

Social Advertising

Catherine Tucker[‡]

December 22, 2011

Abstract

In social advertising, ads are targeted based on underlying social networks and their content is tailored with information that pertains to the social relationship. This paper explores the effectiveness of social advertising using data from field tests of different ads on Facebook. We find evidence that social advertising is effective. However, social advertising is *less* effective if the advertiser explicitly states they are trying to promote social influence in the text of their ad. This suggests that advertisers must avoid being overt in their attempts to exploit social networks in their advertising.

*

[‡]Catherine Tucker is Associate Professor of Marketing at MIT Sloan School of Management, Cambridge, MA. and Faculty Research Fellow at the NBER. Thank-you to Google for financial support and to an anonymous non-profit for their cooperation. Thank-you to Ann Kronrod and seminar participants at the George Mason University Roundtable on the Law and Economics of Internet Search and the University of Rochester for valuable comments. All errors are my own.

1 Introduction

Recent advances on the internet have allowed consumers to interact across digital social networks. This is taking place at unprecedented levels: Facebook was the most visited website in the US in 2010, accounting for 20% of all time spent on the internet, a higher proportion than Google or Yahoo! (Comscore, 2011). However, it is striking that traditional marketing communications have been at the periphery of this explosion of social data despite the documented power of social influence on purchasing behavior.

Much of the emphasis on marketing in social media, so far, has been on the achievement of ‘earned reach,’ whereby a brand builds its subscriber base organically and also hopes that this will influence others organically through sharing links with their social networks (Corcoran, 2009). However, recent research by Bakshy et al. (2011) has emphasized that this kind of organic sharing is far rarer than previously supposed, and that there are very few examples of a commercial message being consistently transmitted across social networks. Further, Tucker (2011a) shows that in order to achieve virality, an advertiser may have to sacrifice the commercial effectiveness of their message.

This means that advertisers may need to use paid advertising to facilitate the sharing of their commercial message through social networks. Both Facebook and LinkedIn have recently introduced a new form of advertising called ‘social advertising.’ A social ad is an online ad that ‘incorporates user interactions that the consumer has agreed to display and be shared. The resulting ad displays these interactions along with the user’s persona (picture and/or name) within the ad content’ (IAB, 2009). This represents a radical technological development for advertisers, because it means that potentially they can co-opt the power of an individual’s social network to target advertising and engage their audience.

This paper asks whether social advertising is effective, and what active steps advertisers themselves should take in their ads to promote social influence.

We explore the effectiveness of social ads using data from a field experiment conducted on Facebook by a non-profit. This field experiment compared the performance of social ads with conventionally targeted and untargeted ads. The social ads were targeted to the friends of ‘fans’ of the charity on Facebook. The ads featured that fan’s name and the fact that they had become a fan of this charity. We find that on average these social ads were more effective than conventionally targeted or untargeted ads. Further, this technique is useful for improving both the performance of targeted and untargeted campaigns. Comparing the performance of these ads that contained the name of the fan with those that were simply targeted to that fan’s social network suggests that their effectiveness stems slightly more from the ability to select a more appropriate targeting group using social relationships than the endorsement of the friend. We present results that suggest that as well as being more effective at gathering clicks, social advertising is also more effective at promoting actual subscriptions to the newsfeed and is more cost-effective.

Given that previous attempts to commercialize social networks have been limited, we then turn to how advertisers should word their social advertising. Through randomized field tests, we investigate the effectiveness of advertisers deliberately promoting social influence in their advertising copy through including a statement that encourages the viewer to, for example, ‘be like their friend.’

We find that attempts by advertisers to explicitly harness or refer to a friend’s actions in their ad copy are rejected by consumers. This result contrasts with previous empirical research that has found consistent benefits to firms from making it obvious to consumers that they might allow social influence or previous consumer actions influence their behavior (Algesheimer et al., 2010; Tucker and Zhang, 2011). This rejection is reasonably uniform across different wording, though slightly less severe for ads that make a less explicit reference to friendship.

We then present additional evidence to rule out two potential explanations for our find-

ings. First, we rule out that the overt mention of social influence simply made people aware they were seeing an ad rather than something organic to the site. Second, to investigate whether it was simply bad advertising copy, we examined how the ads perform on a group of Facebook users who have shown a visible propensity for social influence. These consumers who explicitly have attachment to ‘Fashion’ goods reacted more positively to the advertiser explicitly co-opting social influence than to a message that did not. This suggests that it was not simply that the message was badly communicated, but instead reflects a taste (or more accurately distaste) for explicit references to social influence among most consumers.

