Social Networks, Personalized Advertising, and Privacy Controls

Catherine Tucker

Social Networks, Personalized Advertising, and Privacy Controls

Catherine Tucker*

April 10, 2011

*Catherine Tucker is Douglas Drane Career Development Professor in IT and Management and Assistant Professor of Marketing at MIT Sloan School of Management, MIT, Cambridge, MA; 617 252-1499; cetucker@MIT.EDU. I thank the Time Warner Cable Research Program on Digital Communications for financial support. I also thank Alessandro Acquisti, Avi Goldfarb, Cait Lambert, Alex Marthews, Martin Peitz, Ken Wilbur and seminar participants at Harvard, New York University, the University of Florida and the University of Mannheim. All mistakes are mine alone.

Social Networks, Personalized Advertising, and Privacy Controls

Abstract

This paper investigates how internet users' perception of control over their personal information affects how likely they are to click on online advertising. The paper uses data from a randomized field experiment that examined the relative effectiveness of personalizing ad copy with posted personal information on a social networking website. The website gave users more control over their personally identifiable information in the middle of the field test. The website did not change how advertisers used data to target and personalize ads. After the introduction of improved privacy controls, users were twice as likely to click on personalized ads. There was no comparable change in the effectiveness of ads that did not make explicit that they used private information when targeting. The increase in effectiveness was larger for ads that used more unique private information to personalize their message.

Keywords: Privacy, Online Advertising, Social Networks

Introduction

The Internet and communication technologies revolution has dramatically increased firms' ability to target digital advertising to specific consumers, and to use consumer information to personalize the content of the advertising. However, as online advertising becomes personalized, firms run the risk that customers will find the advertising intrusive and invasive of their privacy, and that 'reactance' will lead them to resist the ad's appeal (White et al., 2008). 'Reactance' is a motivational state when consumers resist something they find coercive by behaving in the opposite way to the one intended (Brehm, 1966; Clee and Wicklund, 1980; Brehm, 1989). This sets up a tension for firms who want to use the huge amounts of data at their disposal to improve advertising, but who want to minimize the potential of consumer resistance.

Nowhere has this tension been more pronounced than on social networking websites like Facebook and MySpace. Social networking websites now account for 23 percent of online display advertising. They have also collated a huge amount of personal data from their users and offer advertisers proprietary ad networks that push the boundaries of tailored advertising. Consumers might see personalized ad content on such sites as more appealing and tied to their interests, but they might conversely see it as 'not only creepy, but off-putting' if they felt that the firm had violated their privacy (Stone, 2010).[1]

To reassure customers about their use of customer data, social networking sites have pioneered new technologies that allow consumers explicit control over how much information about them is publicly available. Theoretically, this could minimize the potential for reactance and improve the performance of online advertising, because behavioral research has emphasized the importance of consumer perceptions of control in mediating reactance (Taylor, 1979). This is the case even if the controls introduced are only tangentially related

[1]As suggested by Stone (2010), 'What a marketer might think is endearing, by knowing a little bit about you, actually crosses the line pretty easily.'

to the area where reactance may be invoked (Rothbaum et al., 1982; Thompson et al., 1993). For example, cancer patients are more likely to comply with restrictive treatment regimes if they are given perceived control over another aspect of their medical care. However, there is always the risk that such introducing privacy controls might sensitize users to privacy concerns, increasing the likelihood of reactance and making advertisers who try to use personal information more unpopular. This paper assesses how these new technologies for giving customers control over their personally identifiable information influence the effectiveness of online advertising on social networking websites.

We use data from a randomized field experiment conducted by a US-based nonprofit to optimize their advertising campaigns on Facebook. These campaigns were shown to 1.2 million Facebook users. The nonprofit's aim was to raise awareness of its work improving education for women in East Africa. The nonprofit randomized whether they explicitly personalized the ad copy to match the user's profile. For example, sometimes the text of the ad explicitly mentioned a celebrity whom the user had listed that they were a fan of. On other occasions, the nonprofit showed the same group of fans an ad that was deliberately generic in the text and made no explicit mention of the celebrity.

