Rational Herding in Social Advertising:
A Large-Scale Randomized Field Experiment

Shan Huang[1]

[1]MIT

Abstract

This study uses data from a large-scale field experiment to investigate how social cues (friends' likes) impact users' public (liking) and private (clicking) responses to social ads. The public responses will become social cues, which are broadcast with ads in social networks and are the source of social influence. The private responses are the main measure for ad engagement. In the experiment, I randomly manipulated the presence and the number of social cues (friends' likes) shown in ads among 37 million users of WeChat Moments ads. The results demonstrate that, on average, showing the first social cue significantly enhances users' liking and clicking propensity, but showing the additional social cues only increases users' tendency to like but does not affect their tendency to click an ad. Although users will always herd in publicly responding to (liking) an ad, I find the evidence of rational herding in users' private engagement with (clicking) an ad. It indicates that users infer the trustworthiness of social cues (likes) by observing the process of generating (liking) them. The first like, generated independent of social conformity, always exhibits significantly positive effects on ad clicking. The unpopularity of brands enables the herding momentum in clicking, as users infer the superior trustworthiness of social cues associated with small brands to justify the herd. Social influence in social advertising may fail if users attribute the herd of social cues to external factors, such as social conformity and popularity of brands.

Contact. shanh@mit.edu.

1. INTRODUCTION

Social advertising displays users' social (public) valuations in the form of social cues on its interface, utilizing social influence (effects of social cues) to deliver marketing communications in social networking sites (e.g. Bakshy et al. 2012, Tucker 2016). Public responses to social ads, such as endorsing (liking) ads, represent users' public (positive social) valuations regarding products or brands, which are always shown and broadcast with ads in social networks. Private responses, such as clicking ads, reflect a user's private engagement with ads and are usually not revealed to others. Displaying social cues in ads can drive private responses to ads, directly increasing social ad engagement[1]. Social cues can also increase users' public responses, leading to additional social cues. The goal of this study is to understand whether and how social influence and herding operate in users' public (liking) and private (clicking) responses to social ads. More specifically, I explore whether herding in liking and clicking is rational and how these underlying behavioral mechanisms inform social advertising strategies.

Rational herding occurs when users attempt to make unbiased inferences from their peers' decisions (active observational learning) in deciding whether to respond to an ad (Zhang and Liu 2012, Simonsohn and Ariely 2008). Irrational herding occurs when users passively mimic their peers' choices (e.g., saliency) or conform to their decisions as social norms (social conformity). Social influence may operate through very different mechanisms in impacting public and private behaviors. Theoretically, the visibility of behaviors allows individuals to signal their compliance with group norms to gain and protect their group memberships and maintain or improve their social status (Veblen 1899, Bagwell and Bernheim 1996, Corneo and Jeanne 1997, Wang and Griskevicius 2014). Passive mimicry is also a possible reason for generating social cues., which is absent from creating ad engagement, since individuals cannot

[1]Private responses are usually closely correlated with product impressions.

mimic others' invisible decisions. Existing theories also predict that consumers are motivated to learn from their friends' endorsements of ads to decide whether to engage with an ad (observational learning) (Huang et al. 2016). As a result, the herding in private responses (e.g., ad clicking) is more likely to be rational than that in public responses (e.g., liking).

It is important to distinguish between irrational and rational herding in public and private responses to social advertising. If irrational herding dominates public responses, the process of generating social cues (e.g., liking ads) will be self-reinforcing. Efforts such as showing the social cues earlier and making them more salient can both increase the number of social cues. However, rational observers care not only about the presence of social cues, but also the reasons that have given rise to them. Therefore, the external efforts that reinforce generating social cues (e.g., likes are driven by previous likes) may dilute the signal contained in the herd of social cues, as users may attribute herding to external efforts rather than to peers' intrinsic preferences. If the rational herding predominates in users private responses and the irrational herding predominates in users' public responses (social-cue generating process), the social cues may lose its effectiveness in lifting actual ad engagement. The external efforts, such as reinforcing the social conformity that drives public responses and the saliency of social cues through viral designs, may hurt the effects of social cues on ad engagement.

Social influence is of key importance and interest in social science (e.g., Deutsch and Gerard 1955, Burnkrant and Cousineau 1975, Sacerdote 2001, Cialdini and Goldstein 2004, Christakis and Fowler 2013). However, evaluating its effect has been difficult, since it is typically confounded with homophily, external stimuli and simultaneity (Manski 1993) and has tiny but nuanced effects in many contexts (e.g., Aral and Walker 2011a, Bakshy et al. 2012). Emerging large-scale online social networking sites, which increasingly connect people and mediate their communications, provide unprecedented opportunities for researchers to observe individuals' behaviors in exogenously manipulated settings at population scale (e.g., Bakshy et al. 2009, Aral et al. 2013, Bond et al. 2012, Muchnik et al. 2013, Tucker 2016).

I therefore designed and analyzed a large-scale randomized field experiment on a world-leading mobile social networking site, WeChat, examining the effect of social influence in public and private ad responses in a random sample of more than 37 million users of WeChat Moments ads, a type of social advertisements embedded in WeChat users' newsfeed. In the experiment, the presence and the number of friends' likes (social cues) displayed in ads were randomly assigned to more than 57 million ad-user pairs. This intervention lasted for 21 days and across 99 ads. Social influence is measured as the degree to which displaying (first-degree) friends' likes impacts users' public (liking) and private (clicking) responses to social ads. The randomized manipulation of treatments in a real context enables me to obtain an unbiased causal estimate of social influence in social advertising. The impression-level data on users' ad responses and the individual-level data on users' demographics and historical behaviors also allow me to explore behavioral mechanisms of influence. The research setting of WeChat Moments ads is particularly suited for contrasting social influence in public and private behaviors. On WeChat, users can signal their preferences and attitudes through public valuations to the *identifiable* audience, making the public behaviors on WeChat an effective signaling device. Unlike Facebook, WeChat shares 100% of users' posts and public social valuations only among their first-degree friends and without filtering or ranking algorithms.

The experimental evidence shows that displaying a friend's like in ads significantly increases users' tendency to (privately) click and (publicly) like an ad. I find that compared to displaying no like, displaying one like caused a 0.98% increase in users' liking propensity and a 0.96% increase in their clicking propensity. This result directly indicates that social influence works in social advertising.

However, I observe that the additional social cues only enhance users' liking propensity; in contrast, their clicking propensity does not change any further when the number of likes increases from one to more. Showing more likes in ads makes users' likelihood of endorsing an ad more probable, leading to more social cues on ads. But the additional social influence from the increased number of social cues cannot lead to significantly more ad clicks in response. As a result, at user-ad level, only the first social cue matters.