This research builds on a literature that has studied the interplay between social networks and word of mouth. Zubcsek and Sarvary (2011) present a theoretical model that examines the effects of advertising to a social network, but assume that a firm cannot directly use the social network for marketing purposes. Instead, firms have to rely on consumers to organically pass their advertising message within the social networks. There has been little work on advertising in social networks. Previous studies in marketing about social networking sites have questioned how such sites can use advertising to obtain members (Trusov et al., 2009), and also how makers of applications designed to be used on social networking sites can best advertise their products (Aral and Walker, 2011) through viral marketing. Bagherjeiran et al. (2010) present a practical application where they use data from instant messaging logs at Yahoo! to improve online advertising targeting, and similarly Provost et al. (2009) show how to use browsing data to match groups of users who are socially similar. Tucker (2011b) explores how privacy controls mediate the effectiveness of advertising on Facebook. However, to our knowledge this is the first academic study of the effectiveness of social advertising.

Managerially, our results have important implications. Social advertising and use of online social networks is effective. However, when advertisers attempt to reinforce this social influence in ad copy, consumers appear less likely to respond positively to the ad. This is, to our knowledge, the first piece of empirical support for emerging managerial theories that

emphasize the need for firms to not appear too obviously commercial when exploiting social media (Gossieaux and Moran, 2010).

2 Field Experiment

The field experiment was run by a small non-profit that provides educational scholarships for girls to attend high school in East Africa. Without the intervention of this non-profit, and other non-profits like them, girls do not attend secondary school because their families prioritize the education of sons. Though the non-profit’s main mission is funding these educational scholarships, the non-profit has a secondary mission which is to inform young people in the US about the state of education for African girls. It was in aid of this secondary mission that the non-profit set up a Facebook page. This page serves as a repository of interviews with girls where they describe the challenges they have faced.

Table 1: Different Ad-Text Conditions

Condition	Ad-Text
Baseline	Help girls in East Africa change their lives through education.
Be like your friend	Be like your friend. Help girls in East Africa change their lives through education.
Don’t be left out.	Don’t be left out. Help girls in East Africa change their lives through education.
Your friend knows	Your friend knows this is a good cause. Help girls in East Africa change their lives through education.
Learn from your friend.	Learn from your friend. Help girls in East Africa change their lives through education.

To launch the field experiment, the non-profit followed the procedure described in ‘A/B Testing your Facebook Ads: Getting better results through experimentation’ (Facebook, 2010) which involved setting up multiple competing campaigns. Table 1 describes the different ad-copy for each condition. Each different type of ad-copy was accompanied by the same picture of an appealing secondary-school student who had benefited from their program.

The different ad conditions were broadly designed to cover the kinds of normative and informational social influence described by Deutsch and Gerard (1955); Burnkrant and Cousineau (1975).¹ We want to be clear that we do not argue that these advertising measures capture all types of social influence or are necessarily successful at distinguishing between the different types of social influence that are possible. The literature on social influence has emphasized that the underlying mechanism is nuanced and complex. Obviously, different types of social influence relate and interact in ways that cannot be teased apart simply with different wording. However, the variation in messages does allow us to study whether explicit advertising messages that attempt to evoke different types of social influence are effective in general.



Figure 1: Sample Ad

Figure 1 displays an anonymized sample ad for a social ad in the ‘be like your friend’ condition. The blacked-out top of the ad contained the non-profit’s name. The grayed-out bottom of the ad contained a supporter’s name, who had ‘liked’ the charity and was a Facebook friend of the person who was being advertised to. It is only with developments in technology and the development of appropriate automated algorithms that this automation

¹Other forms of social influence studied in the literature involve network externalities where there is a performance benefit to multiple people adopting (Tucker, 2008). However, that does not seem to be relevant here.

of the friend’s name when pertinent is possible.

These ad campaigns was targeted to three different groups as shown in Table 2. The first was a broad untargeted campaign for all Facebook users aged 18 and older in the US. The second group were people who had already expressed interest in other charities. These people were identified using Facebook’s ‘broad category targeting’ of ‘Charity + Causes.’ The third group were people who had already expressed an interest in ‘Education + Teaching.’ Previously, the charity had tried such reasonably broad targeting with little success and was hopeful that social advertising would improve the ads’ performance (Tucker, 2011b). For each of these non socially-targeted campaigns, we ran the baseline variant of the ad text which, as shown in Table 1, simply says ‘Help girls in East Africa change their lives through education.’ The non-profit could not run the other four conditions that refer to others’ actions, because federal regulations require ads to be truthful and they did not want to mislead potential supporters.

Then for each of these campaigns, the non-profit launched a socially targeted variant. These ads employed the Facebook ad option that meant that they were targeted only to users who were friends of existing fans of the charity. This also meant that they were eligible to receive a ‘social endorsement’ where the name of the friend was shown at the bottom of the ad as shown in Figure 1. These socially targeted ads displayed all five variants of the advertising message depicted in Table 1.

Table 3 describes the demographics of the roughly 1,500 fans at the beginning of the campaign. Though the initial fans were reasonably spread out across different age cohorts, they were more female than the average population, which makes sense given the nature of the charity. At the end of the experiment, the fans were slightly more likely to be male than before. The way that Facebook reports data means that we have access to the demographics only of the fans of the charity, not of those who were advertised to.

