 0.96% ($p < 0.01$) more likely to click on an ad during their first ad impressions.

Therefore, on average, social influence can indeed have *directly* and significantly enhanced the ad responses in terms of both liking and clicking, thus lending support to the growing popularity of social advertising. That showing one like in an ad significantly increased the users' liking propensity implies that the social influences underlying liking responses are self-reinforcing. That is, one social cue can lead to even more social cues, thereby amplifying the effectiveness of social advertising. Note that the effects of one social cue are similar in the two dimensions of user response, namely the liking and clickthrough rates. Next, we show that this is no longer the case when we consider the scenarios in which more than one social cue can be displayed in an ad.

3.2. Effects of Multiple Social Cues

We now analyze the effects of multiple social cues on liking and clicking and examine the growth pattern of social influence. This amounts to comparing the control group without any social cues and Treatment Group 2 with the organic number of likes displayed. Figure 3 illustrates the model-free evidence for the different impacts of increasing the number of social cues on the liking and clickthrough rates. We calculate, for every number of social cues, the difference in the average response rates (i.e., liking or clickthrough rate) between the control group and Treatment Group 2. There is a clear increasing pattern in the liking rate as the number of social cues increases. This trend, however, is less pronounced for the clickthrough rate.

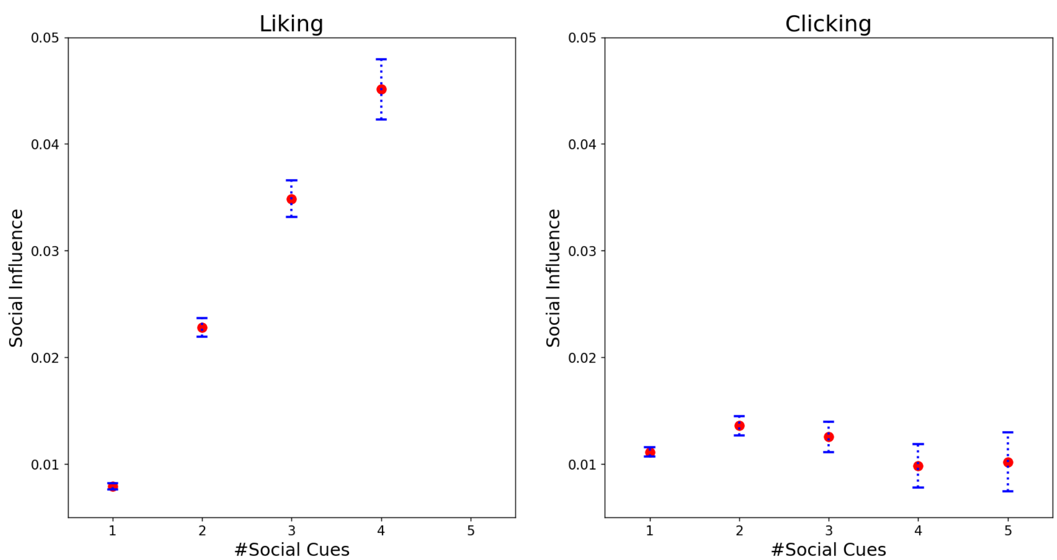

Figure 3: Model-Free Evidence of Social Influence on Liking and Clicking

Note: This figure plots the average effects of social influence, defined as the difference between Treatment Group 2 and the control group, on the liking and clickthrough rates, as a function of the number of social cues. The error bars correspond to the 95% confidence intervals of the mean differences between the control group and the treatment group.

Next we formally test this diverging pattern with regression analysis. Because the decisions on whether to like and click on an ad can be jointly made by the users, the two dependent variables may be correlated. To capture this dependency, we use the bivariate probit model (Greene 2003). User i can respond by expressing a like for ad j, denoted by the binary variable Y_{ij1}, and/or clicking on the ad, denoted by Y_{ij2}. Recall that upon receiving a new ad j, user i was randomly assigned to the control group with no social cue displayed, Treatment Group 1 with a maximum of one like displayed, and Treatment Group 2 with the organic likes displayed. Here, we focus on comparing the control group and Treatment Group 2. Thus, we let $S_{ij} = 1$ denote that the ad-user pair is in Treatment Group 2, and $S_{ij} = 0$ denote that it is in the control group. The bivariate probit model is a two-equation binary outcome model with correlated errors, specified as follows:

$$\begin{cases} Y^*_{ij1} = \alpha_1 + \beta_{11} S_{ij} + \beta_{21} N_{ij} + \pi_1 (N_{ij} \times S_{ij}) + C_{ij}\theta_{11} + (C_{ij} \times S_{ij})'\theta_{21} + \epsilon_{ij1}, \\ Y^*_{ij2} = \alpha_2 + \beta_{12} S_{ij} + \beta_{22} N_{ij} + \pi_2 (N_{ij} \times S_{ij}) + C_{ij}\theta_{21} + (C_{ij} \times S_{ij})'\theta_{22} + \epsilon_{ij2}, \end{cases} \tag{3.1}$$

where Y^*_{ij1} and Y^*_{ij2} are the latent utilities that user i can derive from liking and clicking on ad j. Because Y^*_{ij1} and Y^*_{ij2} are unobservable, we observe $Y_{ijr} = 1$ if and only if $Y^*_{ijr} > 0$, $\forall r = \{1, 2\}$. The variable N_{ij} indicates the number of organic social cues associated with user-ad pair ij. Note that in the treatment condition, N_{ij} is the number of likes *displayed* for user i, whereas in the control condition, it is the number of likes *hidden* from the user. Because we filtered the data based on the condition that there was at least one organic social cue in each observation (see Subsection 2.3), $N_{ij} \geq 1$. We use C_{ij} to represent a vector of control variables, including the users' gender, age group, number of friends, city of residence, and a set of ad dummies.[6] The error terms ϵ_{ij1} and ϵ_{ij2} are assumed to be distributed as bivariate

[6]The city of residence indicates whether the users live in a major, mid-size, or other cities. All of the controls are measured using the data from November 2015, one month before the start of the experiment.

normal with mean zero and unit variance. The correlation $\rho = \text{Corr}(\epsilon_{ij1}, \epsilon_{ij2})$ captures the possibility that some unobservable shocks may influence the like and click decisions together but are uncorrelated with the explanatory variables.

We estimate the bivariate probit regression model in Equation 3.2 as follows,

$$\begin{cases} Pr(Y_{ij1} = 1) = \Phi(\alpha_1 + \beta_{11} S_{ij} + \beta_{21} N_{ij} + \pi_1 (N_{ij} \times S_{ij}) + C_{ij}\theta_{11} + (C_{ij} \times S_{ij})'\theta_{21}), \\ Pr(Y_{ij2} = 1) = \Phi(\alpha_2 + \beta_{12} S_{ij} + \beta_{22} N_{ij} + \pi_2 (N_{ij} \times S_{ij}) + C_{ij}\theta_{21} + (C_{ij} \times S_{ij})'\theta_{22}), \end{cases} \tag{3.2}$$

where Φ denotes the cumulative normal distribution. Our main focus is on the coefficients π_1 and π_2 of the interaction terms, which capture the degree to which showing an additional social cue in an ad changes the degree of social influence conditional on the number of organic social cues. The coefficients β_{11} and β_{12} capture the raw impacts of displaying social cues in ads on users' responses, holding other variables constant. The coefficients β_{21} and β_{22} capture the tendency of users with more organic social cues to spontaneously respond to an ad in the *absence* of social influence ($S_{ij} = 0$). In other words, β_{21} and β_{22} indicate the degree to which the number of organic social cues predicts users' correlated latent preference to respond to an ad without social influence.