I also explore the behavioral mechanism underlying herding in public (liking) and private (clicking) responses. The results show that the presence of herding in liking is not dependent on the characteristics of ads or endorsers (the friends shown in ads). In contrast, I find the evidence of rational herding in users' private engagement with (clicking) an ad. It seems that users infer the trustworthiness of social cues (likes) by observing the process of generating (liking) them. The first like, created independent of social conformity, always exhibits significantly positive effects on ad clicking. The unpopularity of brands (e.g., small brands) enables the herding momentum in clicking, as users infer the superior trustworthiness of social cues to justify their herding. It seems that social influence (the effects of social cues) on ad engagement in social advertising may be diluted if users attribute the herd of social cues to observational factors, such as social conformity and popularity of brands. Finally, I eliminated the possibility that the signals contained in social cues are weakened by the friends' attributes instead of the characteristic of social-cue generating process. I found that the additional social cues are not able to increase users' clicking propensity even if all of them are created by the friends with greater product expertise than the users.

I proceed first by describing the research setting, experimental design, and data. I then present the findings and discuss their implications for theory, research, and practice.

2. FIELD EXPERIMENT

2.1. Experimental Setting: WeChat Moments Ads

I use data from a mobile-based field experiment in WeChat Moments ads. WeChat is one of the largest messaging apps, with over 938 million monthly active users spending ,on average, more than 90 minutes a day on the app. WeChat Moments, like Facebook's newsfeeds, supports posting images and texts, as well as sharing music, articles, and short videos. WeChat Moments ads, similar to Facebook ads, appear on the timeline of Moments and were launched in the spring of 2015. Users can click, endorse (like) and comment on ads in WeChat Moments (See Figure 1). The field experiment started in Dec of 2015. 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.

Several features of WeChat Moments ads make them distinct from Facebook ads and particularly suitable for contrasting social influence in users' public and private responses to social ads. First, public responses to WeChat Moments ads, such as liking, can be an effective social signaling device, while the private ad responses, such as clicking, cannot. WeChat Moments show 100% of users' contents, including their posts, likes, and comments, only to their first-degree friends. For example, once a WeChat user endorses an ad, they will expect that their endorsements will be shared (only) among their first-degree friends. In contrast, Facebook users, on average, receive less than 10% of the organic feeds. Their public behaviors can be shown to strangers, such as some of their second-degree friends. These ambiguities of the audience will significantly reduce Facebook users' motivation to socially signal their preferences and opinions through the public behaviors on the newsfeed.

Second, the design of WeChat Moments' ads allows researchers to identify the marginal social influence: comparing the effectiveness of the social ads between the groups with and without social cues shown in ads. As far as I know, the current Facebook experiments did not compare ads with and without social cues, holding anything else equal. For example, Bakshy et al. (2012) conducted an experiment on Facebook to identify the effects of social cues in social ads. However, they only compared ads with the different number of social cues and compared ads with the regular social cues (associated with particular friends' names) and with global social cues (the number of the endorsers). Tucker (2016) studied social ads effectiveness by comparing socially targeted ads with social cues and demographically targeted ads. It well addresses the effectiveness of social ads, but it does not distinguish the effectiveness of social targeting from the effectiveness of social influence (the effect of social cues). The extant study provides the first large-scale experimental evidence that identifies the marginal effect of social influence in social ads.

2.2. Experimental Design

Social influence in the experiment is measured as the effects of social cues (likes) that represent peer (first-degree friends) endorsements on users' public and private responses to social ads. I, therefore, randomized the presence and the number of social cues shown in ads. In the experiment, as users received a new ad, they were randomly assigned into three experimental groups: without any social cue (control group), with maximum one like (treatment group 1) and with organic likes shown in ads (treatment group 2), or outside the experiment (See figure 2). This randomization would happen again whenever users received a new ad. Users, therefore, could be in a different experimental group or outside the experiment for different ads. The randomization was at user-ad level. Every ad stayed in users' newsfeed only for maximum 48 hours. After 48 hours, the old ads would disappear, and a new ad would be received. In this way, users saw only one ad at one time in WeChat Moments during the experiment.

I can manipulate only the maximum number of social cues instead of the exact number of them, since all the social cues have to be organic and fake social cues are not allowed on WeChat. The group with no social cue and maximum one social cue enable me to identify the marginal social influence in social ads - the effect of one social cue. The variation in the number of social cues in the group with organic social cues allows me to estimate the social influence from a different number of friends. There are two types of potential social cues on WeChat Moments ads: likes and comments. Since comments have both positive and negative sentiments, I focus on likes as the only social cues, which uniformly represent friends' endorsements of an ad, and hid (controlled for) all the comments from the interface of ads during the experiment.

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 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 impression, reducing the likelihood that they have a different experience than (and therefore talk to their friends about) their experience. 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.

2.3. Data Collection

I collected four kinds of data to examine the magnitude and mechanisms of social influence in public and private ad response. First, the impression-level data on the number of social cues shown in ads is essential for the identification of social influence. I also collected data on the number of the social cues hid from the ads interface due to the experimental treatment to indicate homophily. The number of peer endorsements correlates not only with the social influence in social ads but also with users' latent preferences for products, brands or ads. The more friends' endorse an ad, the more likely a user favors this ad, since similar people tend to associate together (homophily). Second, the dependent variables for this study are users' public and private responses to social ads. I recorded the impression-level data on users' binary responses to an ad - whether to like" and whether to click an ad, and the response time. I counted any click on a given ad as long as they click on the profile page, links to the landing page or product photos (See figure 1). Third, to explore the heterogeneous effects of social influence and construct the control variables, I collected data on users' and their affiliated friends' characteristics, such as their demographics (age, gender, city) and historical behaviors on WeChat. [2] Fourth, I also recorded the name of the products and brands associated with the ads participated in the experiments.

2.4. Descriptive Statistics

The experiment was conducted over a 21-day period starting on Dec 22th of 2015, 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 cue. 19,174,955 user-ad pairs were randomly assigned to the treatment group with a maximum of one like. 19,137,118 user-ad pairs were randomly assigned to the treatment group with organic likes. I dropped 17 ads with invalid data and finally got 82 ads [3].

The maximum number of social cues that we can display is limited by the number of organic likes posted by friends of the ad viewers. To guarantee that at least one social cue was shown in ads in treatment groups, I further filtered the data on the condition that there was at least one organic like posted before the ad viewers receive ads. I finally got a sample of 5,571,116 user-ad pairs and 4,884,070 distinct users in total for our analysis from three groups. Zero like was displayed to the users in the control group, and one like and organic likes (at least one) were correctly displayed to the users in treatment group 1 and treatment group 2 (see Table 1). There are ,on average, 1.672 likes shown in ads in the treatment group 2, which is significantly greater than one ($p < 0.01$). There is no economically significant mean difference among these three different groups, regarding users' age, gender, network degree (number of WeChat friends), and level of WeChat Moment's activity (log-in days in

[2] *Affiliated friends* are the ones 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 treatments.

