

Social Advertising

Catherine Tucker[‡]

October 20, 2011

Abstract

Social networks, such as Facebook, have grown at an explosive rate. The quantity of information they provide about users' social connections could potentially be very valuable for marketers. For example, this data can be used to improve paid advertising, either by being used to target ads based on underlying social networks or by using social relationships to tailor the content of advertising. This paper explores the effectiveness of such 'social' advertising using data from field tests of different ads on Facebook. We find evidence that using consumers' social networks to target ads and personalizing advertising with information about friends' actions are both very effective, relative to customary methods of targeting. However, advertising like this is less effective if the advertiser explicitly emphasizes the potential for social influence in the text of their ad. Speculatively, this suggests that one of the main advances of social networks for advertising may be to harness the power of social networks by using computer-based algorithms to serve tailored ads that would otherwise feel presumptuous to consumers.

*MIT Sloan School of Management, MIT, Cambridge, MA. and NBER.

[‡]Thank-you to Google for financial support and to an anonymous non-profit for their cooperation. Thank-you to Avi Goldfarb and Juanjuan Zhang for valuable comments. All errors are my own.

1 Introduction

An individual's purchasing behavior is affected by social influence. Typically, such social influence has been organic and not easily stimulated or enhanced through actual advertising. There have been attempts to evoke social networks in advertising copy. For example, one famous public service announcement is, 'Friends don't let friends drive drunk'; 'Over 2,000,000 customers can't be wrong' was a persuasive advertising slogan for Dell's Inspiron brand. However, advertising of this kind has not been personalized or targeted to reflect a target consumer's actual social network or the actions of their own friends.

This has changed with the evolution of online social networks such as Facebook.com and LinkedIn.com. Both websites 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 both to target advertising and potentially to engage their audience.

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 and 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. Comparing the performance of these ads that contained the name of the fan and those that were simply targeted to that fan's social network suggests that their effectiveness stems almost equally from the ability to select a more appropriate targeting group using social relationships and because of 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 ultimate conversion and is more cost-effective.

We then ask what kind of advertising copy an advertiser should use when attempting to harness social influence through advertising. Social influence can, broadly, take two forms (Deutsch and Gerard, 1955; Burnkrant and Cousineau, 1975). *Normative* social influence is where consumers gain utility from emulation or conformity if they do the same as their friends. *Informational* social influence describes a process of observational learning and informational cascades where consumers absorb information about unobserved product quality from observing their friends' choices.¹ Through randomized field tests, we investigate the effectiveness of advertisers deliberately intending stimulate both types of social influence explicitly in their advertising copy through including a statement that encourages the viewer to, for example, be like their friend.

We find that, surprisingly, attempts by advertisers to explicitly harness or refer to a friend's actions in their ad copy are rejected by consumers. This rejection is reasonably uniform across different wording, though slightly less severe for ads that were slightly less personal. To investigate whether it was simply bad advertising copy, we also 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, which suggests that it was not simply that the message was badly communicated. Instead, the evidence suggests that the majority of individuals react poorly when an advertiser is seen to explicitly attempt to propagate social influence.

¹For certain categories of products, there can also be network externalities, where there is a direct performance benefit for the consumer if they conform to the product choice of their friends (for example, Skype). See Tucker (2008) for a description of how social networks affects network externalities for such products.

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, the major purpose the social network serves is to transmit word of mouth, and online advertising and word of mouth are seen as substitutes. Godes and Mayzlin (2009) use field-test data to show that firm-created word of mouth is most effective between acquaintances and for less loyal customers. Our results demonstrate something new - social advertising may be most effective when targeted towards groups of consumers who are not obvious customers for the product.

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. Tucker (2011) 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 on such websites.

Managerially, our results have two important implications. It appears that the average Facebook user reacts positively to social advertising as served by the standard Facebook algorithm. However, when advertisers attempt to emulate or reinforce this social influence, consumers appear less likely to respond positively to the ad. The only exception to this are Facebook users who have already signalled a willingness to follow ‘Fashion’ and conform to social influence. This suggests that part of the future power of social networking websites will be as venues that make social targeting more acceptable to a broader group of users, when it would more usually, if not an automated feature of the site, be seen as presumptuous.