We control for the users' demographic variables (i.e., age, gender, city) and network characteristics (number of WeChat friends)[7], ads dummies, and their interactions with the treatment condition $C_{ij} \times S_{ij}$. Although we can evaluate how the number of social cues moderates social influence (the effects of social cues), we cannot make causal claims about this estimate because, in the experiment, we only randomized the *display* of social cues, rather than directly manipulating the number of the social cues shown in the ads. Moreover, some factors can lead to variations in the number of organic social cues. For example, users

[7]To reduce the data skewness, we log transform the number of friends in our analysis.

with more friends may have more likes on their ads, because WeChat users can only see their friends' likes. The ads with more ad viewers usually have more endorsements and organic social cues. Therefore, C_{ij} and $C_{ij} \times S_{ij}$ in the model can help control for these factors.

Table 2: Effects of Increasing Numbers of Social Cues on Social Influence

	Ad Like	Ad Click
Displaying Social Cues (β_1)	0.345***	0.570***
	(0.0184)	(0.0353)
#Social Cues (β_2)	0.0419***	0.0147***
	(0.00347)	(0.00400)
Displaying Social Cues * #Social Cues (π)	0.0329***	0.00251
	(0.00365)	(0.00585)
ρ	0.316	
Observations	3,697,658	
Log pseudolikelihood	-872194.12	

Note. This table reports the estimation results of the bivariate probit model in Equation 3.2, using all of the ad-user pairs in the control group and Treatment Group 2. The dependent variables are whether a user liked the ad and whether he or she clicked on the ad.

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

Table 2 summarizes the estimation results.[8] We observe that the coefficient for the interaction term is significantly positive for the liking rate ($\pi_1 = 0.0329$, $p < 0.01$), but insignificant for the clickthrough rate ($\pi_2 = 0.00251$, $p > 0.1$). It is important to note that in a nonlinear model with an interaction term, the cross-partial effect is not constant across the covariates and thus the coefficient for the interaction term is insufficient to determine the cross-partial effect (Ai and Norton 2003). Thus, we calculate the average effects of social influence separately for liking and clicking and for *different* numbers of social cues. That is, we use the coefficient estimates to derive the difference in predicted probabilities $\hat{Pr}(Y_r = 1|S = 1, N) - \hat{Pr}(Y_r = 1|S = 0, N)$, with other covariates set to mean, for every N and $r \in \{1, 2\}$. We then adopt the graphical approach (Greene 2010) to depict the patterns of social influence for the two types of response in Figure 4. The figure shows that the social influence on liking significantly increases with the number of social cues shown in the ads,

[8] We exclude 57 user-ad pairs with extreme values of organic likes (> 70), and 76 user-ad pairs of the ads with less than 30 observations.

validating the significant positive interaction effect for the liking response. In contrast, the social influence for the clicking response remains almost flat as the number of social cues increases, which is consistent with the insignificant coefficient for the interaction term in Table 2. Thus, these results confirm that the number of social cues has a positive impact on the degree of social influence on liking but not so much on clicking.

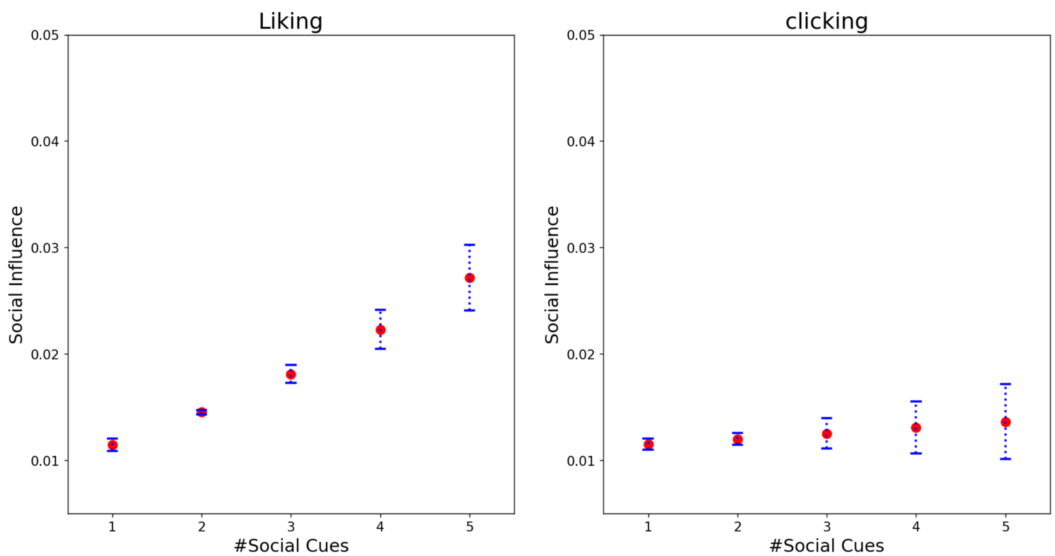

Figure 4: Model Estimates of Social Influence on Liking and Clicking

Note: This figure plots the model estimates of the effects of social influence on liking and clicking as a function of different numbers of social cues. The error bars correspond to the 95% confidence intervals of the estimated social influence.

Our model also allows us to assess whether, in the absence of social influence, users respond to the increasing numbers of social cues spontaneously simply based on their latent preferences. Homophily predicts that the more friends endorse an ad, the more likely a user is to respond to it (Mcpherson et al. 2001). Consistent with this prediction, the coefficients for the main effects for both liking ($\beta_{21} = 0.0419, p < 0.01$) and clicking ($\beta_{22} = 0.0147, p < 0.01$) are significantly positive, indicating that the number of friends' likes is correlated with the tendency to like and click on an ad even in the absence of social influence, $S_{ij} = 0$. We then estimated the average marginal effects of the numbers of likes on users' liking and clicking when no likes are shown in the ads. The results suggest that ads with one more organic like are, on average, associated with a 0.0136 ($p < 0.01$) increase in the liking rate and 0.00125 ($p < 0.05$) increase in the clickthrough rate. The implication is that even without showing social cues, targeting users who have more friends that have endorsed an ad will improve both the liking rate and the clickthrough rate of the ads. However, the predictive power of the number of friends' likes for an ad is much larger for the liking rate than the clickthrough rate.

3.3. Robustness Check: Effects of Social Cues in Multiple Ad Impressions

In the main analysis, we focus on the effects of social cues on users' responses in their first ad impressions, which is a reasonable measure of advertising effectiveness. However, users may go beyond the first ad impression and respond to an ad even after they have seen it multiple times. Therefore, we now examine the alternative measure of advertising effectiveness based on the *entire* span of users' ad impressions. Specifically, we construct the new dependent variables by counting any clicks or likes on a given ad at any time the ad was displayed on the users' WeChat Moments. As we describe in Subsection 2.1, the ads remained on WeChat Moments for no more than 48 hours during the experiment.

We estimated the same bivariate probit model in Equation 3.2 with the new dependent variables. Table A1 in the Appendix reports the estimated coefficients. Figure A2 in the Appendix illustrates the corresponding patterns of social influence for liking and clicking based on the estimates. Although we observe larger effects of the social cues on both liking and clicking compared with those when focusing on the first ad impressions, Figures 4 and A2 show very similar growth patterns of social influence for the two types of ad response.

Hence, our findings that social influence (the effects of social cues) is significantly positive for both the liking and clicking responses and that only the liking response increases with the numbers of social cues displayed in the ads are robust with respect to the alternative ways to measure the ad responses.

4. MECHANISM: INFORMATIONAL VERSUS NORMATIVE SOCIAL INFLUENCE

The results in the previous section show that the increasing numbers of social cues can lead to quite different dynamics in the two different dimensions of user responses, which suggests that they are likely driven by different motives of the users. In this section, we explore this possibility by exploiting the notable difference between a liking response and a clicking response, namely that liking is publicly observable to others in the network, whereas clicking is privately known to the users. When users choose to *like* an ad, they understand that this response will be observed by their friends, and may then feel the urge to conform to the positive expectations of others (Deutsch and Gerard 1955, Asch 1956, Mouton et al. 1956, Argyle 1957, Levy 1960, Insko et al. 1983, Insko et al. 1985, Campbell and Fairey 1989). If a user sees that more of her friends have liked the ad, then her incentive to conform to this behavior may be reinforced. This mechanism is commonly known as normative social influence. In contrast, when deciding whether to *click* on an ad, users do not need to worry about the effects their actions may have on the beliefs of others, who are unable to observe the click behavior. Hence, the decision to click is mainly driven by the user's private interest, the desire to be accurate. That is, the user is more likely to be subject to informational social influence.