[3] We dropped ten old ads, which were left over from the pre-experiment period and another seven 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 underperforming ad campaigns.

November of 2015) (See Table 2).These pieces of evidence taken together confirm the integrity of the randomization procedure.

3. EMPIRICAL ANALYSIS

I use data from the large-scale field experiment to investigate how social influence operates in social advertising through its impact on public and private ad responses. I apply bivariate probit models in analysis. The decisions of whether to like and whether to click an ad are always jointly decided by individuals. The bivariate probit models are thus appropriate for simultaneously estimating these two (potentially) correlated decisions. I focus on users' behaviors in their first impressions of ads in the main analysis [4].

3.1. Modeling the Effects of One Social Cue

I first model users' decisions to respond to an ad. Upon receiving a new ad, user-ad pair i was randomly assigned to the control group with no social cue, the treatment group 1 with one like and the treatment group 2 with organic likes. The user then made a decision about whether to like (Y_{i1}) and whether to click (Y_{i2}) an ad, which are assumed to be determined by the unobserved latent variables Y_{i1}^* and Y_{i2}^*. In this section, I analyze how the presence of one (marginal) social cue impacts users' public and private ad responses using the data from the control group and the treatment group 1 with one like. In particular, I estimate the average treatment effects of showing one friends' like on users' propensity to like and click an ad.

$$\begin{cases} Y_{i1}^* = \alpha_1 + \beta_{11} S_i + C_i \theta_1 + \epsilon_{i1} \\ Y_{i2}^* = \alpha_2 + \beta_{12} S_i + C_i \theta_2 + \epsilon_{i2} \end{cases} \tag{3.1}$$

[4] I also compared the social influence in the first impressions and entire impressions for ads in robustness checks.

Only Y_{ij} are observed, which equals 1 only if $Y^*_{ij} > 0$, implying that the user chooses to respond. Y_{ij} equals zero if the user chooses not to. S_i is a dummy variable denoting whether user-ad pair i was assigned into the treatment group with one social cue or not. C_i represents a vector of control variables, including users' demographic controls (specifically, indicator variables for whether users are male, whether users belong to specific age groups and whether users live in important/big/others cities) and week dummies.

Assuming that ϵ_{i1} and ϵ_{i2} are distributed as bivariate normal with mean zero, unit variance, and $\rho = \mathrm{Corr}(\epsilon_{i1}, \epsilon_{i2})$. I estimate the bivariate probit regression model in Equation 3.2, which is appropriate when $\rho \neq 0$ (Greene 2003).

$$Pr(Y_{ij} = 1) = \Phi(\alpha_j + \beta_{1j} S_i + \theta_j C_i) \tag{3.2}$$

where ϕ is the cumulative normal distribution function. The coefficient β_{1j} captures the causal effects of S_i on $Pr(Y_{ij} = 1)$, $j = 1, 2$. Since user-ad pairs were randomly assigned to the control and treatment groups, S_i is orthogonal to ϵ_{ij}, $j = 1, 2$. β_{1j} thus captures the causal effects of one social cue on users' tendency to like and to click an ad. The control variables should primarily improve the estimation efficiency.

3.2. Modeling the Effects of Different Number of Social Cues

I then identify the effects of different number of social cues in liking and clicking and examine the growth pattern of social influence, using the data from the control group without social cue and the treatment group 2 with organic likes shown in ads. In particular, I estimate the bivariate probit regression model in Equation 3.3 to identify how the effects of social cues (likes) in liking (Y_{i1}) and clicking (Y_{i2}) vary across different number of them shown in ads.

$$Pr(Y_{ij} = 1) = \phi(\alpha_j + \beta_{1j}SS_i + \beta_{2j}N_i + \pi_j N_i \times SS_i + \theta_j C_i + \Delta_j C_i \times SS_i) \quad (3.3)$$

where $j = 1, 2$, N_i indicates the number of organic social cues associated with user-ad pair i in the first ad impressions, and SS_i is a dummy denoting whether user-ad pair i is in the treatment group 2. Note that N_i is the number of social cues displayed for user-ad pair i in the treatment group, while N_i is the number of social cues hid from the user-ad pair i in the control group. Since the data was filtered by the condition that there was at least one organic social cues in all the observations (See Section 2.4), $N_i >= 1$.

The coefficient π_j on the interaction term is of my main interest and indicates the degree to which showing additional one social cue in ads changes the social influence - the effects of social cues. The coefficient β_{1j} captures the raw impact of displaying social cues in ads on users' responses to ads, holding constant heterogeneity across other variables. The coefficient β_{2j} captures the tendency of users with more organic social cues to spontaneously respond to an ad in the absence of influence ($SS_i = 0$). β_{2j} indicates the degree to which the number of organic social cues predicts users' correlated latent preference to adopt an ad without social influence.

I control for not only the user' demographics C_i but also their interactions with social influence $C_i \times SS_i$. While we can evaluate how the *number* of social cues moderates social influence (the effects of social cues), we cannot make casual claims about this estimate. Because in the experiment, I only randomized the *presence* of one and organic social cues and did not directly manipulate the *number* of the social cues shown in ads. As a result, I can causally identify social influence but cannot causally estimate the moderating effects of the number of social cues on social influence. The groups with different number of organic social cues may be systematically different in other dimensions. For example, the number of organic social cues can be an indicator for users' intrinsic preference for ads. The correlated preferences among friends (homophily) predict that users are more likely to favor an ad, when they have more friends endorsing it. Users' intrinsic preferences for ads may also impact social influence, confounding with the number of social cues. Therefore, I add $C_i \times SS_i$ in the model (See Equation 3.3) to control for the effect of users' characteristics on social influence. I revisit this issue in manipulation checks as well.

4. RESULTS

4.1. Effects of One Social Cue

Columns 1,2,5 and 6 in Table 3 report the coefficients estimated from the bivariate probit model specified in equation 3.2 without any controls. Columns 3,4,7 and 8 in Table 3 repeat the analysis, including controls for users' demographics and week dummies. All the controls are measured using the data in Nov of 2015 - one month before the experiment started. The results hold after adding these controls, which reassures that any unevenness across different users does not drive the results. It confirms the success of the randomization in the experiment and the robustness of the results. $\rho = 0.252, p < 0.01$ shows that liking and clicking are positively correlated decisions, indicating that the bivariate probit model is appropriate for our data.