In the middle of the field experiment, Facebook announced a large and well-publicized shift in their privacy policy. The aim was to reassure users, given mounting media criticism, about how their data were used, by giving them more control over their privacy settings and the extent to which their personally identifiable data could be tracked or used by third parties. This change did not, however, affect the underlying algorithm by which advertising was displayed, targeted and personalized, since the advertising platform used anonymous data. The nonprofit had not anticipated there would be such a change when it launched its field test of the ads. However, the fact that this occurred mid-way through the field experiment is valuable for measuring the effect of a change in privacy policies on advertising effectiveness, while circumventing the usual endogeneity issues.

We have data on the number of times each ad was shown to a unique user and the number of times it was clicked on for each ad for a five-week period spanning the introduction of the new privacy controls. Empirical analysis of both campaign-level and individual-level click-through data suggests that personalized advertising was over twice as effective at attracting users to the nonprofit's Facebook page, after the shift in Facebook policy that gave users more control over their personal information. There was no significant change in advertising that was shown to the same people but used a generic message over the period. This is to be expected, because such ads do not make clear to consumers whether their private information is being used to target.

Identification comes from the assumption that there were no underlying changes in user behavior that coincided with the introduction of privacy controls but were not directly attributable to the introduction of these controls. To ensure the robustness of this assumption, we check that there was no significant change in the ads shown, the user composition of Facebook, use of the website, or advertiser behavior during the period we study. We also control for the amount of publicity surrounding privacy issues at the time of the introduction of privacy controls. Controlling for media attention either by including direct controls or excluding the days at the height of the media storm leads us to estimate a smaller, though still economically significant effect. Last, we show that there was no change in how likely people were to signup for the nonprofit's news feed, suggesting that our result is not an artifact of stimulated curiosity.

To explore the underlying mechanism, we build on existing research that documents that 'reactance' to personalized advertising is greatest when the information used is more unique (White et al., 2008). We explored whether the positive effect of improved privacy controls was greatest for ads that used more unique information. Though some celebrities in our test, like Oprah Winfrey, have as many as two million fans on Facebook, some of the celebrities or undergraduate institutions were unusual enough that their potential reach was only in the thousands. We found that personalization was relatively more effective for personalized ads that used unusual information after privacy controls were enhanced. This provides evidence that indeed consumers were concerned that the information being used in the ads was simply too personal to be used in an ad without a corresponding sense of control over their data.

We confirm this interpretation with evidence from an online survey that tested consumer reactions to different online ads that were associated with either unique or not at all unique private information, in contexts where respondents either felt they had control over their personal information or not. The results from this experiment confirm our earlier findings and, by explicitly measuring stated reactance, provide support for a behavioral mechanism where reactance is reduced for highly personal advertising if consumers perceive they have control over their privacy.

Contribution

These findings contribute at four levels. To our knowledge, this is the first paper that studies an instance where a firm gave web users better control over how their personal information was shared and how that affected advertising outcomes. The finding that there are positive effects for advertisers, in this instance, of addressing users' privacy concerns, is potentially useful to advertising-supported websites. Turow et al. (2009) found that 66 percent of Americans do not want marketers to tailor advertisements to their interests. Fear of such resistance has led advertisers to limit their tailoring of ads (Lohr, 2010). However, our results suggest that there are benefits to the advertising-supported internet of reassuring users explicitly about how their private information is shared and used.

This has implications for public policy governing advertising. Currently, proposed regulations governing behavioral advertising in the US are focused around the mechanics of how firms implement opt-in and opt-out use of cookies and other tracking devices. Previous empirical research suggests that this approach, by limiting the use of data by firms, reduces ad effectiveness (Goldfarb and Tucker (2011b)). By contrast, the results in this paper show that in this setting when a social networking website allowed customers to choose how personally identifiable information about them is shared and used, there was no negative effect on advertising performance. The current staff-discussion draft of US privacy legislation proposed by Representatives Boucher and Stearns exempts individually-managed preference profiles (P.17, Sec. 3(e)). This provision may be an important way of ensuring that the advertising-supported internet can continue to thrive.