To isolate these two different mechanisms, we first identify scenarios with varying degrees of either informational social influence or normative social influence, and then evaluate their impact on the two dimensions of advertising effectiveness. We expect that when informational social influence becomes more important, the impact of social cues on the clickthrough rate will be stronger than when it is expected to play a weaker role. However, we expect that this difference will be less pronounced for the liking rate because it is more subject to normative social influence. In a similar vein, we expect that when normative social influence becomes more prominent, the impact of social cues on the liking rate will be stronger than when normative influence is weaker. However, we expect that this difference will be less pronounced for the clickthrough rate, which is more subject to informational social influence. Next, we test each of these two hypotheses.

4.1. Well-Known versus Lesser-Known Brands

As some advertised products or brands are more reputable or familiar than others, consumers do not need to gather as much information on the brands or their products. Hence, informational social influence is less important for these brands. Following this argument, we use the list of the top 100 global brands published by Interbrand in 2015 to determine whether an advertiser is a well-known brand (Lovett et al. 2014). We call the brands belonging to the Interbrand list "well-known" brands, whereas the remaining brands are referred to as "lesser-known" brands. Examples of well-known brands in our sample include BMW, Pepsi, and Prada, while the lesser-known brands are mostly local brands or relatively new brands.

We estimate the same regression models in Equation 3.2 separately for the well-known and lesser-known brands. Table A3 in the Appendix reports the coefficients. Again, because the interaction effects are not constant across the covariates, we use the graphical representation in Figure 5 to illustrate the impacts of displaying social cues on both public (i.e., liking) and private responses (i.e., clicking) as a function of the increasing numbers of social cues. First, with respect to the private response of clicking (see the lower panels of 5), as the well-known brands are associated with more social cues, the impact of social influence on the clickthrough rate is non-increasing and it even exhibits a slight decreasing pattern, as shown in the lower left-hand panel of Figure 5. This sharply contrasts with the increasing pattern observed for the lesser-known brands (see the lower right-hand panel of Figure 5). This result is consistent with the prediction that when the desire for information is greater (i.e., for lesser-known brands), users are more likely to incorporate information from others (i.e., social cues). The more social cues in an ad for a lesser-known brand, the more the users can make positive inferences about the brand. Thus, we observe a stronger increasing effect of the number of social cues on the clickthrough rates for lesser-known brands.

However, this pattern may disappear when normative social influence plays a stronger role, such as in the case of the public response of liking. Indeed, we observe a stronger increasing impact of more social cues on the liking rate for both well-known and lesser-known brands ($p < 0.01$), as shown in the upper panels of Figure 5. If the effect of social cues on the liking rate were purely driven by informational social influence, we would observe a pattern similar to the case of the clickthrough rate. Figure 5 also shows that although both the well-known and lesser-known brands exhibit the same trend as the numbers of social cues increases, the social influence of the well-known brands is generally greater than that of the lesser-known brands. This is likely because normative social influence is stronger for well-known brands than for lesser-known brands, because users expect that more of the friends will like the well-known brands and thus tend to conform with this expectation. If clicking were mainly driven by normative social influence, we would expect a stronger increasing pattern for the well-known brands' clickthrough rates. However, this is not supported by the data pattern,

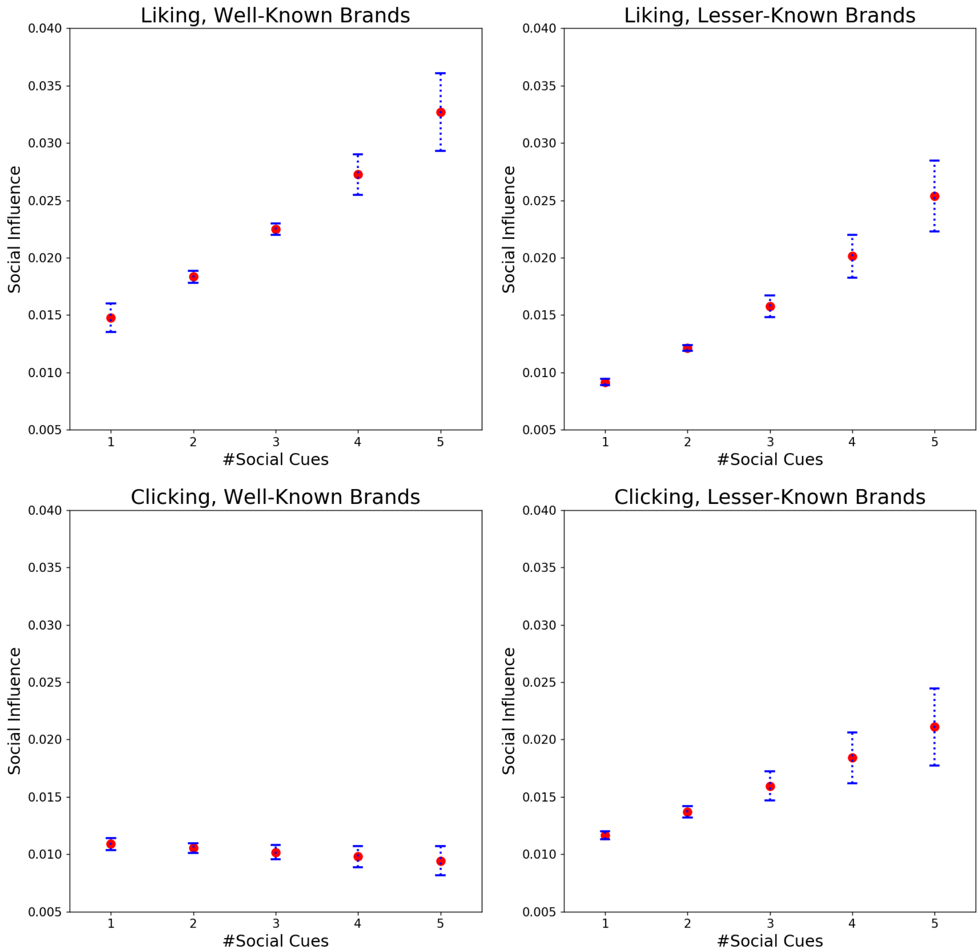

Figure 5: Model Estimates of Social Influence on Liking and Clicking (Well-Known versus Lesser-Known Brands)

Note: This figure plots the model estimates of the effects of social influence on liking and clicking as a function of different numbers of social cues, separately for well-known and lesser-known brands. The error bars correspond to the 95% confidence intervals of the estimated social influence.

which again supports the proposition that the clicking response is more likely subject to informational social influence.

Our results for the clicking response echo the findings of Tucker and Zhang (2011), which suggest that popularity information benefits narrow-appeal products more than broad-appeal products. Their mechanism builds on information learning. That is, for the same level of popularity, a narrow-appeal product must generate more good quality signals than a broad-appeal product, because the narrow-appeal product matches the tastes of fewer consumers. Similarly, we find that less familiar or lesser-known brands benefit more from social cues than well-known brands in terms of the clickthrough rates. Although the driver is also an informational one, we offer a quite different but complementing perspective here. First, the advertised products in our study are mostly mainstream products with a broad appeal. Our brand categorization is based more on the level of familiarity or quality uncertainty than on the breadth of appeal. Thus, our results focus on the heterogeneity in the value of information across brands. Second, our results reveal the different patterns for the clicking and liking responses and suggest that normative influence plays a role in the difference. This is not identified in Tucker and Zhang (2011).