Speculatively, this could reflect a greater acceptance by consumers of social targeting when it is obviously done by a computer algorithm than when it is done with the complicity of a more identifiable entity such as an advertiser.

Second, our our results suggest that social advertising may be most useful when advertising to consumers who would otherwise be outside the product’s natural or obvious market segment. This is important because the majority of advertising technologies in the web such as search advertising are limited by whether a consumer takes a certain action, like searching for a specific product. Similarly, behavioral targeting relies on a consumer taking a certain action like visiting a certain category of website. Therefore, social advertising is potentially very attractive to advertisers as a means of extending their ability to target to a broader group of consumers.

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 their main mission is focused on 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 (and often videos) where they describe the challenges they have faced.

The ad campaign was originally targeted to three different groups. 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 targeting would improve the ads' performance (Tucker, 2011).

Table 1: Different Conditions

Condition	Target Group	Ad-Text
Baseline: Targeted	Friends of Supporters	Help girls in East Africa change their lives through education.
Normative Social Influence (Emulation)	Friends of Supporters	Be like your friend. Help girls in East Africa change their lives through education.
Normative Social Influence (Exclusion)	Friends of Supporters	Don't be left out. Help girls in East Africa change their lives through education.
Informational Social Influence (Unstated)	Friends of Supporters	Your friend knows this is a good cause. Help girls in East Africa change their lives through education.
Informational Social Influence (Explicit)	Friends of Supporters	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). 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 student who had benefited from their program.

The different ad conditions are broadly designed to follow Deutsch and Gerard (1955); Burnkrant and Cousineau (1975)’s distinction between normative and informational social influence. The idea behind the ‘informational’ conditions is broadly to evoke a feeling that there is knowledge or learning to be gained from that friend’s subscription to the non-profit’s news feed. The two conditions differ in how explicit they make this. The idea behind the ‘normative’ conditions is to evoke a sense that there could be utility from conformity or from emulating the friend. One condition emphasizes more the positive benefits of conformity while the other condition emphasizes the positive benefits of emulation.

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 very 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 displays an anonymized sample ad for the campaign for the Emulation condition. The blacked-out top of the ad contained the non-profit’s name. The greyed 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 this latter element that is unique and different about social advertising.




Figure 1: Sample Ad

Table 2 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.

After the experiment had run for a few weeks, and empirical regularities were becoming apparent, the charity agreed to expand the demographics that they were targeting in order to help tease apart the effects. Specifically, the charity agreed to run test conditions for the people who expressed affinity with ‘Fashion’ goods. The idea behind testing this additional category is that according to previous research, such people are more likely to be subject to social influence.

Table 2: Fan Demographics

Age	Before Experiment		After Experiment	
	Male	Female	Male	Female
18-24	5.2	13	8.1	14
25-34	4.9	14	6.4	14
35-44	5.6	17	6.0	16
45-54	3.3	13	3.3	13
55+	3.4	10	3.5	10

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. However, given the central research question of the study which is whether, on average, different types of social advertising are more effective, the aggregate nature of the data is sufficient. Table 3 reports daily summary statistics for the campaigns in our data.

Table 3: Summary Statistics

	Mean	Std Dev	Min	Max
Average Impressions	13780.0	14854.2	1	110427
Average Clicks	4.86	5.26	0	37
Unique Clicks	4.85	5.24	0	36
Click Rate	0.12	0.18	0	2.27
Click Rate (Multiple)	0.042	0.047	0	0.50
Cost Per Click	0.98	0.42	0.31	3.90
Cost Per 1000 views	0.51	1.40	0	24.5
Ad-Reach	6198.9	6919.8	1	72283
Frequency	2.82	2.99	1	27.3

620 observations of ad campaigns at the daily-level.

The data reassuringly suggests that there were only 5 occasions where someone clicked twice on the ads. Therefore, 99.8% of cases capture a single individual clicking on the ad. The Reach measure is the number of individual people who saw ads. This is different than the Impressions measure, which captures the number of times an ad was shown on a page the user was browsing.