On the academic side, this paper's focus on advertising complements research that has focused on more general questions of information sharing and privacy in social networks (Acquisti and Gross, 2006; Golder et al., 2007; Caverlee and Webb, 2008). Early research on privacy tended to simply describe privacy as a matter of giving users control over their data (Miller, 1971). However, more recent research in information systems has challenged this and has shown how individual-level control over even tangential information can mediate privacy concerns, even if access to the data remains unchanged (Spiekermann et al., 2001; Xu, 2007; Brandimarte et al., 2010). The psychology literature provides a theoretical explanation for these findings. When consumers feel that firm behavior is intrusive of their privacy, this can lead to reactance (Clee and Wicklund, 1980). It has been documented, particularly in the health literature, that one way to overcome such reactance is to reinforce perceptions of control even if the controls do not actually give full control over the domain under threat (Rothbaum et al., 1982; Thompson et al., 1993). Therefore, firms are able to reduce reactance to potentially intrusive marketing activities by improving perceived consumer control.

The paper also contributes to the online advertising literature. It appears that personalizing ads using user-disclosed information in the ad copy increases their appeal if accompanied by appropriate privacy controls. This was studied from a theoretical perspective by Anand and Shachar (2009), who pointed out that the signaling power of a targeted ad in the traditional ad-signaling framework (Kihlstrom and Riordan, 1984), could be strengthened by personalizing the ad, making consumers more likely to assume there is a match between them and the product. The majority of the empirical work on targeting and social networks has studied offline methods (see for example Manchanda et al. (2008)). 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, 2010) through viral marketing. Outside of social networks, Goldfarb and Tucker (2011a) have shown that privacy concerns can influence ad effectiveness.

There have been no studies, however, to our knowledge, that examine advertising on social networks by external firms. This is an important topic, because social networking sites are attractive media venues that are growing rapidly in importance. They have a youthful and passionate following: The average Facebook user in the United States spent 6.5 hours on Facebook over the course of December 2009, which was more than twice as long as the next leading web brand (Nielsen, 2010). Facebook doubled its U.S. audience from 54.5 million visitors in December 2008 to 111.9 million visitors in December 2009, and now accounts for 7% of all time spent online in the U.S (Lipsman, 2010); worldwide, its membership passed 500 million in July 2010. In late 2010, as shown in Table A-1, Facebook became the top display advertising venue on the internet, accounting for 23 percent of all display advertising impressions within the US, totalling over a billion impressions each year (Cormier, 2010). In 2011, it is projected to receive $4 billion in advertising revenues (Kerr, 2011). However, social networking websites have previously been perceived as being problematic venues for advertising because of extremely low click-through rates (Holahan, 2007). This research suggests that if such sites are successful at reassuring consumers that they are in control of their privacy, firms can use personalization of ads to generate higher click-through rates.

Institutions and Data

The Nonprofit

The nonprofit running the experiment provides educational scholarships in East Africa that enable bright girls from poor families to go to or stay in high school. Part of the nonprofit's mission involves explaining its work in Africa to US residents and also engaging their enthusiasm and support for its programs. In order to do this, the nonprofit set up a Facebook 'page' which explained its mission, and also allowed people who were interested to see photos, read stories and watch videos about the girls who had been helped by the program.

To attract people to become fans of its Facebook page, the nonprofit started advertising using Facebook's own advertising platform. Initially, they ran an untargeted ad campaign which displayed an ad in April 2010 to all users of Facebook that live in the US and are 18 years and older. This campaign experienced a very low click-through rate and attracted fewer than five new 'fans' to the website. The disappointing nature of this campaign led them to want to see if they could engage further with their potential supporters by both targeting and personalizing ad content.

Randomized Campaign

The nonprofit designed two separate campaigns with two separate target populations. The aim of the campaign was to encourage users to click on the ad and become a fan of the nonprofit's website. The first target population were college graduates from 20 small liberal arts colleges that had a reputation of emphasizing the benefits of education for the community. Facebook started as a college-based social network, so it explicitly facilitates the identification of such graduates, and most users indicate what educational institutions they have attended and whether they are a current student or a graduate.