4.2. More versus Less Socially Engaged Users

Turning to the case in which normative social influence can be varied, we note that this type of social influence tends to operate when individuals wish to build and maintain social relationships, and that the strength of these desires may vary across individuals. Hence, we postulate that users who actively engage in social interactions on WeChat Moments are more likely to be those who use the platform to build social relationships and seek social approvals from others, and thus are more likely to act to meet the expectations of others or, more generally, to maximize their social outcomes. Therefore, more socially engaged users are more susceptive to normative social influence.

A reasonable measure of a user's social interaction is his or her frequency of endorsing (i.e., liking) others' posts on WeChat Moments. We thus construct a measure of social engagement based on users' total likes of their friends' posts divided by their total number of friends in the month before the experiment (November 2015). We then conduct a mean split to categorize the users as either more socially engaged users or less socially engaged users. We run the same regression models as in Equation 3.2 separately on these two groups of users. Table A6 in the Appendix reports the coefficients. Again, we focus on the results in Figure 6 to illustrate the effects of increasing social cues on both the liking and clickthrough rates for these two groups of users.

First, we observe that the effect of the social cues on the liking rate significantly increases with the number of social cues displayed in ads for both groups of users, as shown in the upper panels of Figure 6. However, social cues have a significantly greater impact and there is a stronger increasing pattern of social influence on liking for the more socially engaged users. This result supports our prediction that normative social influence operates among the more socially engaged users, thus leading to the greater impact of increasing social cues on liking.

However, this difference between the more and less socially engaged users may disappear when informational social influence plays a significant role in determining private responses (i.e., clicking on an ad). Indeed, we find a similar insignificant increasing pattern in clicking for both groups of users (see the lower panels of Figure 6), which suggests that informational social influence has a strong effect on the clickthrough rate. If the effect of social cues on the clickthrough rate were purely driven by normative social influence, we would observe a similar pattern to the case of the liking rate (the upper panels of Figure 6). However, this is not supported by the data pattern that we observe, which again accords with the proposition that clicking is more likely to be subject to informational social influence.

One limitation of using the frequency of likes on friends' posts to control for the normative social influence is that the correlation between liking an ad and liking friends' posts is

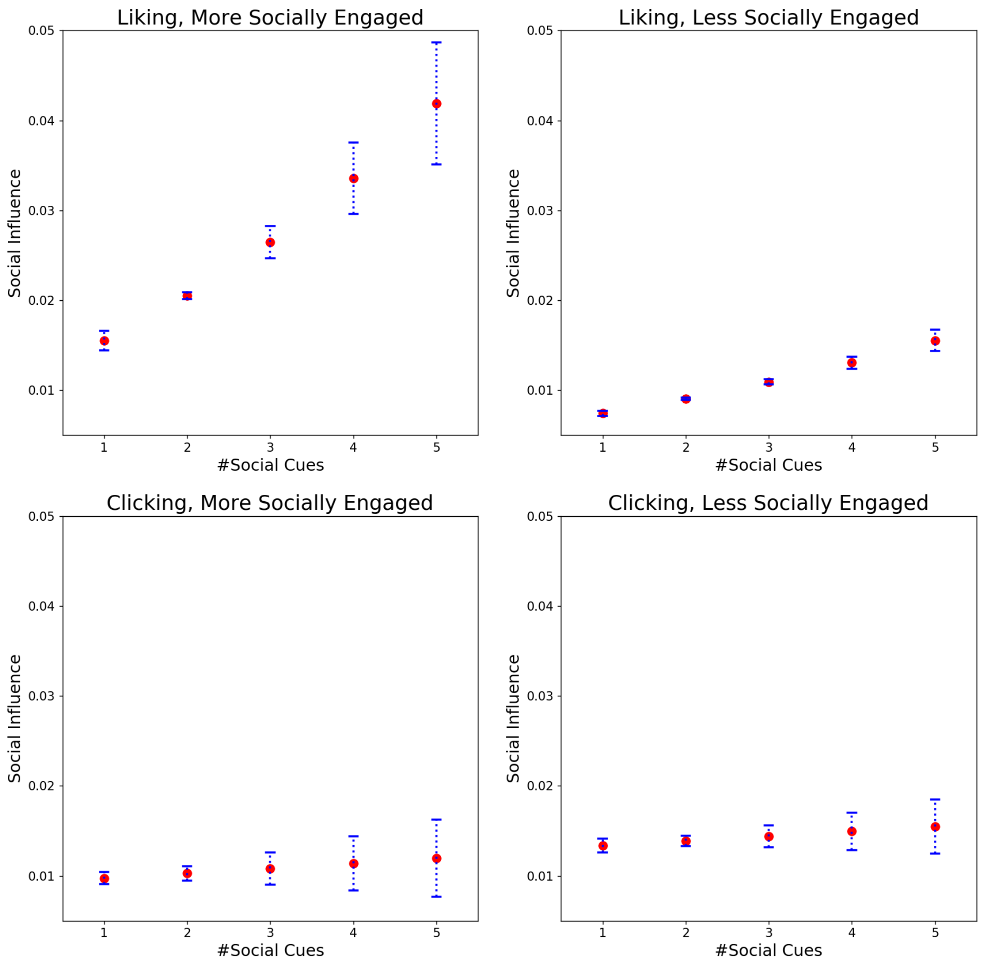

Figure 6: Model Estimates of Social Influence on Liking and Clicking (More versus Less Socially Engaged Users)

Note: This figure plots the model estimates of the effects of social influence on liking and clicking as a function of different numbers of social cues, separately for the more and the less socially engaged users. Social engagement is measured based on the numbers of likes on friends' posts on WeChat Moments. The error bars correspond to the 95% confidence intervals of the estimated social influence.

potentially driven by the similarity of the two actions, rather than social influence. Therefore, we explore an alternative but related measure, the frequency of commenting on others' posts, to capture the degree of social engagement. We first count each user's total number of comments on friends' posts divided by the number of friends in the month prior to the experiment, and then group the users into more or less socially engaged users using a mean split. We observe a very similar data pattern to that found when using the number of likes to measure social engagement, as shown in Figure 7. Specifically, the effect of increasing social cues on liking is significantly stronger for more socially engaged users who comment more frequently on their friends' posts, while there is a similar non-increasing impact of increasing social cues on clicking for both groups of users. Therefore, our conclusion remains unchanged when we use this alternative measure of social engagement to capture the variation in normative social influence.

4.3. Robustness Checks: Alternative Measures

We further consider alternative measures of well-known brands and social engagement to ensure the robustness of our findings. First, we used Interbrand's list of top 100 global brands in 2016 instead of that in 2015 to measure well-known brands, as our experiment was conducted from December 2015 to January 2016. The difference between the two lists is that the 2016 list includes Dior that is excluded from the 2015 list. We apply the same analysis using this new measure of well-known brands and find very similar results. We again estimate the same regression models in Equation 3.2 separately for the re-defined well-known and lesser-known brands. Figure A3 in the Appendix shows that for the well-known brands, the impact of increasing social cues on the clickthrough rate is non-increasing and even exhibits a slightly decreasing pattern, whereas the effect of increasing social cues on clicking on an ad is increasing for lesser-known brands. However, the lesser-known and well-known brands both exhibit an increasing trend for the effects of social influence on liking, although the effect is stronger for the well-known brands than for the lesser-known brands. These results corroborate our main findings in Subsection 4.1 and our theoretical prediction.