There are two forms of click-through rates reported in Table 3. Due to the relatively small number of clicks, these 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. The first type of click-through rates is the proportion of people who clicked on an ad that day. The second is the click-through rate 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. However, we do show robustness to using the click-through rate per impression measure. The data also contains an alternative means of measuring advertising success. The Connection rate measures ‘The number of people who liked your Facebook page within 24 hours of seeing this sponsored story or ad’ relative to 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.




Figure 2: Social advertising is effective

4 Results

4.1 Does Social Advertising Work?

Figure 2 displays the basic comparison of average daily 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 and highly statistically significant ($t = 6.50$). To check the robustness 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 at demographic group k as:

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

$SocialTargeting_j$ is an indicator for whether or not the campaign was socially targeted.

γ_k is a vector of fixed effects for the different demographic groups targeted. These control for underlying systematic differences in how likely people within that target 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, but they also control for the fact the field experiment was ran in a staggered manner over multiple weeks. We estimate the specification using ordinary least squares. Though we recognize that theoretically a click-through rate is bounded at one, click-through rates in our data are never close to this upper bound. This means there is unlikely to be any distortions from assuming a linear functional form for a dependent measure bounded between 0 and 1 (or, given the percentage point conversion, 0 and 100). Following evidence presented by Bertrand et al. (2004), we cluster standard errors at the ad-campaign level to avoid artificially understating our standard errors due to the fact that the experiment ran over multiple days.

Table 4 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 increased by almost half the average daily click-through rate. 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 the fact the experiment was run in two waves over the course of several weeks is not driving our result. 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 targeted at one of the customer groups initially selected by the non-profit - Educational and Charity supporters. It suggests that indeed, as expected, a demographically targeted campaign was more effective at attracting click-throughs than untargeted campaigns. However, it was not the case that such targeting was as effective as social targeting. Column (4) repeats the estimation but adds indicator variables for each of the different demographic groups. This largely controls for the potential for one (particularly persuadable) demographic group to have unexpectedly received more

Socially Targeted ads than others. However, little changes when we add these controls, suggesting that randomization between the groups was effective.

Table 4: Social Targeting is Effective

	(1)	(2)	(3)	(4)
	Click Rate	Click Rate	Click Rate	Click Rate
SocialTargeting	0.0491*** (0.0141)	0.0593*** (0.0166)	0.0687*** (0.0160)	0.0601*** (0.0162)
Target Demo			0.0267** (0.0125)	
Constant	0.0756*** (0.00692)			
Date Controls	No	Yes	Yes	Yes
Demo Controls	No	No	No	Yes
Observations	620	620	620	620
Log-Likelihood	184.7	229.2	230.4	236.6
R-Squared	0.0114	0.391	0.394	0.406

OLS Estimates. Dependent variable is the percentage points of people who click on the ad. Robust standard errors clustered at ad-level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5 checks the robustness of the finding that social targeting is 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 targeting is more effective, though the efficacy 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, just reflect noise introduced into the process by someone refreshing their browser multiple times.

An obvious question is what explains the success of social targeting. 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 and disentangle the power of social targeting to enable homophily and the power of personal endorsements. Column (2) 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. A comparison of Column (2) and Column (4) in Table 4 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. It appears that, roughly, the endorsement accounted for half of the persuasive effect and the ability to use social networks to target the ad accounted for the other half of such ads’ efficacy.

The results in Table 4 suggest that social targeting is effective despite privacy concerns. An obvious worry with any form of social advertising is that by being too intrusive it becomes ineffective. This was explicitly addresses by Provost et al. (2009) who study practical methods of using social targeting that preserve the complete privacy of its users. However, the results in this study suggest that consumer privacy concerns or the intrusiveness of such ads

do not seem to outweigh the appeal of social targeting for consumers. This may be because Facebook users find it reassuring that they have control over their privacy Tucker (2011) or because these ads, though narrowly targeted, are not overly visually intrusive (Goldfarb and Tucker, 2011). 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. To explore this, we estimate a specification where the dependent measure was the proportion of conversions to the people who saw the ad, that is the number of people who after clicking on the ad became fans of the charity. The results are reported in Column (3). We see that again social targeting appears to be more effective, though the size of the effect is smaller, suggesting that as one would expect there is attrition between conversion and clicks. This is reassuring evidence that people are not clicking on social targeted 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.