Table 2: Different Groups Targeted

Condition	Baseline: Only Shown Baseline text	Social Variant: Shown all 5 texts from Table 1
Untargeted	All people in US over age of 18	All people in US over age of 18 who are friends of the non-profit’s supporters
Charity	All people in US over age of 18 who state affinity with charities on their Facebook profile	All people in US over age of 18 who state affinity with charities on their Facebook profile who are friends of the non-profit’s supporters
Education	All people in US over age of 18 who state affinity with education on their Facebook profile	All people in US over age of 18 who state affinity with education on their Facebook profile who are friends of the non-profit’s supporters

Table 3: Demographics of the non-profit’s fans before and after the field experiment

Age	Before Experiment		After Experiment	
	Male	Female	Male	Female
18-24	5	13	8	14
25-34	5	14	6	14
35-44	6	17	6	16
45-54	3	13	3	13
55+	3	10	4	10
Total	22	67	27	67

The ‘Total’ row does not add up to 100% because fans who are below 18 years of age are omitted.

3 Data

The data that Facebook shares with advertisers is both anonymous and aggregate. This means that we cannot trace the effects of social advertising on the friends of any one individual. It also means that we cannot examine heterogeneity in the degrees of influence across individuals, as is studied, for example, by Godes and Mayzlin (2009) in their study of offline firm-sponsored communications. However, given that the central research question of the study is whether, on average, different types of social advertising are more effective, the aggregate nature of the data is sufficient. Table 4 reports daily summary statistics for the campaigns in our data. Over a 5-week period, there were 630 observations. There were 18 campaigns in total that consisted of a) The three baseline conditions that were demographically targeted to everyone, charity-lovers and education-supporters and used the baseline text, and b) The fifteen social ad conditions that had all the five different types of text, and socially targeted separately to everyone, charity-lovers and education-supporters.

Table 4: Summary Statistics

	Mean	Std Dev	Min	Max
Average Impressions	13815.7	13898.6	1	98037
Average Clicks	5.06	5.17	0	37
Connections	2.70	3.52	0	24
Unique Clicks	5.04	5.14	0	36
Click Rate	0.11	0.10	0	1.27
Click Rate (Multiple)	0.045	0.047	0	0.50
Cost Per Click (USD)	0.98	0.40	0.31	3.90
Cost Per 1000 views (USD)	0.52	1.37	0	24.5
Ad-Reach	6165.7	6185.0	1	60981
Frequency	2.32	0.82	1	9.70

630 observations of 18 ad variants at the daily level.

There are two click-through rates reported in Table 4. The first click-through rate is the proportion of people who clicked on an ad that day. The denominator here is the Ad-Reach measure that captures the number of people exposed to an ad each day. The

second click-through rate is per ad impression. We focus on the former in our econometric analysis, because impressions can be a function of person refreshing their page or using the back button on the browser or other actions which do not necessarily lead to increased exposure to the ad. We show robustness subsequently to using this click-through rate per impression measure. Due to the relatively small number of clicks, these click through rates are expressed as percentage points or sometimes as fractions of a percentage point. In our regression analysis we also use this scaling in order to make our coefficients more easily readable.²

The data also contains an alternative means of measuring advertising success. The connection rate measures the number of people who liked a Facebook page within 24 hours of seeing a sponsored ad, where the denominator is the ad's reach that day. We compare this measure to clicks in subsequent analysis to check that the click-through rate is capturing something meaningful. We also use the cost data about how much the advertiser paid for each of these ads in a robustness check.

²The data reassuringly suggests that there were only five occasions where someone clicked twice on the ads. Therefore, 99.8% of the click-through rate we measure captures a single individual clicking on the ad.

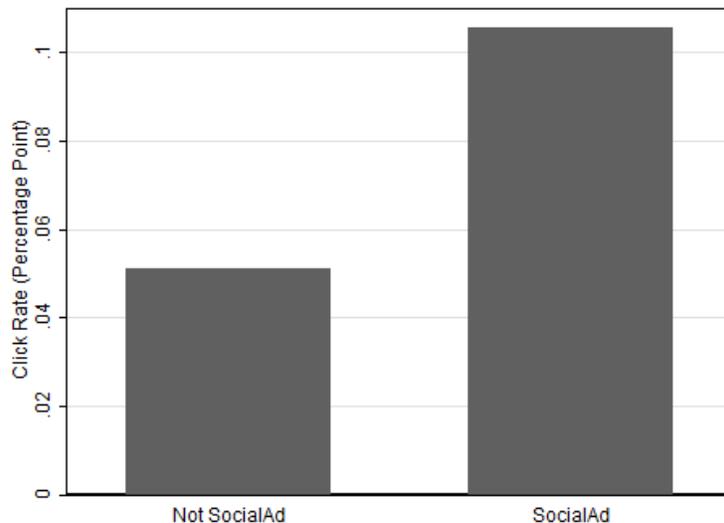


Figure 2: Social advertising is effective

4 Results

4.1 Does Social Advertising Work?

First, we present some simple evidence about whether social advertising is more effective than regular display advertising. Figure 2 displays the basic comparison of aggregate (that is, across the whole five-week period) click-through rates between non-socially-targeted ads and ads that were socially targeted. These are expressed as fractions of a percentage point. It is clear that social advertising earned far larger click-through rates.