Second, we consider the frequency of sharing (e.g., posting photos, articles, and texts) on

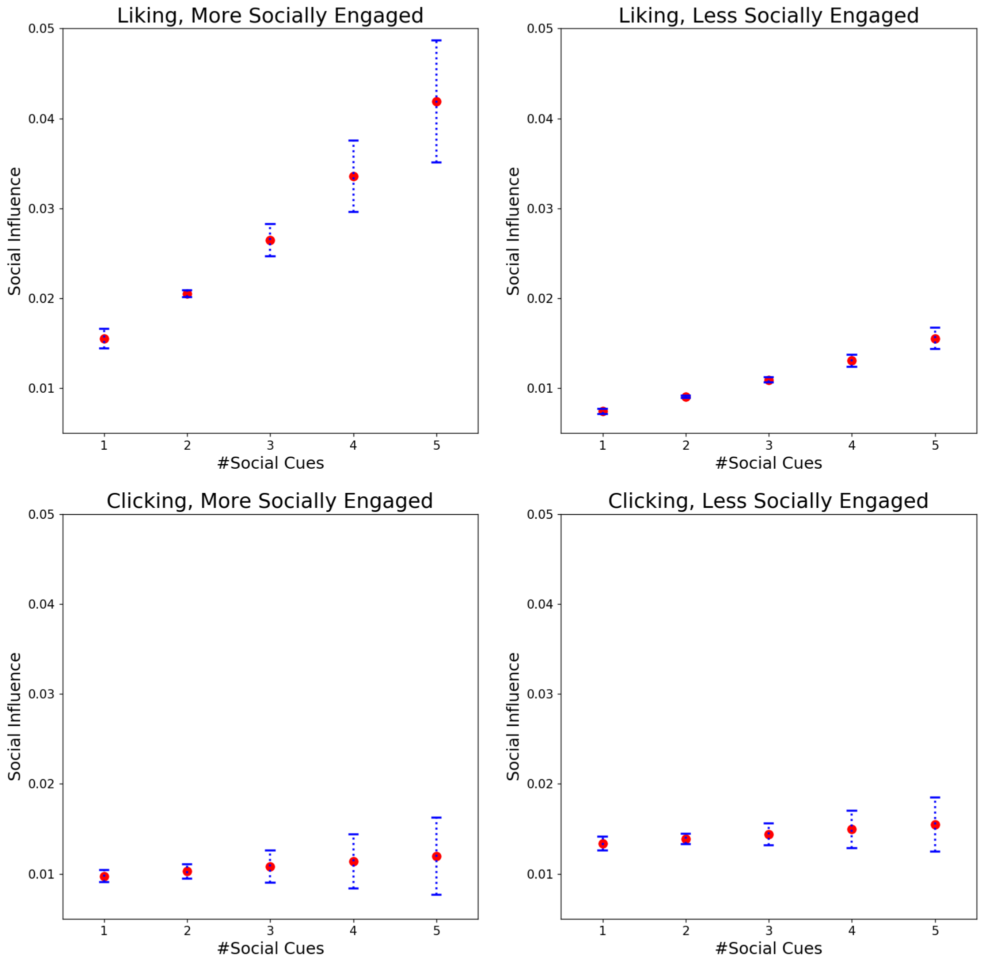

Figure 7: Model Estimates of Social Influence on Liking and Clicking (More versus Less Socially Engaged Users, Defined by Commenting Activity)

Note: This figure plots the model estimates of the effects of social influence on liking and clicking as a function of different numbers of social cues, separately for the more and less socially engaged users. Social engagement is defined based on users' frequency of commenting on others' posts. The liking and clicking responses are defined over users' first ad impressions. The error bars correspond to the 95% confidence intervals of the estimated social influence.

WeChat Moments as another alternative measure of social engagement. Here, the idea is that individuals who more actively update their status and share personal moments with their friends through broadcasting on social media exhibit stronger incentives to build positive relationships with other people, which leads to conformity (DeWall and Bushman 2011).

As a result, users who post more often on WeChat Moments are more likely to be subject to normative social influence. We categorize the users into two groups based on whether their total number of posts on WeChat Moments in November 2015, the month before the experiment, is above (more socially engaged) or below the mean (less socially engaged). As shown in Figure A4 in the Appendix, the pattern of the effects of increasing social cues is very similar to those in Figure 6 and Figure 7. Specifically, more socially engaged users, who posts more frequently on WeChat Moments, are more affected by the increasing number of social cues, in terms of their liking propensity, than the less socially engaged users. In contrast, we observe a similar non-increasing pattern of social influence in terms of the clickthrough rates for all users. These findings are consistent with the results of our main analysis in Subsection 4.2.

5. DISCUSSION AND IMPLICATION

Our results provide evidence for the co-existence of both informational and normative social influence in social advertising. We show that the two types of social influence disproportionately affect individuals' liking and clicking responses, such that informational social influence has a greater impact on the clickthrough rate, whereas normative social influence has a greater effect on the liking rate. These findings provide a coherent explanation for the diverging pattern of increasing social cues on the liking and clickthrough rates documented in Subsection 3.2, that is, more social cues lead to higher liking rates but do not increase the clickthrough rate. On the one hand, the liking response is more likely driven by normative social influence, which results in the snowballing effect whereby the more friends like an ad, the more social influence they exert, and thus the more other users tend to like it as well.

On the other hand, if the clicking response is driven more by informational social influence, then individuals will try to infer the quality of an advertiser from the likes (social cues) it receives. They may reason that the social cues are driven by both informational and normative social influence. The more weight of normative social influence in generating the social cues is, the less informative these social cues become. Consequently, in this case, more social cues will not convince the users that it is worthwhile to click on the ad.

This mechanism has rich implications for the social advertising policies of advertisers and social media platforms. First, it is important that advertisers have a clear objective when deciding to advertise on social media. If an advertiser has a performance-focused objective, such as driving product sales, then the clickthrough rate will be a more important metric, but it may not be effective if normative social influence dominates the social network and social-cue generating process. However, if the objective is to build brand awareness and/or image, then advertisers should focus on the liking rate to exploit the snowballing effect of social cues in generating more likes. Choosing a suitable social media platform and targeting consumer segments in which normative social influence prevails can effectively fulfill this goal.

Second, social media platforms clearly have flexibility of determining how many social cues users can observe for an ad. Although full disclosure is generally optimal for the liking rate, it can be suboptimal for the clickthrough rate. In some cases, a social media platform may only need to display one social cue to maximally impact the clickthrough rate. Hence, social media platforms need to evaluate the nature of the social influence at play in social advertising and carefully design the appropriate policies in relation to social cues.

Lastly, our tests illustrate two approaches for social media platforms to manage the effectiveness of social advertising. One approach is to design a brand-specific policy because brands may be heterogeneous in terms of their familiarity to consumers, and thus vary in terms of their degree of informational social influence. For more familiar, reputable brands, it is more desirable to exploit normative social influence in social advertising, whereas for less familiar brands, social advertising is more effective when based on informational social influence. The second approach is to design a targeting based, consumer-specific policy that builds on users' historical behavior. This requires accurate assessments of each user's susceptibility to informational and normative social influence. For example, users' frequency of endorsing and commenting on friends' posts, as well as sharing on social media can be the indicators of the degree of susceptibility to normative social influence. Clearly, the latter approach has to be carefully evaluated in relation to users' privacy concerns.

6. CONCLUDING REMARKS

Although social medial have become an increasingly popular advertising channel among firms, few studies have empirically evaluated the effectiveness of social media ads, partly because distinguishing social influence from homophily is empirically challenging. In this study, we use a large-scale experiment to identify the impact of social influence on social advertising. Unlike prior studies, we distinguish two dimensions of effectiveness, namely the public responses (i.e.,liking rate) and private responses (i.e., clickthrough rate). We argue that firms with different advertising objectives should weigh these dimensions differently and show that they are subject to different dynamics of social influence. Although the first social cue always enhances the liking and clickthrough rates, more social cues can only increase the liking rate but have little effect on the clickthrough rate. Our results highlight the coexistence of two distinct mechanisms of social influence, namely normative social influence and informational social influence. Although they can simultaneously affect users' liking and clicking responses, normative (informational) influence has a greater effect on the liking (clicking) response, because this response is publicly observable (privately known).

The effectiveness of more social cues on the private responses (i.e., clicking) may depend on the relative weight of normative social influence in generating these social cues. These findings have rich implications for advertisers and social media platforms in managing social advertising.