A final question is whether ads that are socially targeted are simply more expensive for advertisers, thereby wiping out their relative effectiveness in terms of return on advertising investment. We explore this in Columns (4) of Table 5. There are several missing observations where there were no clicks that day and consequently there was no price recorded. In Column (4) 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 social targeted. This suggests that Facebook is not charging an artificial 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).

Table 5: Social Targeting is Effective: Checking robustness to different dependent variables

	(1)	(2)	(3)	(4)
	Click Rate (Multiple)	Non-Social Click Rate	Connection Rate	Cost Per Click
SocialTargeting	0.0122** (0.00607)	0.0242*** (0.00817)	0.0278*** (0.00917)	-0.171* (0.0925)
Date Controls	Yes	Yes	Yes	Yes
Demo Controls	Yes	Yes	Yes	Yes
Observations	620	620	620	523
Log-Likelihood	1076.2	854.7	833.0	-209.3
R-Squared	0.550	0.449	0.481	0.886

OLS Estimates. Dependent variable is the click-through rate (expressed as a fraction of a percentage point) for impression in Column (1). Dependent variable is the fraction of percentage points of people who click on the ad who did not see the endorsement in Column (2). Dependent variable is the fraction of percentage points for the connection-conversion rate in Column (3). Dependent variable is cost per click in Column (4).

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

4.2 What Kind of Social Advertising Messages Work?

We then go on to explore what kind of advertising message an advertiser should use when using social targeting. 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 (as shown in Figure 1, and ads that explicitly and intentionally refer to this endorsement in their ad copy.

We use the additional binary indicator variable *Intentional_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 therefore measures the incremental advantage or disadvantage of mentioning the friend or the potential for social influence in the ad. Column (1) of Table 6 reports the results. Surprisingly, it appears that the explicit reference to the potential for social influence 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. Column (2) in Table 6 reports the results of a specification where we explicitly break up *Intentional* 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 striking that the one message which was not statistically significant and had a smaller point estimate than the others did not refer to the friend explicitly but instead referred to the friend’s action only obliquely, though of course the point estimate here is not statistically different from the others.

Column (3) repeats the exercise for the non-social click through rate investigated in Table 5. This specification is essentially a falsification check. 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. In this case, we do not see a negative and significant effect of the ‘intentional’

advertising message which referred to a friend. Column (4) repeats this falsification exercise for the different advertising messages and again, as would be expected, finds no significant differences. The results in Columns (3) and (4) suggest that what is harmful is the combination of an advertiser making it explicit they are trying to harness social influence and the algorithmic social advertising message.

Table 6: Social Targeting is Less Effective if an Advertiser is Too Explicit

	(1) Click Rate	(2) Click Rate	(3) Non-Social Click Rate	(4) Non-Social Click Rate
SocialTargeting	0.0906*** (0.0271)	0.0902*** (0.0271)	0.0268** (0.0119)	0.0265** (0.0120)
Intentional		-0.0485** (0.0235)	-0.00403 (0.0110)	
Normative: Exclusion		-0.0389 (0.0269)		0.00119 (0.0115)
Normative: Emulation		-0.0412* (0.0236)		0.00318 (0.0119)
Informational: Explicit		-0.0597** (0.0249)		-0.0126 (0.0109)
Informational: Unexplicit		-0.0532* (0.0270)		-0.00702 (0.0175)
Date Controls	Yes	Yes	Yes	Yes
Demo Controls	Yes	Yes	Yes	Yes
Observations				
Log-Likelihood	241.2	241.6	854.9	856.5
R-Squared	0.414	0.415	0.449	0.452

OLS Estimates. Dependent variable is the percentage points of people who click on the ad in Columns (1)-(2). Dependent variable is the percentage point daily click-through rate of ads that did not display the endorsement in Columns (3)-(4).