I find that the marginal effects of showing one like in ads on users' tendency to like and click an ad are significantly positive. Compared to no social cue, displaying a like causes a user 0.98% ($p < 0.01$) more likely to like an ad and 0.96% ($p < 0.01$) more likely to click an ad. Three main observations arise from consideration of these results. First, social influence *directly* and significantly enhances social ad engagements (ad clicks). To the best of my knowledge, this is the first experiment that identifies the marginal effect of social cues on ad engagement. Second, social influence also leads to additional social cues. Showing one like in ads significantly increases users' liking propensity. Third, the effects of one social cue are not significantly different between liking and clicking ($p > 0.1$).

I then transformed the two dependent variables into three mutually exclusive behaviors: pure clicking, pure liking, and clicking & liking. The marginal effects of showing one social cue on these three ad responses, which are all significantly positive. I observe that showing a social cue impacts most the users' propensity to purely like an ad. Pure liking is a behavior that is possible because there can be an inconsistency between individuals' public and private responses to social ads. This result indicates that showing a social cue greatly increases users' tendency to publicly endorse (like) an ad without privately reading (clicking) it.

4.2. Effects of Different Numbers of Social Cues : Cue-Response Curve

In Table 4, I present the evidence that reflects whether the number of social cues moderates their effects on users' propensity to like and click an ad. The coefficient on the interaction term is significantly positive for liking ($\pi_1 = 0.0375$, $p < 0.01$) but not significant for clicking ($\pi_2 = -0.0012$, $p > 0.1$). This result suggests that social influence in liking significantly grows with the number of social cues shown in ads, while social influence in clicking is not affected by the number of the social cues.

In considering the interpretation of these results, it is important to note that the distribution of marginal effects on interaction terms is not constant across covariates in nonlinear models (Ai and Norton 2003). To thus home in on the interaction effects $N_i \times SS_i$, I estimate social influence - effects of social cues - in liking and clicking for the user-ad pairs associated with different numbers of likes shown in ads. Figure 4 then displays the corresponding distributions of estimated marginal effects of social cues. It shows that the social influence in liking significantly increases with the number of likes shown in ads, and thus shows that the significantly positive interaction effect holds in liking. It also indicates that social influence in clicking stays the same as the number of likes grows, which is consistent with the insignificance of the coefficients of the interaction term for clicking in Equation 3.3. Now it is safe to conclude that the number of social cues only positively impacts the degree of social influence in liking not in clicking. It seems that showing increasing numbers of likes in ads makes users linearly more likely to like an ad but does not change users' tendency to click it.

I then transformed the two dependent variables into three mutually exclusive behaviors: pure clicking, pure liking, and clicking & liking. Figure 5 shows that the social influence grows with the number of social cues most in pure liking, but not in pure clicking. This result indicates that more social cues increases their effects on users' tendency to publicly endorse (like) an ad without privately reading (clicking) it. Two implications arise from these observations. First, the first social cue makes users significantly more likely to click an ad, directly increasing ad effectiveness. However, the subsequent social cues only increase pure likes, and therefore can not generate additional social influence on ad clicks. Second, social influence greatly enhances pure likes, making like rate discrepant from click rate. Like rate, the public performance of ads, may not be a relatable indicator for users' actual attention paid to ads and ads effectiveness.

The model specification also enables me to measure how the number of likes predict users' latent preferences and spontaneous responses to an ad, in the absence of social influence. Homophily indicates that similar people tend to associate with one another. Therefore, the more friends endorse an ad, the more likely a user respond to an ad. Consistent with the theory, the coefficient β_{21} is significantly positive ($p < 0.01$), indicating that the number of friends' endorsements (friends' likes) predicts users' tendency to like an ad. Predictors of spontaneous adoptions can inform social targeting: users with more friends' who have liked an ad are more likely to endorse the ad. However, the coefficient β_{22} is insignificant ($p > 0.1$), indicating that the number of friends' likes does not predict users' tendency to click an ad. This result implies that targeting the users with more friends, who endorsed an ad, won't significantly improve their probability to click the ad. This result also suggests that users' may have different motivation for liking and clicking ads.

4.3. Economic Implications

Sections 4.1 and 4.2 establish that social cue significantly increases ad engagement. However, we have not yet established that social influence (the effects of social cues) is economically meaningful. Columns 8 in Table 3 suggests that seeing one social cue at users' first ad impressions increases click intention by 0.96%, compared to seeing no social cue. Figures 4 and 5 show that seeing multiple social cues does not make users more likely to click than seeing one social cue. These results together indicate that showing (at least one) social cues in an ad enhances its effectiveness by 0.96% clicking propensity in users' first impressions of the ad. Assuming that the price difference between social ads and non-social ads is Δ/impression[5], this means that social advertising pays off if increasing one percentage point of clicking propensity is worth more than $\frac{\Delta}{0.96\%}$ to the firm. The market price of Δ is around 1.5 cents in China, and this means that social advertising pays off if increasing one percentage point of clicking propensity is worth more than \$1.63 to the firm.

[5]Here social ads refer to the ads with social cue shown in ads, while non-social ads mean no social cue shown in ads.

4.4. Underlying Behavioral Mechanisms : Irrational vs. Rational Herding

The result that only the first social cue significantly impacts users' clicking propensity provides some evidence that herding in clicking is likely to be rational. That is because, unlike additional social cues, the first social cue is generated independently of social conformity and represents users' intrinsic preferences. This distinct social-cue generating process increases the trustworthiness of the first social cue for observers.

In this section, I further explore whether herding in clicking is rational, through comparing the cue-response curves between small and big brands. If the unpopularity of brands amplifies the herding momentum, rational herding may be present. That is because the social cues for small brands signal superior trustworthiness and rational observers will be more affected by them. I also eliminate the possibility of the other competing behavioral mechanism : the signals contained in social cues are weakened by the characteristics of friends instead of the social-cue generating process. I test whether the additional social cues would increase users' tendency to click, if the friends who generated those social cues posses greater product expertise. To validate whether herding in liking is irrational, I examine whether the presence of herding in liking depends on observable characteristics of endorsers (friends shown in ads) and ads.

Similar regressions as in Equation 3.3 were run on different subgroups: small/big brands and friends with greater/less product expertise. Based on Interbrands list of the top 100 brands for 2016, I code a binary variable to indicate whether the brand in an ad is a big brand or not(Lovett et al. 2014). 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 algorithm has been well applied in WeChat daily operations.

Table 5 shows that the coefficients on the interaction term are significantly positive for liking ($p < 0.01$), whether the ads are for small or big brands or the social cues in ads are created by the friends with greater product knowledge or not. This suggests that the presence of herding in liking does not rely on any observable characteristics that I examined, and therefore the herding in liking is presumably irrational.