Although our study delivers the important message that the effectiveness of social advertising should be evaluated in multiple dimensions and provides exploratory evidence that both normative and informational social influence can affect the effectiveness of social media ads in different ways, questions remain in terms of how these two types of influence affect social advertising and the factors that can moderate their impacts. These questions are clearly worthy of further investigation, both empirically and theoretically. Perhaps the more challenging question is how to measure the relative effects of normative and informational social influence. The answer to this question would shed further light on how to manage social advertising more efficiently.

The authors gratefully acknowledge the existence of the Journal of Irreproducible Results and the support of the Society for the Preservation of Inane Research.

References

Chunrong Ai and Edward C. Norton. Interaction terms in logit and probit models. *Economics Letters*, 80(1):123–129, 2003.

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

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

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

Michael Argyle. Social pressure in public and private situations. *The Journal of Abnormal and Social Psychology*, 54(2):172, 1957.

Johan Arndt. Role of product-related conversations in the diffusion of a new product. *Journal of marketing Research*, 4(3):291–295, 1967.

Solomon E Asch. Opinions and social pressure. *Scientific American*, 193(5):31–35, 1955.

Solomon E Asch. Studies of independence and conformity: I. a minority of one against a unanimous majority. *Psychological monographs: General and applied*, 70(9):1, 1956.

Eytan Bakshy, Menlo Park, Cameron Marlow, and S I Feb. The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web*, pages 519–528. ACM, 2012.

Abhijit V Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3): 797–817, 1992.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992.

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

Jennifer D Campbell and Patricia J Fairey. Informational and normative routes to conformity: The effect of faction size as a function of norm extremity and attention to the stimulus. *Journal of Personality and Social Psychology*, 57(3):457, 1989.

Judith A Chevalier and Dina Mayzlin. The effect of word of mouth on sales: Online book reviews. *Journal of marketing research*, 43(3):345–354, 2006.

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

Robert B Cialdini and Noah J Goldstein. Social Influence: Compliance and Conformity. *Annual Review of Psychology*, 55(1974):591–621, 2004.

Robert B Cialdini and Melanie R Trost. Social influence: Social norms, conformity and compliance. 1998.

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

C Nathan DeWall and Brad J Bushman. Social acceptance and rejection: The sweet and the bitter. *Current Directions in Psychological Science*, 20(4):256–260, 2011.

David Godes and Dina Mayzlin. Using online conversations to study word-of-mouth communication. *Marketing science*, 23(4):545–560, 2004.

William Greene. Testing hypotheses about interaction terms in nonlinear models. *Economics Letters*, 107(2):291–296, 2010. ISSN 01651765.

William H Greene. *Econometric Analysis*, volume 97. 2003.

Shan Huang, Sinan Aral, Yu Jeffrey Hu, and Erik Brynjolfsson. Social Advertising Effectiveness across Products: A Large-scale Randomized Filed Experiment. *Marketing Science*, forthcoming, 2020.

Chester A Insko, Sarah Drenan, Michael R Solomon, Richard Smith, and Terry J Wade. Conformity as a function of the consistency of positive self-evaluation with being liked and being right. *Journal of Experimental Social Psychology*, 19(4):341–358, 1983.

Chester A Insko, Richard H Smith, Mark D Alicke, Joel Wade, and Sylvester Taylor. Conformity and group size: The concern with being right and the concern with being liked. *Personality and Social Psychology Bulletin*, 11(1):41–50, 1985.

Leo Levy. Studies in conformity behavior: a methodological note. *The Journal of Psychology*, 50(1): 39–41, 1960.

Mitchell Lovett, Renana Peres, and Ron Shachar. A Data Set of Brands and Their Characteristics. *Marketing Science*, 33(4):609–617, 2014.

Charles F. Manski. Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3):531–542, 1993.

Miller Mcpherson, Lynn Smith-lovin, and James M Cook. Birds of a Feather : Homophily in Social Networks. *Annual Review of Sociology*, 27:415–444, 2001.

Jane S Mouton, Robert R Blake, and Joseph A Olmstead. The relationship between frequency of yielding and the disclosure of personal identity. *Journal of Personality*, 1956.

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

Catherine Tucker and Juanjuan Zhang. How does popularity information affect choices? a field experiment. *Management Science*, 57(5):828–842, 2011.

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

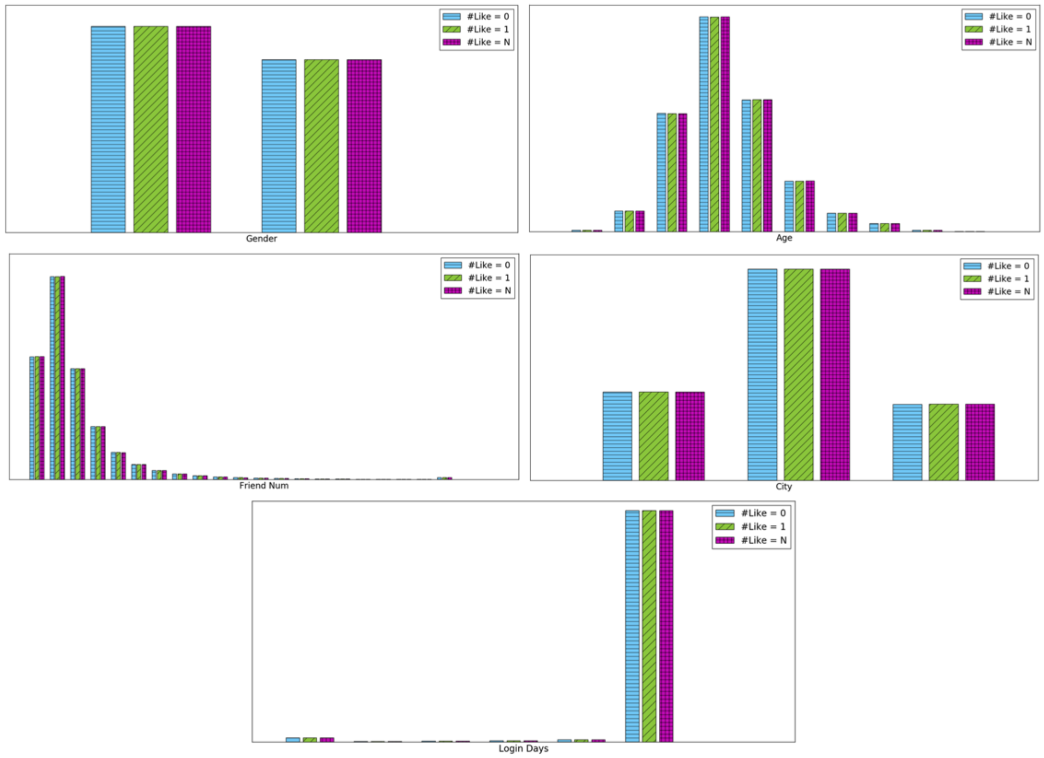

Figure A1: Distributions of User Characteristics Across Control and Treatment Groups

Note: This figure shows the distributions of user characteristics (i.e., users age, gender, network degree (WeChat friend number), and the number of login days in November 2015) across the control group and the two treatment groups. Due to the Non-Disclosure Agreement(NDA), we cannot reveal coordinates on x-axis and y-axis.)