The difference between the two bars is quite striking. To check the robustness and statistical significance of this relationship, we turn to econometrics. The econometric analysis is relatively straightforward because of the randomization induced by the field tests. We model the click-through rate of campaign j on day t targeted to demographic group k as:

$$ClickRate_{jt} = \beta SocialTargeting + Endorsement_j + \gamma_k + \delta_t + \epsilon_j \quad (1)$$

$SocialTargeting + Endorsement_j$ is an indicator for whether or not this campaign variance was socially targeted and displayed the endorsement. Since Facebook does not allow the testing of these different features separately, this is a combined (rather than separable) indicator. γ_k is a fixed effect that captures whether this was the untargeted variant of the ad. This controls for underlying systematic differences in how likely people within that target and untargeted segment were to respond to this charity. We include a vector of date dummies δ_t . Because the ads are randomized, δ_t and γ_k should primarily improve efficiency. We estimate the specification using ordinary least squares. Though we recognize that theoretically a click-through rate is bounded at one hundred since it is measured in percentage points, click-through rates in our data are never close to this upper bound.

Table 5 reports our initial results. Column (1) presents results for the simple specification implied by equation (1) but without the date and demographic controls. The point estimates suggest that social targeting and a friend’s endorsement increased the average daily click-through rate by around half. Column (2) repeats the analysis with the controls for date. It suggests that after controlling for date, the result holds. This is reassuring and suggests that any unevenness in how ads were served across days does not drive our results. It also suggests that our result is not an artifact of a failure of randomization. Column (3) adds an extra coefficient that indicates whether that campaign was untargeted rather than being targeted to one of the customer groups identified as being likely ‘targets’ by the non-profit - Educational and Charity supporters. It suggests that indeed, as expected, an untargeted campaign was weakly ineffective, though the estimate is not significant at conventional levels. We speculate that the apparent weakness of demographic targeting may be because target markets of charity and educational supporters is reasonably broad, and consequently may have contained many individuals who would not support an international charity.

An obvious question is what explains the success of social advertising. One explanation is that the endorsement of a friend is informative. Another explanation is that social targeting

uncovers people who will be more likely to be interested in their charity as they are similar, in unobserved ways, to their friends who are already fans of the charity. Manski (1993) pointed out that this particular issue of distinguishing homophily (unobserved characteristics that make friends behave in a similar way) from the explicit influence of friends on each other is empirically problematic.

Ideally, to address this we would simply randomize whether users saw the endorsement or not. However, Facebook’s advertiser interface does not allow that. What we can do is take advantage of the fact that sometimes ads are shown to people without the endorsement if that fan has selected a privacy setting which restricts the use of their image and name. The interface which users use to do this is displayed in Figure A1; all users do is simply select the ‘No One’ rather than the ‘Only my friends’ option. Of course, this will not represent perfect randomization. It is likely that the fans who select stricter privacy settings differ in unobserved ways from those who do not, and that therefore their social networks may differ as well. However, despite this potential for bias, this does represent a useful opportunity to try to disentangle the power of social targeting to enable homophily and the power of personal endorsements. Column (4) displays the results of a specification for equation (1) where the dependent variable is the conversion rate for these socially targeted but not socially endorsed ads. Here for ads that were being shown to friends, the click-through rate was only calculated for occasions when the endorsement was not shown. A comparison of Column (3) and Column (4) in Table 5 makes it clear the ads that were displayed to friends of fans but lacked a clear endorsement were less effective than those that had a clear endorsement. However, they were still measurably more effective than non-socially-targeted ads. It appears that, roughly, the endorsement accounted for less than half of the persuasive effect and the ability to use social networks to target the ad accounted for slightly more than half of such ads’ efficacy.

Columns (5) and (6) of Table 5 estimate the specification separately by whether the cam-

paign was targeted or untargeted. Though the point estimate for the targeted campaigns is higher, it is notable that social advertising improved the performance of both targeted and untargeted campaigns. Given the widely reported lack of efficacy of untargeted campaigns (Reiley and Lewis, 2009; Tucker, 2011b), the increase in effectiveness allowed by social advertising appears large for untargeted campaigns.