However, although I show that, on average, the additional social cues do not impact users' clicking propensity, they further increase users' tendency to click for small brands. The coefficient on the interaction term for clicking is significantly positive for small brands ($p < 0.01$, See Model 2 in Table 5 and Figure 6). However, the additional social cues even decrease users' clicking propensity for big brands ($p < 0.01$, See Model 1 in Table 5 and Figure 6). These results indicate the rational herding in users' (private) engagement with ads (clicking).

Finally, I eliminate the possibility that, instead of observing the social-cue generating process, users are observing the friends shown in ads and then judging whether to learn from them. I focus on the two unique subgroups, in which the friends shown in ads had greater or less product expertise than the users (the ad viewers). I find that the additional social cues still cannot further increase users' tendency to click even if the additional social cues are created by the friends with greater product expertise ($p > 0.1$, See Model 3 and 4 in Table 5 and Figure **??**)

4.5. Robustness: Social Influence in Entire Ad Impressions

I show the robustness to measuring the effects of social cues on users' responses in their entire impression of ads. In the main analysis, I focus on how showing social cues in ads impacts users' responses in their first impressions of an ad. This is the cleanest measure for influence identification. However, advertisers are also interested in whether users respond to their ads during the entire ad impressions. Once users engage with an ad at any ad impressions, the product impressions will be created. If first impressions play a significant role in generating social influence, if pays to maximize the influence in users' first impressions to ads. I, therefore, examine how social influence affects users' responses in their entire impressions of ads. I construct a measure for users' responses to ads in multiple ads impressions, counting any click or like on a given ad as long as users click or like the ad at any impression when it is displayed in users' WeChat Moments. In the experiment, users, on average, had three impressions for an ad. I replicate the previous analysis and the results in general hold.

5. CONCLUSION

This study moves the research frontier from identifying the presence of social influence in social ads to identifying and comparing the influence mechanisms in different ad responses (public and private responses), using a large-scale networked experiment. Recent developments in digital experimentation and large-scale social networking sites have provided new opportunities to identify social influence and its heterogeneous effects (Aral and Walker 2011b, Taylor et al. 2015, Huang et al. 2016). Although previous research has demonstrated the effectiveness of social ads over non-social ads (e.g., Bakshy et al. 2012, Tucker 2016), little is known about the potential different underlying mechanisms of social influence in different responses to social ads.

There are four novel findings. First, social influence operates in both public and private responses to social ads. Showing social cues (friends' likes) in ads significantly raises users' probability to both like and click an ad. This directly indicates the effectiveness of social ads. Second, although the first social cue increases both the liking and clicking propensities, the additional social cues only increase users' liking propensity but do not impact users' clicking propensity. As a result, in each ad impression, only the first social cue contributes to ad engagement. Third, I find the evidence of rational herding in ad engagements. Instead of passively following friends' choices and taking them as social norms, users actively learn and attempt to make unbiased inferences from the social cues. Only for small brands, users' tendency to privately engage with (clicking) an ad significantly increases with the number of social cues in ads. The unpopularity of brands signals the superior trustworthiness of signals contained in the herd. Fourth, since the herding in public responses (the social-cue generating process) is more irrational, the irrationality inherent in the herd of social cues may dilute its signal quality and effectiveness in enhancing ad engagement. The rational observers care not only about the presence of the herd (the social cues) but also about its generating process.

The managerial applications are two-fold. First, it is evident that social cues are useful in enhancing ad effectiveness. However, this effect is more salient in increasing the number of social cues than in raising private ad engagement. Second, users are rational observers in deciding whether to privately engage with an ad. Especially when the herding in public responses is irrational, the quality of signals contained in the herd of social cues will be weakened by the observable irrationality in the social-cue generating process. External efforts, such as viral product designs to increase social conformity that leads to social cues, or to enhance the saliency of social cues, may even reduce the effectiveness of social ads. Users (ad viewers) are likely to attribute the herding in public responses (generating social cues) to these external factors and thus be less affected by it. This paper identifies a potential trade-off between increasing the social cues by external efforts and raising ad engagement through social influence in social advertising.

References

Chunrong Ai and Edward C. Norton. Interaction terms in logit and probit models. *Economics Letters*, 80(1):123–129, 2003. ISSN 01651765. doi: 10.1016/S0165-1765(03)00032-6.

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, 2011a.

Sinan Aral and Dylan Walker. Identifying social influence in networks using randomized experiments. *IEEE Intelligent Systems*, 26(5):91–96, 2011b.

Sinan Aral, Lev Muchnik, and Arun Sundararajan. Engineering Social Contagions: Optimal Network Seeding and Incentive Strategies. *Network Science*, 1(2):125–153, 2013.

Laurie Simon Bagwell and B. Douglas Bernheim. Veblen effects in a theory of conspicuous consumption. *American Economic Review*, 86(3):349–373, 1996.

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

Eytan Bakshy, Brian Karrer, and Lada a Adamic. Social influence and the diffusion of user created content. In *Electronic Commerce*, pages 325–334, 2009.

Robert M. Bond, Christopher J. Fariss, Jason J. Jones, Adam D. I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295–298, 2012.

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

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*, 55:591–621, 2004.

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

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.

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

Shan Huang, Sinan Aral, Y. Yu (Jeffrey) Hu, and Erik Brynjolfsson. Social Influence across Products. In *Conference on Digital Experimentation (CODE)*, pages 1–24, Boston, 2016.

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, 1993.

Lev Muchnik, Sinan Aral, and Sean J Taylor. Social influence bias: a randomized experiment. *Science (New York, N.Y.)*, 341(6146):647–51, 2013.

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

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

Uri Simonsohn and Dan Ariely. When Rational Sellers Face Nonrational Buyers: Evidence from Herding on eBay. *Management Science*, 54(9):1624–1637, 2008.

Sean J Taylor, Eytan Bakshy, Dean Eckles, and Sinan Aral. Heterogeneous social influence from personalized social cues. 2015.

C E Tucker. Social advertising: How advertising that explicitly promotes social influence can backfire. 2016.

Thorstein Veblen. *The theory of the leisure classes: An economic study of institutions*. Macmillan, 1899.

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

Juanjuan Zhang and Peng Liu. Rational Herding in Microloan Markets. *Management Science*, 58 (5):892–912, 2012.