The second target population were Facebook users who had expressed appreciation for 19 celebrities and writers who in the past had made statements supporting the education of girls in Africa or African female empowerment in general.[2] Examples could be Oprah Winfrey, who has set up a girls' school in South Africa, or Serena Williams, who was a supporter of 'Build African Schools.' The target group was identified by whether they mention they 'like' such a person in their likes or interests section on their Facebook profile. We refer to whether someone was a fan of these 19 celebrities or 20 undergraduate institutions as the information which constitutes 'the targeting variable'. Using the Facebook advertising interface, we also verified that there was very little overlap in fans across these different groups.

However, it was unclear to the nonprofit whether they should also personalize the ad content that these users saw. They thought that personalization might improve their ad's appeal, but they also did not want their ad to be unattractively intrusive or make potential supporters feel that their privacy had been violated. In order to establish whether Facebook user data should be used merely to target ads, or should in addition be used to personalize the content of the advertising appeal, they decided to experiment with two different ad formats. Table 1 summarizes the different conditions used. In the personalized condition, the ad explicitly mentioned the undergraduate institution or the celebrity's name. In the targeted but non-personalized case, the ad was similar in content but did not explicitly mention the undergraduate institution or the celebrity's name. In both cases, the baseline or 'non-personalized' message was not completely generic, but instead alluded to some kind of very broad user characteristic. Therefore, it would be precise to interpret our estimates as reflecting the incremental benefit of personalized ad-content that has specific and concrete personal information relative to ad content that uses non-specific and non-concrete information. In each case, the ad was accompanied by the same picture of a girl who had been helped by the program. Based on the work of Small and Verrochi (2009), this girl had a slightly mournful expression. Figure A-1 contains a screenshot of the ad-design interface.

[2]The nonprofit is eager to protect the privacy of its supporters, and consequently has asked the authors to not reveal either the names of the celebrities or of the schools that were used in this advertising campaign.

Table 1: Campaign appeals in different conditions

Information used to target ad	College	Interest
Personalized	As a [undergraduate institution name] graduate you had the benefit of a great education. Help girls in East Africa change their lives through education.	As a fan of [name of celebrity] you know that strong women matter. Help girls in East Africa change their lives through education.
Non-Personalized	You had the benefit of a great education. Help girls in East Africa change their lives through education.	You know that strong women matter. Help girls in East Africa change their lives through education.

In addition to these two campaigns, the nonprofit also continued to use as its baseline, an untargeted campaign which reached out to all adult US Facebook users simultaneously. This provided an additional baseline control for advertising effectiveness over the course of the study. The text of this baseline and untargeted ad read "Support [Charity Name]. Help girls in East Africa change their lives through education." This ad and the two targeted campaigns were restricted to Facebook users who live in the US, and were 18 years and older. The charity set a daily maximum spending cap on advertising campaigns that was significantly below the $250-a-day maximum spending cap mandated by Facebook. It also agreed to pay at most $0.50 for each click produced by the different advertising campaigns.

The Introduction of Improved Privacy Controls

What was unique and potentially valuable about this field experiment was that on May 24 2010 (after the field experiment was planned and initiated and the first data collected), Mark Zuckerberg, the CEO of Facebook, announced that the company would be simplifying and clarifying their privacy settings as well as rolling back some previous changes that had made Facebook users' information more public. Studying this change was not the initial purpose of the randomized field experiment, but it fortuitously presented a unique opportunity to study how a change in user privacy controls in social networking sites can change consumer responses to advertising, since the nonprofit tested the ads using the same randomization technique before and after the change in the privacy-control interface.

The background to the introduction of an improved privacy interface was that Facebook had been heavily criticized because its privacy settings were very granular and difficult to access. For example, Bilton (2010) pointed out that the 5,850 words of Facebook's privacy policy were longer than the United States Constitution, and that users wanting to manage their privacy settings had to navigate through 50 settings with more than 170 options. As detailed by Table A-2 in the appendix, Facebook had previously acted to reduce the amount of control users had over their data and had attracted negative publicity for doing so. As well as bad press, Facebook faced legal challenges. In December 2009, ten privacy groups filed a complaint with the Federal Trade Commission[3] over changes to Facebook's privacy policy, which included default settings that made users' status updates available potentially to all Internet users, as well as making users' friend lists publicly available.