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

We next explored whether this finding that attempts by advertisers to explicitly harness

social influence in their ad text damaged the effectiveness of social targeting differed by the target group selected. Column (1) 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. Given the widely reported lack of efficacy of untargeted campaigns (Reiley and Lewis, 2009; Tucker, 2011), the increase in effectiveness allowed by social targeting appears large. Though the estimates are not precise, the extent to which advertisers being explicit about their intent to harness social influence in their ads cancels the effectiveness of social targeting is striking for this untargeted group. Column (2) 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 targeting and the damage done by the advertiser being overly explicit about social influence across Columns (1) and (2). That is, social targeting appears to be able to offer as large a lift to ad efficacy for an untargeted population as a targeted one. This is particularly notable because, as reported in Table 4 the ads targeted at the right demographics were, as expected, already more effective.

As described earlier, at the request of the authors the charity added an additional set of campaigns where the charity advertised to consumers who were considered to be in the ‘Fashion’ category by Facebook, based on their interests and likes displayed in their Facebook profile. The Fashion category of users were chosen because typical models of social influence focus 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. Interestingly, this group of people exhibits a very different pattern to that exhibited by the general population or the consumers in the target market for the charity. They appear to

respond somewhat positively to social targeting, 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. Importantly, this result also helps us rule out alternative explanations for the results in Table 6. Specifically, it allows us to rule out explanations such as the 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 devotees. In general, the results of Table 7 suggests 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.

Table 7: Efficacy of Social Targeting Varies by Group Targeted

	Untargeted (1)	Targeted (2)	Fashion (3)
	Click Rate	Click Rate	Click Rate
Intentional	-0.0765* (0.0402)	-0.0415** (0.0196)	0.0510** (0.0232)
SocialTargeting	0.0903* (0.0545)	0.0606** (0.0280)	0.0126 (0.0218)
Date Controls	Yes	Yes	Yes
Log-Likelihood	4.917	198.2	75.82
R-Squared	0.358	0.648	0.677

OLS Estimates. Dependent variable is the percentage points 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). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

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 initial evidence that social advertising is indeed very effective.

This is important, as for the past few years social networking websites have 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.

Strikingly, we find that the average Facebook user appears to find social targeting 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. The only exception to this are Facebook users who have already signalled a willingness to follow ‘fashion’ and conform to social influence. This suggests that part of the future power of social networking websites will be as venues that make social targeting, a powerful but also potentially presumptuous and consequently off-putting form of advertising, acceptable to a broader group of users.

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 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 also suggest that social targeting is a particularly useful technique when advertising to consumers outside the product's natural or obvious market segment. This is important because the majority of advertising technologies in the web such as search advertising have been limited, since they only focus on a narrow set of people who are at a certain (usually late) stage in the purchase process. Therefore, an effective advertising method that offers a greater scope than existing targeting online is very attractive for advertisers.

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

- 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.
- Barefoot, D. and J. Szabo (2008, April). Facebook ad click-through rates are really pitiful. *Friends with Benefits: A social media handbook Blog*.
- Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics* 119(1), 249–275.
- 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.
- 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. 30, 389–404. *Marketing Science*.
- 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*.
- 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. (2011). Social Networks, Personalized Advertising, and Privacy Controls. *Mimeo, MIT*.

Zubcsek, P. and M. Sarvary (2011). Advertising to a social network. *Quantitative Marketing and Economics* 9, 71–107.

Facebook Ads

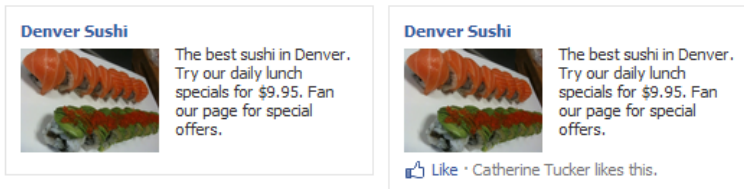
Ads and friends

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:



Denver Sushi
The best sushi in Denver. Try our daily lunch specials for \$9.95. Fan our page for special offers.

Denver Sushi
The best sushi in Denver. Try our daily lunch specials for \$9.95. Fan our page for special offers.

Like · Catherine Tucker likes this.

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