Table 5: Social Targeting and Endorsement is Effective

	All		Untargeted		Targeted	
	(1)	(2)	(3)	(4)	(5)	(6)
	Click Rate	Click Rate	Click Rate	No Endorsement Click Rate	Click Rate	Click Rate
SocialTargeting + Endorsement	0.0386*** (0.0123)	0.0385*** (0.0108)	0.0386*** (0.0125)		0.0297*** (0.00755)	0.0376*** (0.00927)
SocialTargeting				0.0287** (0.0143)		
Untargeted			-0.000275 (0.0122)	0.0132 (0.0166)		
Constant	0.0794*** (0.0116)					
Date Controls	No	Yes	Yes	Yes	Yes	Yes
Observations	630	630	630	630	210	420
Log-Likelihood	542.1	610.3	610.3	427.8	187.7	452.3
R-Squared	0.0221	0.212	0.212	0.119	0.317	0.228

OLS Estimates. Dependent variable is the percentage point of people who click on the ad. Dependent variable in Columns (3) for social ads is the percentage point daily click-through rate of ads that did not display the endorsement. Robust standard errors clustered at ad level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.2 Robustness

Table 6 checks the robustness of the finding that social targeting and endorsement are effective, to different definitions of the dependent variable. Column (1) reports the results of using a dependent measure which is the percentage click-through per impression. Again, we find that social advertising is more effective, though the effectiveness is less pronounced and less precisely estimated than before. This suggests that the appeal of social advertising is not necessarily enhanced by multiple exposure. It could also, of course, merely reflect noise introduced into the process by someone refreshing their browser multiple times.

The results so far suggest that consumer privacy concerns or the intrusiveness of such ads do not seem to outweigh the appeal of social advertising for consumers.³ There is always the possibility of course that people clicked on the ads because they were annoyed or wanted to understand more the extent of privacy intrusion rather than because the ads were actually effective. To explore this, we estimate a specification where the dependent measure was the proportion of clicks that became subscribers of the newsfeed. The results are reported in Column (2). We see that again social advertising appears to be more effective at encouraging Facebook users to take the intended action as well as simply clicking. This is evidence that people are not clicking on social ads due to annoyance at their intrusiveness but instead are clicking on them and taking the action the ads intend to encourage them to take. Untargeted ads are less likely to lead to conversions than those targeted at appropriate demographics. This makes sense - these people are being targeted precisely because they are the kind of people who have signed up for such news feeds in the past.

A final question is whether ads that are socially targeted and display endorsements are more expensive for advertisers, thereby wiping out their relative effectiveness in terms of return on advertising investment. We explore this in Column (3) of Table 6. There are

³This may be because Facebook users find it reassuring that they have control over their privacy (Tucker, 2011b) or because these ads, though narrowly targeted, are not overly visually intrusive (Goldfarb and Tucker, 2011).

several missing observations where there were no clicks that day and consequently there was no price recorded. In Column (3), we report the results of a specification where our explanatory variables is the relative price per click. The results suggest that advertisers pay less for these clicks that are socially targeted. This suggests that Facebook is not charging a premium for this kind of advertising. Though Facebook shrouds in secrecy the precise pricing and auction mechanism underlying their advertising pricing, this result would be consistent with a mechanism whereby advertisers pay less for clicks if they have higher click-through rates. In other words, prices paid benefit from an improved ‘quality-score’ (Athey and Nekipelov, 2011). The results also suggest that advertisers pay less for demographically untargeted clicks which is in line with previous studies such as Beales (2010).

Table 6: Social Advertising is Effective: Checking robustness to different dependent variables

	(1)	(2)	(3)
	Click Rate (Multiple)	Clicks to Connections Rate	Cost Per Click (USD)
SocialTargeting + Endorsement	0.0108** (0.00501)	0.433*** (0.0997)	-0.195*** (0.0480)
Untargeted	0.00526 (0.00582)	-0.321*** (0.0768)	-0.177*** (0.0520)
Date Controls	Yes	Yes	Yes
Observations	630	554	559
Log-Likelihood	1086.5	-467.5	-129.0
R-Squared	0.150	0.163	0.426

OLS Estimates. Dependent variable is the click-through rate (expressed as a fraction of a percentage point) for impressions in Column (1). Dependent variable in Column (2) is the clicks to conversions rate.

Dependent variable in Column (3) is cost per click.

Robust standard errors clustered at ad level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.3 What Kind of Social Advertising Messages Work?

We then go on to explore what kind of advertising message works in social ads. We distinguish between ads that rely simply on the Facebook algorithm to promote social influence by featuring the automated endorsement at the bottom of their ad, and ads that explicitly refer to this endorsement in their ad copy.

Table 7: Social Advertising is Less Effective if an Advertiser is Too Explicit

	All (1)		(2)		(3)		Untargeted (4)		Targeted (5)	
	Click Rate		Click Rate		No Endorsement	Click Rate	Click Rate		Click Rate	
SocialTargeting + Endorsement	0.0577*** (0.0139)		0.0571*** (0.0113)				0.0498** (0.0245)		0.0527*** (0.0130)	
SocialTargeting					0.0333** (0.0168)					
SocialTargeting + Endorsement \times Explicit	-0.0287*** (0.00886)				-0.0101 (0.0124)		-0.0303* (0.0167)		-0.0284** (0.0124)	
Untargeted					0.0161 (0.0169)					
SocialTargeting + Endorsement \times Don't be left out					-0.0136 (0.0115)					
SocialTargeting + Endorsement \times Be like your friend					-0.0189* (0.01000)					
SocialTargeting + Endorsement \times Learn from your friend					-0.0378*** (0.0115)					
SocialTargeting + Endorsement \times Your friend knows					-0.0429*** (0.0144)					
Date Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	630	630	630	630	630	210	420	210	420	420
Log-Likelihood	615.4	618.1	618.1	429.5	189.6	461.0	461.0	189.6	461.0	461.0
R-Squared	0.225	0.232	0.232	0.124	0.329	0.260	0.260	0.329	0.260	0.260

OLS Estimates. Dependent variable is the percentage points of people who click on the ad. Dependent variable in Columns (3) adjusted for social ads so that is the percentage point daily click-through rate of ads that did not display the endorsement.