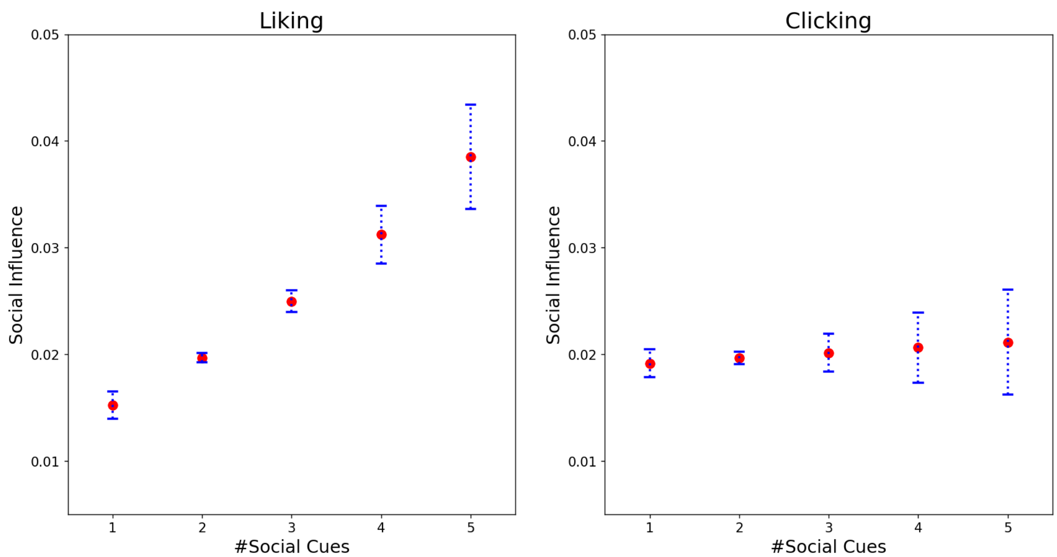

Figure A2: Model Estimates of Social Influence on Liking and Clicking (Multiple Ad Impressions)

Note: This figure plots the model estimates of social influence for liking and clicking as a function of different numbers of social cues. The liking and clicking responses are defined over users' entire ad impressions. The error bars correspond to 95% confidence intervals of the estimated social influence.

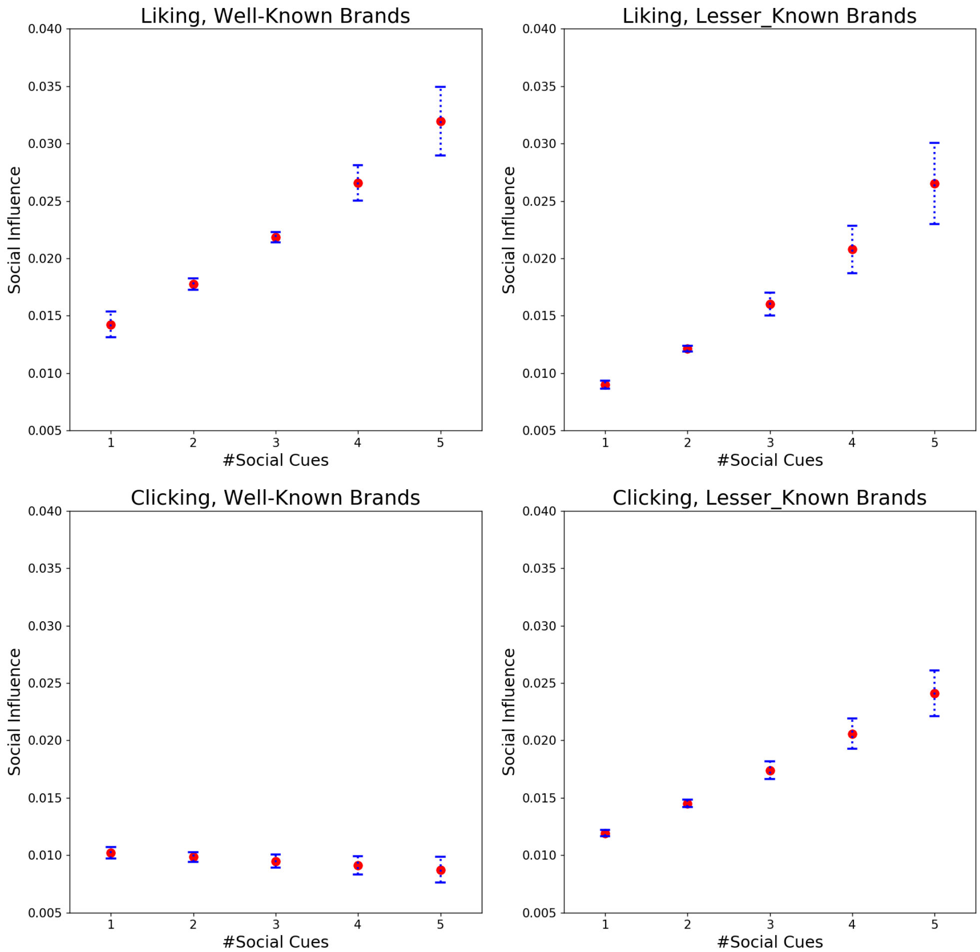

Figure A3: Model Estimates of Social Influence on Liking and Clicking (Well-Known versus Lesser-Known Brands, Defined by Interbrand 2016)

Note: This figure plots the model estimates of the effects of social influence on liking and clicking as a function of different numbers of social cues, separately for well-known and lesser-known brands. The categorization of brands is based on Interbrand's best global brands in 2016. The liking and clicking responses are defined over users' first ad impressions. The error bars correspond to the 95% confidence intervals of the estimated social influence.

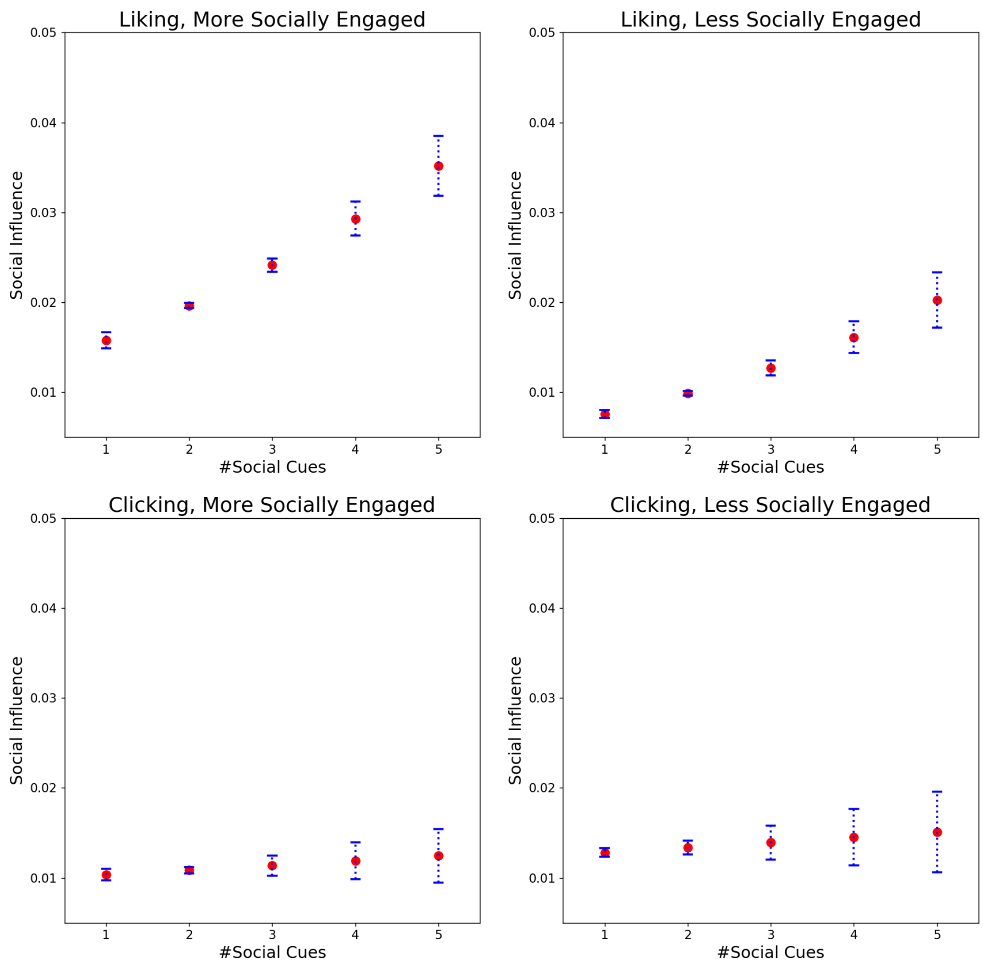

Figure A4: Model Estimates of Social Influence on Liking and Clicking (More versus Less Socially Engaged Users, Defined by Posting Activity)

Note: This figure plots the model estimates of the effects of social influence on liking and clicking as a function of different numbers of social cues, separately for more and less socially engaged users. Social engagement is defined based on users' posting activity. The liking and clicking responses are defined over users' first ad impressions. The error bars correspond to the 95% confidence intervals of the estimated social influence.