6. FIGURES AND TABLES

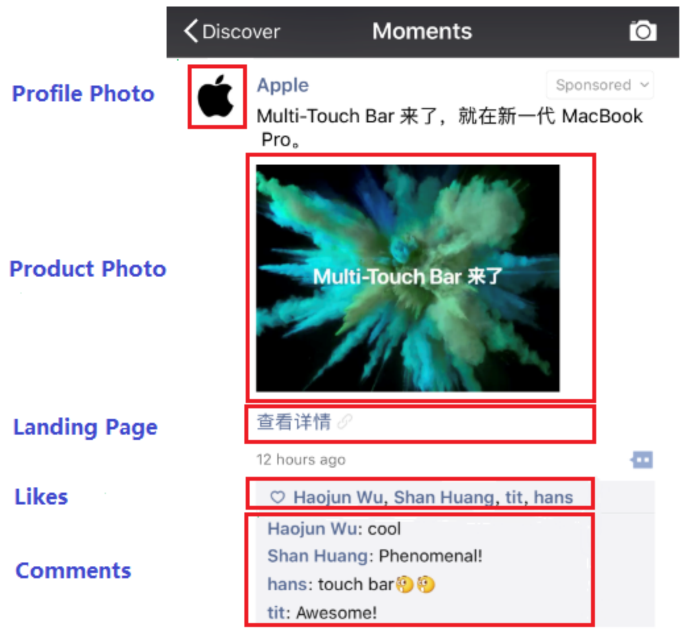

Figure 1: Example of WeChat Moments Ads

Note: This figure provides an example of WeChat Moments Ads, where the experiment was conducted. WeChat Moments supports posting images and texts, as well as sharing music, articles, and short videos. WeChat Moments ads, similar to Facebook ads, appear on the timeline of Moments and were launched in the spring of 2015. Users can click, endorse (like) and comment on the ads in WeChat Moments.

Figure 2: Experimental Treatments

Note: This figure provides an example that illustrates the experimental design used in this study. In the experiment, as users received a new ad, they were randomly assigned into three groups: without any social cue (control group), with maximum one like (treatment group 1) and with organic likes shown in ads (treatment group 2), or outside the experiment. This randomization would happen again whenever users received a new ad. Users could thus be in a different treatment group or outside the experiment for different ads. The randomization was at user-ad level. Every ad stayed in users' newsfeed only for maximum 48 hours. After 48 hours, the old ad would disappear, and a new ad was received. In this way, users saw only one ad at a time in WeChat Moments during the experiment. Since comments may be positive or negative about the ads or products, I focused, in this paper, exclusively on the effect of likes and hid all friends' comments on ads during the experiment.

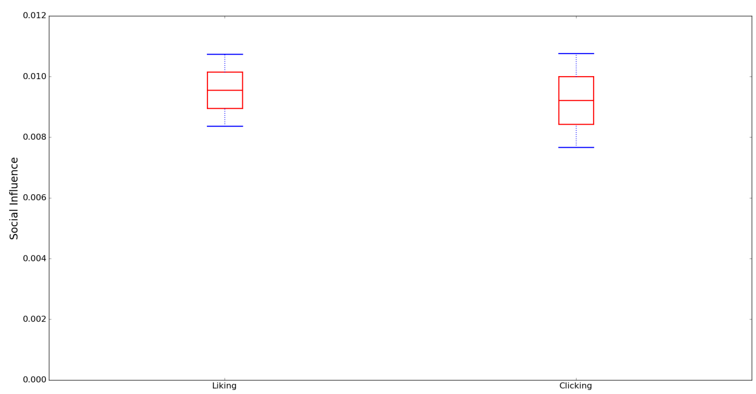

Figure 3: Effects of One Social Cue in Liking and Clicking

Note: This figure presents the marginal effects of one social cue on users' public (liking) and private (clicking) responses to social ads, which are shown with SEs (boxes) and 95% confidence intervals (whiskers). This figure is based on user-ad observations; the coefficients are estimated from the bivariate probit regressions. The sample includes the user-ad pairs of control group and treatment group 1. The effects of one social cue (marginal social influence) on users' tendency to like and click an ad are significantly positive.

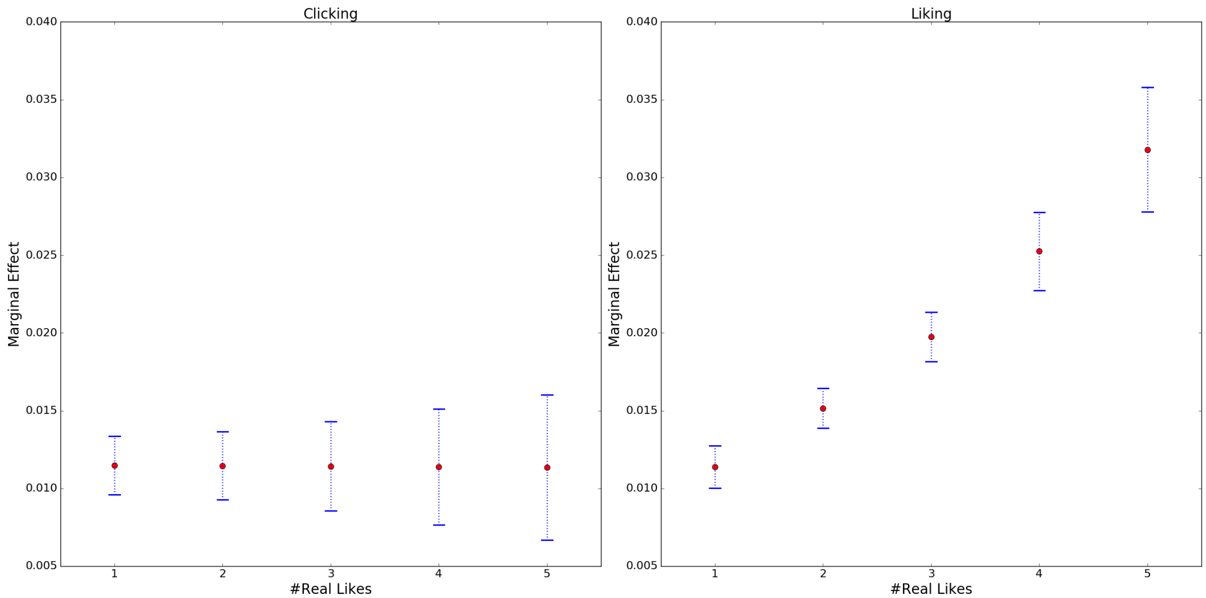

Figure 4: Cue-Response Curves of Liking and Clicking

Note: This figure plots the relationship between the marginal effects of social cues in liking and clicking and the number of social cues shown in ads. On the x axis is the number of social cues. The social influence associated with the different number of social cues is shown with SEs (boxes) and 95% confidence intervals (whiskers). This figure is based on user-ad observations; the coefficients are estimated from the bivariate probit regressions. The sample includes the user-ad pairs from two groups: the control group without any social cue and the treatment group 2 with organic social cues. The effect of social cues significantly grows in liking, as the number of social cues increases. However, the effect of social cues in clicking does not change at the same time.