Robust standard errors clustered at ad level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

We use the additional binary indicator variable $Explicit_j$ to indicate when the advertiser uses a message that evokes social influence explicitly in their ad copy, in addition to the social endorsement automated by the Facebook algorithm. This covers all the non-baseline conditions described in Table 1. We interact this with the $SocialTargeting + Endorsement_j$, meaning that $SocialTargeting + Endorsement_j$ now measures the effect of the baseline effect, and the interacted variable measures the incremental advantage or disadvantage of mentioning the friend or the potential for social influence in the ad.

Column (1) of Table 7 reports the results. The negative coefficient on the interaction between $Explicit$ and $SocialTargeting + Endorsement_j$ suggests that explicit reference to a social influence mechanism in the ad affected the performance of the ad negatively. That is, when the advertiser themselves were explicit about their intention to harness social influence, it backfires. Further, the large point estimate for $SocialTargeting + Endorsement_j$ suggests that the baseline message is even more effective than the estimates of Table 5 suggested.

Column (2) in Table 7 reports the results of a specification where we break up $Explicit$ by the different types of ‘social influence’-focused advertising messages featured in Table 1. It is striking that all measures are negative. It is also suggestive that the one message that was not statistically significant and had a smaller point estimate than the others did not refer to the friend explicitly but instead referred obliquely to the friend’s action. This is speculative, since the point estimate here is not statistically different from the others due to its large standard error.

Column (3) repeats the exercise for the click-through rate for the ads that did not display an endorsement that we investigated in Table 5. Since these ads did not display the friend’s name at the bottom, it should not be so obvious to a viewer that the firm is explicitly trying to harness the social influence that results from the friend being a fan of the charity. We recognize that there may of course be some confusion at the mention of a friend when no name is displayed, but this confusion should work against us rather than for us. In this case,

we do not see a negative and significant effect of the ‘Explicit’ advertising message which referred to a friend. This suggests that it was the combination of the friend’s name and the mention of social influence which was particularly off-putting. The results in Column (3) suggest that what is damaging is the combination of an advertiser making it explicit they are trying to harness social influence and the algorithmic social advertising message.

We next explored whether this finding that attempts by advertisers to explicitly harness social influence in their ad text damaged the effectiveness of social advertising differed by the target group selected. Column (4) presents the results for the campaign that was targeted at friends of fans who were simply over 18 years old and based in the US. Column (5) presents the results for the group of users whom the charity selected as being in the target ‘demographic’ groups for the campaign - that is users whose Facebook profile revealed their support for other educational and charitable causes. What is striking is the similarity of the estimates for the efficacy of social advertising and the damage done by the advertiser being overly explicit about social influence across Columns (4) and (5). Again, similar to the results reported in Table 5 social advertising appears to be able to offer as nearly as large a lift to ad efficacy for an untargeted population as a targeted one.

4.4 Behavioral Mechanism

We then collected additional data to help rule out alternative explanations of our finding that the explicit mention of social influence was undesirable in social ads.

One obvious potential explanation is that what we are measuring is simply that people are unaware that what they are seeing is actually an ad, rather than part of Facebook. When a non-profit uses a message such as ‘Be like your friend’ then it becomes obvious that this is an ad, and people respond differently. To test this, we persuaded the non-profit to run a subsequent experiment that allowed us to explicitly tease this apart. In this experiment we compared the performance of ads that said ‘Please read this ad. Help girls in East Africa

change their lives through education.’, and ads that simply said ‘Help girls in East Africa change their lives through education.’.⁴

If it is was the case that Facebook users were simply mistaking socially targeted ads for regular content and the explicit appeals to social influence stopped them making this mistake, we would expect to also see a negative effect of wording that made it clear that the message was an ad. However, it appears that adding ‘Please read this ad’ actually helped ad performance, which suggests that it was not the case that Facebook users were simply mistaking socially targeted ads for content if there is no explicit message. Obviously, though, the sample size here is very small, making more definitive pronouncements unwise.