Table A1: Effects of Increasing Numbers of Social Cues on Social Influence (Multiple Ad Impressions)

	Ad Like	Ad Click
Displaying Social Cues (β_1)	0.422***	0.514***
	(0.0242)	(0.0279)
#Social Cues (β_2)	0.0434***	0.0177***
	(0.00342)	(0.00356)
Displaying Social Cues * #Social Cues (π)	0.0323***	-0.00197
	(0.00417)	(0.00614)
ρ	0.303	
Observations	3,697,658	
Log pseudolikelihood	-1200536.3	

Note. This table reports the estimation results of the bivariate probit model in Equation 3.2, using all ad-user pairs in the control group and Treatment Group 2. The dependent variables are whether a user liked the ad and whether she clicked on the ad at any ad impressions. $*p < 0.10$;$**p < 0.05$;$***p < 0.01$.

Table A2: Effects of Increasing Number of Social Cues on Social Influence for Well-Known and Lesser-Known Brands

	Ad Like		Ad Click	
	Well-known Brand	Lesser-known Brand	Well-known Brand	Lesser-known Brand
Displaying Social Cues (β_1)	0.111***	0.412***	0.553***	0.569***
	(0.0509)	(0.0330)	(0.0630)	(0.0546)
#Social Cues (β_2)	0.0400***	0.0510***	0.0103*	0.0269***
	(0.00398)	(0.00425)	(0.0256)	(0.0333)
Displaying Social Cues $\times$ Social Cues (π)	0.0310***	0.0364***	-0.0123**	0.0173***
	(0.00440)	(0.00545)	(0.00413)	(0.00371)
Observations	1,805,757	1,891,901	1,805,757	1,891,901
Log pseudolikelihood	-389395.37	-478896.79	-389395.37	-478896.79

Note. This table reports the estimation results of the bivariate probit model in Equation 3.2, separately for well-known and lesser-known brands, using all ad-user pairs in the control group and Treatment Group 2. The dependent variables are whether a user liked the ad and whether she clicked on the ad at users' first ad impressions.. $*p < 0.10$;$**p < 0.05$;$***p < 0.01$.

Table A3: Effects of Increasing Numbers of Social Cues on Social Influence for Well-Known and Lesser-Known Brands (2016)

	Ad Like		Ad Click	
	Well-known Brand	Lesser-known Brand	Well-known Brand	Lesser-known Brand
Displaying Social Cues (β_1)	0.0924**	0.434***	0.565***	0.564***
	(0.0476)	(0.0302)	(0.0628)	(0.0560)
#Social Cues (β_2)	0.0402***	0.0555***	0.00979*	0.0294***
	(0.00362)	(0.00521)	(0.00230)	(0.00328)
Displaying Social Cues × Social Cues (π)	0.0308***	0.0391***	-0.0128***	0.0204***
	(0.00389)	(0.00566)	(0.00352)	(0.00199)
Observations	1,974,251	1,723,407	1,974,251	1,723,407
Log pseudolikelihood	-418517.03	-449563.37	-418517.03	-449563.37

Note. This table reports the estimation results of the bivariate probit model in Equation 3.2, separately for well-known and lesser-known brands, using all ad-user pairs in the control group and Treatment Group 2. The brands are categorized based on the Interbrand's list of best global brands in 2016. The dependent variables are whether a user liked the ad and whether she clicked on the ad at users' first ad impressions. $*p < 0.10$;$**p < 0.05$;$***p < 0.01$.

Table A4: Effects of Increasing Numbers of Social Cues on Social Influence for More or Less Socially Engaged Users (#Likes)

	Ad Like		Ad Click	
	More Socially Engaged	Less Socially Engaged	More Socially Engaged	Less Socially Engaged
Displaying Social Cues (β_1)	0.393***	0.329***	0.619***	0.548***
	(0.0283)	(0.0360)	(0.0437)	(0.0383)
#Social Cues (β_2)	0.0509***	0.0346***	0.0164***	0.0134***
	(0.00449)	(0.00276)	(0.0435)	(0.0389)
Displaying Social Cues × Social Cues (π)	0.0391***	0.0275***	0.0234	0.0219
	(0.00649)	(0.00224)	(0.00679)	(0.00532)
Controls	Yes	Yes	Yes	Yes
Observations	1,850,953	1,846,705	1,850,953	1,846,705
Log pseudolikelihood	-494282.9	-365620.84	-494282.9	-365620.84

Note. This table reports the estimation results of the bivariate probit model in Equation 3.2, separately for more and less socially engaged users, using all ad-user pairs in the control group and Treatment Group 2. Social engagement is defined based on users' endorsement activity on others' posts. The dependent variables are whether a user liked the ad and whether she clicked on the ad at users' first ad impressions. $*p < 0.10$;$**p < 0.05$;$***p < 0.01$.

Table A5: Effects of Increasing Numbers of Social Cues on Social Influence for More or Less Socially Engaged Users (#Comments)

	Ad Like		Ad Click	
	More Socially Engaged	Less Socially Engaged	More Socially Engaged	Less Socially Engaged
Displaying Social Cues (β_1)	0.382***	0.337***	0.609***	0.612***
	(0.0311)	(0.0359)	(0.0538)	(0.0345)
#Social Cues (β_2)	0.0506***	0.0353***	0.0141**	0.0146***
	(0.00548)	(0.00256)	(0.00481)	(0.00368)
Displaying Social Cues × Social Cues (π)	0.0408***	0.0248***	0.0394	0.0168
	(0.00568)	(0.00236)	(0.00792)	(0.00467)
Controls	Yes	Yes	Yes	Yes
Observations	1,850,117	1,847,541	1,850,117	1,847,541
Log pseudolikelihood	-467738.05	-386313.91	-467738.05	-386313.91

Note. This table reports the estimation results of the bivariate probit model in Equation 3.2, separately for more and less socially engaged users, using all ad-user pairs in the control group and Treatment Group 2. Social engagement is defined based on users' commenting activity on others' posts. The dependent variables are whether a user liked the ad and whether she clicked on the ad at users' first ad impressions. $*p < 0.10$;$**p < 0.05$;$***p < 0.01$.

Table A6: Effects of Increasing Numbers of Social Cues on Social Influence for More or Less Socially Engaged Users (#Moments Posts)

	Ad Like		Ad Click	
	More Socially Engaged	Less Socially Engaged	More Socially Engaged	Less Socially Engaged
Displaying Social Cues (β_1)	0.374***	0.376***	0.527***	0.602***
	(0.0233)	(0.0335)	(0.0386)	(0.0473)
#Social Cues (β_2)	0.0381***	0.0557***	0.0118**	0.0210***
	(0.00355)	(0.00418)	(0.00359)	(0.00566)
Displaying Social Cues × Social Cues (π)	0.0326***	0.0306***	0.00377	-0.0000763
	(0.00341)	(0.00477)	(0.00548)	(0.00667)
Controls	Yes	Yes	Yes	Yes
Observations	1,882,677	1,814,981	1,882,677	1,814,981
Log pseudolikelihood	-504956.67	-361612	-504956.67	-361612

Note. This table reports the estimation results of the bivariate probit model in Equation 3.2, separately for more and less socially engaged users, using all ad-user pairs in the control group and Treatment Group 2. Social engagement is defined based on users' posting activity on WeChat Moments. The dependent variables are whether a user liked the ad and whether she clicked on the ad at users' first ad impressions. $*p < 0.10$;$**p < 0.05$;$***p < 0.01$.