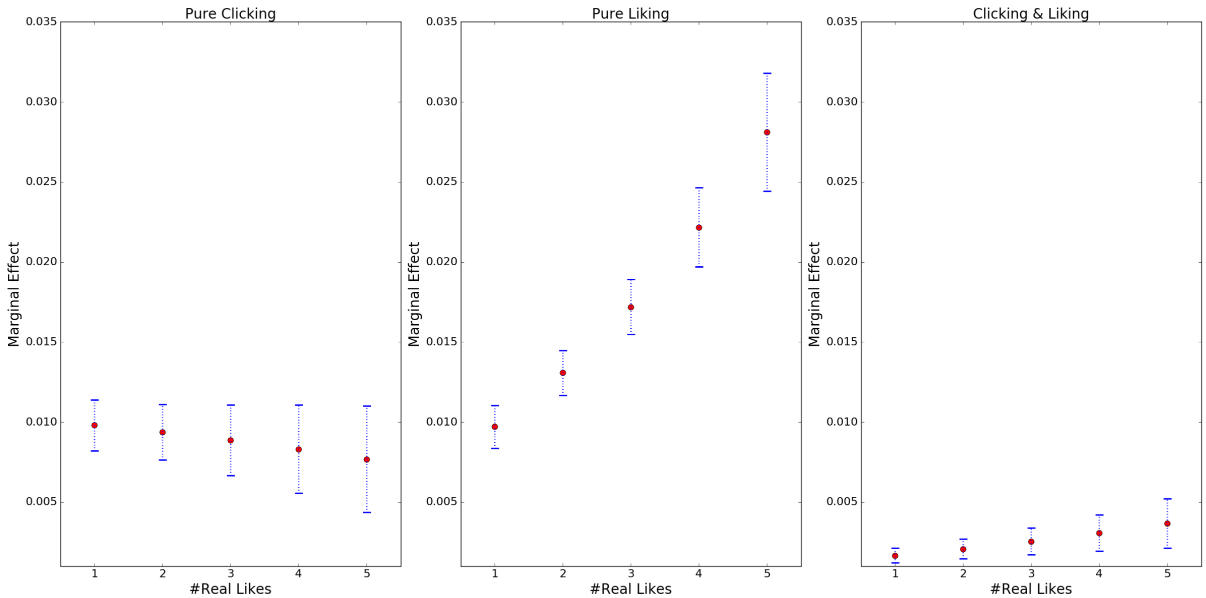

Figure 5: Cue-Response Curves of Pure Liking, Pure Clicking and Liking&Clicking

Note: This figure plots the relationship between the marginal effects of social cues in pure liking, pure clicking and liking&clicking, and the number of social cues shown in ads. On the x axis is the number of social cues. The social influence associated with the different number of social cues are shown with SEs (boxes) and 95% confidence intervals (whiskers). This figure is based on user-ad observations; the coefficients are estimated from the bivariate probit regressions. The sample includes the user-ad pairs from two groups: the control group without any social cue and the treatment group 2 with organic social cues. The effect of social cues significantly grows most in pure liking, as the number of social cues increases. The effect of social cues in pure clicking does not change at the same time.

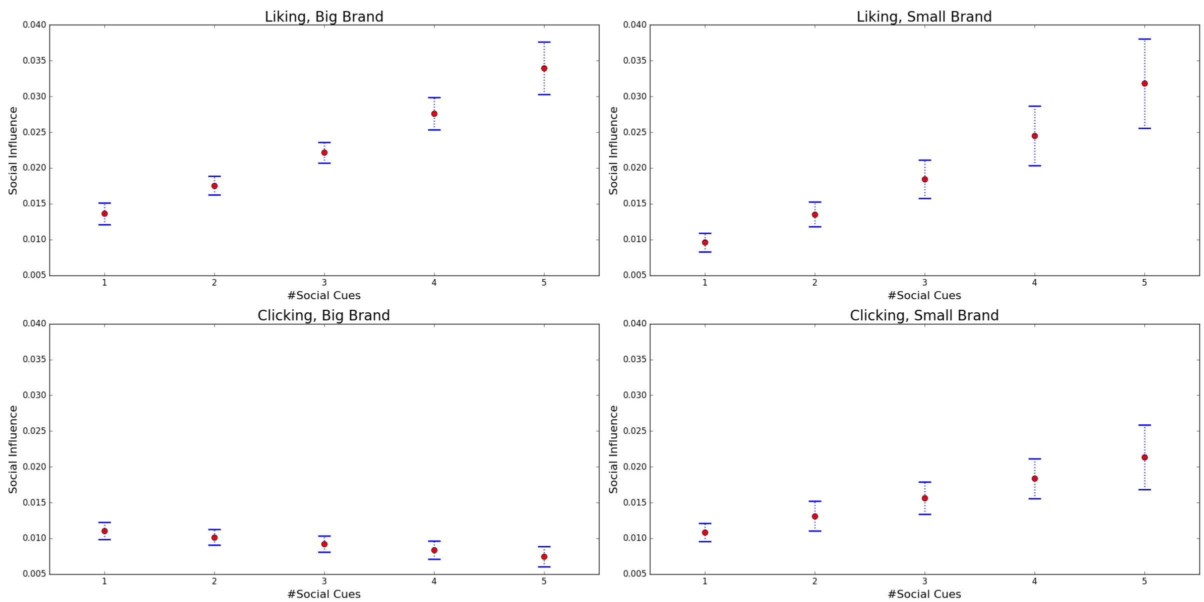

Figure 6: Cue-Response Curves for Ads of Small and Big Brands

Note: This figure focuses on the subsamples for small and brands. It plots the relationship between the marginal effects of social cues in liking and clicking and the number of social cues shown in ads. On the x axis is the number of social cues. The social influence associated with the different numbers of social cues is shown with SEs (boxes) and 95% confidence intervals (whiskers). This figure is based on user-ad observations; the coefficients are estimated from the bivariate probit regressions. The sample includes the user-ad pairs from two groups: the control group without any social cue and the treatment group 2 with organic social cues. The effect of social cues significantly grows in liking for big and small brands, as the number of social cues increases. In clicking, the effect of social cues significantly increases for small brands but decreases for big brands, as the number of social cues increases.