Table 8: Efficacy of Social Advertising Varies by Group Targeted

	Knowledge (1) Click Rate	Fashion (2) Click Rate	(3) Click Rate
SocialTargeting + Endorsement	0.0312* (0.0160)	0.0194 (0.0208)	0.0182 (0.0208)
SocialTargeting + Endorsement × Explicit	0.0114 (0.0288)	0.0376* (0.0221)	
SocialTargeting + Endorsement × Don’t be left out			0.0449* (0.0254)
SocialTargeting + Endorsement × Be like your friend			-0.00448 (0.0218)
SocialTargeting + Endorsement × Learn from your friend			0.0172 (0.0254)
SocialTargeting + Endorsement × Your friend knows			0.127** (0.0584)
Date Controls	Yes	Yes	Yes
Observations	20	60	60
Log-Likelihood	55.43	91.77	103.7
R-Squared	0.916	0.267	0.508

OLS Estimates. Dependent variable is the percentage point of people who click on ad that day. Robust standard errors (Standards errors are not clustered at the advertising campaign level, unlike in previous tables, as there are too few clusters). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Another alternative explanation for our findings is that the messages referring to the friend were poorly-written or unappealing. To test whether this was the case, we selected

⁴Recent research has questioned the use of the imperative in advertising copy, which is why we used ‘please’ (Kronrod et al., 2012)

an alternative set of users whom might be expected to react in an opposite way to potential presumptions of social influence. Specifically, the charity agreed to run test conditions identical to those in Table 1 for the people who expressed affinity with ‘Fashion’ goods on their Facebook profiles. The Fashion category of users were chosen because typical models of social influence have focused on fashion cycles (Bikhchandani et al., 1992). These models emphasize the extent to which people who participate in Fashion cycles receive explicit utility from conformity, even when this conformity is provoked by a firm. In other words, they may find advertiser-endorsed social influence more persuasive and advertiser attempts at emphasizing the power of social influence more acceptable than the general population does. This group of users exhibits a very different pattern to that exhibited by the general population. They appear to respond somewhat positively to social advertising, though this estimate is imprecise and the point estimate is smaller than for the other conditions. However, strikingly, they reacted particularly positively to advertising messages that emphasized social influence and the actions of the friend in the ad copy. In other words, social advertising for this group worked even when the advertiser explicitly embraced the potential for social influence. This result suggests that there may be heterogeneity in consumer responses to the wording of social advertising messages depending on their previous consumption patterns. This is evidence against an alternative explanation for our results in Table 7 based on these advertising messages which explicitly refer to the potential for social influence being confusing or overly wordy, since they were effective for this group of Fashion fans. In general, the results of Tables 7 and 8 suggest that there is heterogeneity in distaste for advertiser attempts to harness social influence given previous consumption patterns, but that for the average person the effects are negative.

5 Implications

How helpful is data on social relationships when it comes to targeting and delivering advertising content? This paper answers this question using field test data of different ads on the large social networking site Facebook. We find evidence that social advertising is indeed very effective.

This is important, as for the past few years social networking websites have often been dismissed by advertisers as venues for ‘paid media’, that is, paid advertising. Instead, the emphasis was on ‘earned’ or organic media whereby social networks were venues for organic word of mouth. This dismissal of paid advertisements was echoed in the popular and marketing press with headlines such as ‘Online Social Network and Advertising Don’t Mix’ and ‘Facebook Ad Click-Through Rates Are Really Pitiful’ (Joel, 2008; Barefoot and Szabo, 2008). Our results suggest, however, that as social advertising develops this will change swiftly. In particular, social networks will be able to exploit their considerable inherent network effects to enlarge their share of advertising dollars.

Strikingly, we find that the average Facebook user appears to find social advertising as done by the standard Facebook algorithm appealing. However, when advertisers attempt to emulate or reinforce this social influence, consumers appear less likely to respond positively to the ad. Speculatively, the results suggest that intrusive or highly personal advertising is more acceptable if done algorithmically by a faceless entity such as a computer than when it is the result of evident human agency. Very speculatively, there is perhaps a parallel with users of web-based email programs accepting an algorithm scanning their emails to serve them relevant ads when the interception of emails by a human agent would not be acceptable.

Our results suggest that social advertising works well for both targeted and untargeted populations, which may mean that social advertising is a particularly useful technique when

advertising to consumers outside the product's natural or obvious market segment since their are less obvious ways of targeting in these settings. The majority of this efficacy appears to be because social targeting uncovers unobserved homophily between users of a website and their underlying receptiveness to an advertising message.

There are of course limitations to our study. First, the non-profit setting may bias our results in ways that we cannot predict. Second, the aims of the non-profit also means the outcome measure we study is whether or not people sign up to hear more about the non-profit, rather than studying the direct effect of advertising on for-profit outcomes such as customers making purchases. Third, we studied this advertising at a time when Facebook was just launching and promoting its social advertising features. It is not clear if the results will be as strong if the advertising market becomes saturated with social ads. Notwithstanding these limitations, we believe that this paper makes a useful contribution in terms of documenting when social advertising is useful and when it is not.

References

- Algesheimer, R., S. Borle, U. M. Dholakia, and S. S. Singh (July/August 2010). The impact of customer community participation on customer behaviors: An empirical investigation. *Marketing Science* 29(4), 756–769.
- Aral, S. and D. Walker (September 2011). Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Science* 57(9), 1623–1639.
- Athey, S. and D. Nekipelov (2011). A structural model of sponsored search advertising auctions. *Mimeo, Berkeley*.
- Bagherjeiran, A., R. P. Bhatt, R. Parekh, and V. Chaoji (2010). Online advertising in social networks. In B. Furht (Ed.), *Handbook of Social Network Technologies and Applications*, pp. 651–689. Springer US.
- Bakshy, E., J. M. Hofman, W. A. Mason, and D. J. Watts (2011). Everyone’s an influencer: quantifying influence on Twitter. In *Proceedings of the fourth ACM international conference on Web search and data mining, WSDM ’11*, New York, NY, USA, pp. 65–74. ACM.
- Barefoot, D. and J. Szabo (2008, April). Facebook ad click-through rates are really pitiful. *Friends with Benefits: A social media handbook blog*.
- Beales, H. (2010). The value of behavioral targeting. *Mimeo, George Washington University*.
- Bikhchandani, S., D. Hirshleifer, and I. Welch (1992, October). A theory of fads, fashion, custom, and cultural change in informational cascades. *Journal of Political Economy* 100(5), 992–1026.

- Burnkrant, R. E. and A. Cousineau (1975, December). Informational and normative social influence in buyer behavior. *Journal of Consumer Research* 2(3), 206–15.
- Comscore (2011). Comscores 2011 social report: Facebook leading, microblogging growing, world connecting. *White Paper*.
- Corcoran, S. (2009, December 16). Defining earned, owned and paid media. *Forrester Research*.
- Deutsch, M. and H. B. Gerard (1955). A study of normative and informational social influences upon individual judgment. *The Journal of Abnormal and Social Psychology* 51(3), 629–636.
- Facebook (2010, September). A/B testing your Facebook ads: Getting better results through experimentation'. *Technical Report*.
- Godes, D. and D. Mayzlin (2009). Firm-Created Word-of-Mouth Communication: Evidence from a Field Test. *Marketing Science* 28(4), 721–739.
- Goldfarb, A. and C. Tucker (2011, May). Online display advertising: Targeting and obtrusiveness. *Marketing Science* 30, 389–404.
- Gossieaux, F. and E. Moran (2010). *2010*. McGraw-Hill.
- IAB (2009). IAB social advertising best practices. *Interactive Advertising Bureau*.
- Joel, M. (2008, February 8). Online social networks and advertising don't mix. *Six Pixels of Separation, Twist Image*.
- Kronrod, A., A. Grinstein, and L. Wathieu (2012, January). Go green! Should environmental messages be so assertive? *Journal of Marketing* 76, 95–102.

- Manski, C. F. (1993, July). Identification of endogenous social effects: The reflection problem. *Review of Economic Studies* 60(3), 531–42.
- Provost, F., B. Dalessandro, R. Hook, X. Zhang, and A. Murray (2009). Audience selection for on-line brand advertising: privacy-friendly social network targeting. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '09, New York, NY, USA, pp. 707–716. ACM.
- Reiley, D. and R. Lewis (2009). Retail advertising works! measuring the effects of advertising on sales via a controlled experiment on Yahoo!". Working Paper, Yahoo! Research.
- Trusov, M., R. E. Bucklin, and K. Pauwels (2009, September). Effects of word-of-mouth versus traditional marketing: Findings from an internet social networking site. *Journal of Marketing* 73, 90–102.
- Tucker, C. (2008). Identifying formal and informal influence in technology adoption with network externalities. *Management Science* 54(12), 2024–2038.
- Tucker, C. (2011a). Ad virality and ad persuasiveness. *Mimeo, MIT*.
- Tucker, C. (2011b). Social Networks, Personalized Advertising, and Privacy Controls. *Mimeo, MIT*.
- Tucker, C. and J. Zhang (2011). How does popularity information affect choices? A field experiment. *Management Science* 57(5), 828–842.
- Zubcsek, P. and M. Sarvary (2011). Advertising to a social network. *Quantitative Marketing and Economics* 9, 71–107.

Facebook Ads

Ads and friends

Everyone wants to know what their friends like. That's why we pair ads and friends—an easy way to find products and services you're interested in, based on what your friends share and like. Learn more about [social ads](#).

Here are the facts:

- Social ads show an advertiser's message alongside actions you have taken, such as liking a Page
- Your privacy settings apply to social ads
- We don't sell your information to advertisers
- Only confirmed friends can see your actions alongside an ad
- If a photo is used, it is your profile photo and not from your photo albums

Here's an example of a Facebook Ad:

 <p>Denver Sushi The best sushi in Denver. Try our daily lunch specials for \$9.95. Fan our page for special offers.</p> <p>Like · Catherine Tucker likes this.</p>	 <p>Denver Sushi The best sushi in Denver. Try our daily lunch specials for \$9.95. Fan our page for special offers.</p>
---	--

This setting only applies to ads that we pair with news about social actions. So, independent of this setting, you may still see social actions in other contexts, like in Sponsored Stories or paired with messages from Facebook. You can learn more about how social ads, Sponsored Stories, and messages from Facebook work in the [Help Center](#).

Pair my social actions with ads for

Figure A1: Control interface for switching off Social Advertising