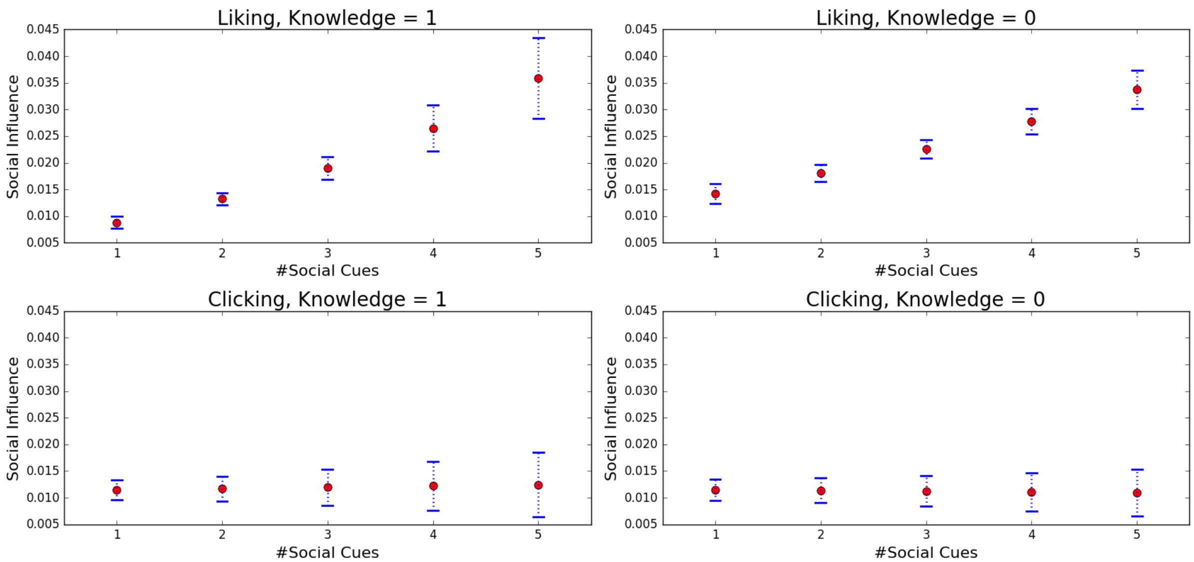

Figure 7: Cue-Response Curves : Friends with Greater or Less Product Knowledge

Note: This figure focuses on the subsamples, in which the social cues were created by the friends shown in ads with greater (knowledge=1) or less (knowledge=0) product knowledge than the users (the ad viewers). It plots the relationship between the marginal effects of social cues in liking and clicking and the number of social cues shown in ads. On the x axis is the number of social cues. The social influence associated with the different number of social cues is shown with SEs (boxes) and 95% confidence intervals (whiskers). This figure is based on user-ad observations; the coefficients are estimated from the bivariate probit regressions. The sample includes the user-ad pairs from two groups: the control group without any social cue and the treatment group 2 with organic social cues. The effect of social cues significantly grows in liking, as the number of social cues increases. The effect of social cues in clicking does not change at the same time.

Table 1: Manipulation Check

#Like	#user-ad pairs	#user	#Displayed likes			
			Mean	Std	Max	Min
0	1,860,622	1,775,820	0.000	0.000	0	0
1	1,873,401	1,787,240	1.000	0.000	1	1
2	1,837,093	1,755,895	1.672	1.742	100	1

Note: This table shows that the manipulation in the experiment is correct. "0" represents the control group. "1" represents treatment group 1. "2" represents treatment group 2. No likes were displayed to users in the control group, and a maximum of one like or the organic number of likes were correctly displayed to the users in the two treatment groups.

Table 2: Mean Comparisons Between Treatment Groups

	$(\lvert\#0-\#1\rvert < \#0 * X\%)$	$(\lvert\#0-\#2\rvert < \#0 * X\%)$	$(\lvert\#1-\#2\rvert < \#1 * X\%)$
	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: This table compares the mean between three groups and tested the hypothesis : mean(a) - mean(b) $<$ X% * mean(b). "0" represents the control group. "1" represents the treatment group 1. "2" represents the treatment group 2. I report 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 three groups. No economically meaningful mean differences is found 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) ($p > 0.1$). As a result, assignment to treatment groups is random.

Table 3: Effects of One Social Cue in Liking and Clicking

	Liking		Liking		Clicking		Clicking	
	1	2	3	4	5	6	7	8
	Coefficient	Marginal Effect	Coefficient	Marginal Effect	Coefficient	Marginal Effect	Coefficient	Marginal Effect
One Social Cue	0.1941***	0.0098***	0.1973***	0.0096***	0.1304***	0.0096***	0.1321***	0.0092***
	(0.0030)	(0.0001)	(0.0105)	(0.0006)	(0.0025)	(0.0002)	(0.0152)	(0.0008)
Controls	No		Yes		No		Yes	
Observations	3,734,023		3,734,023		3,734,023		3,734,023	
Log-Likelihood	-1486250.10		-1462622.40		-1486250.10		-1462622.40	

Note. This table compares the control group with no like and treatment group 1 with one like and shows that showing one (the first) like in ads significantly increases users' liking and clicking propensity.
* $*p < 0.10; **p < 0.05; ***p < 0.01$.

Table 4: Moderating Effects of the Number of Social Cues on Social Influence

	Liking	Clicking
	1	2
	Coefficient	Coefficient
Social Cues	0.1710***	0.1125***
	(0.0206)	(0.0327)
#Social Cues	0.0522***	0.0033
	(0.0049)	(0.0056)
Social Cues * #Social Cues	0.0375***	-0.0012
	(0.0038)	(0.0059)
Controls	Yes	Yes
Observations	3,697,715	
Log-Likelihood	-931223.36	

Note. This table uses the data of the control group with no like and treatment group 2 with organic likes to examine how the number of social cues moderates their effects.
$*p < 0.10; **p < 0.05; ***p < 0.01$.

Table 5: Behavioral Mechanism

		Model 1 Big Brand	Model 2 Small Brand	Model 3 Knowledge = 1	Model 4 Knowledge = 0
Liking	SC	0.0823*** (0.0127)	0.1777*** (0.0346)	0.1412*** (0.0302)	0.1881*** (0.0188)
	#Real Likes	0.0481*** (0.0049)	0.0580*** (0.0039)	0.0667*** (0.0104)	0.0439*** (0.0043)
	SC*#Real Likes	0.0335*** (0.0036)	0.0420*** (0.0070)	0.0552*** (0.0048)	0.0320*** (0.0036)
Clicking	SC	-0.0243 (0.0549)	0.1293*** (0.0325)	0.1209*** (0.0332)	0.1015*** (0.0336)
	#Real Likes	0.0049 (0.0089)	0.0250*** (0.0065)	-0.0008 (0.0079)	-0.0016 (0.0058)
	SC*#Real Likes	-0.0177*** (0.0043)	0.0183*** (0.0041)	0.0034 (0.0092)	-0.0013 (0.0054)
Observations		1,666,313	2,031,402	1,940,504	1,757,211

Note. This table uses the data of the control group with no like and treatment group 2 with organic likes in different subgroups to reveal the underlying mechanisms of social influence in liking and clicking.
$*p < 0.10; **p < 0.05; ***p < 0.01$.