There were three major components to Facebook's change in privacy interface. The first was that all privacy settings were aggregated into one simple control. Users no longer had to deal with 170 granular options. As depicted in appendix Figure A-2, this interface was far more approachable and easily adjustable than before. Second, Facebook no longer required users' friends and connections to be visible to everyone. Third, Facebook made it easier to opt out with a single click from third-party applications from accessing users' personal information. Generally, these changes were received favorably. For example, the chairman of the American Civil Liberties Union, Chris Conley, wrote 'The addition of simplified options (combined with the continued ability to fine-tune your settings if you wish) and user control over Facebook's 'connections' are significant improvements to Facebook's privacy.'

This change in privacy settings did not change how the banner ads that were served on Facebook were targeted, or whether advertisers could use user information to personalize ads. Display advertising was treated separately because, as Facebook states, 'Facebook's ad targeting is done entirely anonymously. If advertisers select demographic targeting for their

[3]http://epic.org/privacy/inrefacebook/EPIC-FacebookComplaint.pdf.

ads, Facebook automatically matches those ads to the appropriate audience. Advertisers only receive anonymous data reports.' To reassure advertisers that the change would not adversely affect them, Facebook sent out an email to its advertisers saying that 'this change will not affect your advertising campaigns' (The full letter is reproduced in the appendix.) This means that though users were given control over how much information was being shared and the extent to which they were being tracked by third parties, the actual mechanism by which the ads tested were targeted and served did not change.

Data

We obtained daily data from the nonprofit on how well each of the ads performed for the duration of the experiment. There were 79 different ad campaigns for which we obtained daily data on the number of times they were shown and the number of clicks. In total these ads were shown to 1.2 million users and they received 1,995 clicks. When a user clicked on the ad, they were taken to the nonprofit's Facebook page. These data spanned 2.5 weeks on either side of the introduction of privacy controls on May 28, 2010. We also check robustness to this time-span in Table 3.

This data included the number of unique impressions (that is, the number of users the ad was shown to) and the number of clicks each ad received. Each of these clicks came from a unique user. It contains information on the date that click was received but does not the time. It also includes data on the cost to the nonprofit per click and the imputed cost per thousand impressions. As shown in Figure A-1, Facebook also offers advertisers an estimate of the potential ad-reach of such targeting when they design their ads. This is the number of Facebook users whom Facebook estimated could be in the target segment for any targeted ad-campaign. We use this ad-reach data in our subsequent regressions to explore the behavioral mechanism driving our results. To protect the privacy of the nonprofit's supporters, we did not receive information about the backgrounds or identities of those who chose to like it, or on any of their actions after they made that choice. We also do not have information about whether these users did indeed change their privacy settings.

Table 2 reports the summary statistics. The average number of clicks relative to ad impressions is small, at two-tenths of one percent. This is even smaller when looking at the daily level, since many campaigns received no clicks on a given day, inflating the appearance of low click-through rates. We use both aggregate and daily measures of click-through rates in our regressions, and find qualitatively similar results. However, this is similar to rates reported by other advertisers for Facebook ads. In their provocatively-titled piece 'Facebook Ad Click-Through Rates Are Really Pitiful', Barefoot and Szabo (2008) reported average click-through rates between .01% and 0.06%.

Table 2 also reports summary statistics that we constructed for two indices, based on data from Google Trends about the number of searches for Facebook and privacy and Google News about the number of newspapers that had stories that contained the words 'privacy' and 'Facebook'. Google Trends only reports search volumes in terms of indices rather than giving aggregate search data, so for comparability we also converted the number of stories reported on Google News into an index, where everything is rebased relative to the largest number of news stories.

The nonprofit considers the campaign to have been an immense success, especially given the relatively small cost of the trial (less than $1,000). In their most recent fundraising campaign, around 6 percent of revenues from new donors came directly from their Facebook page. Compared to its peer nonprofits, it now has a far broader and deeper social media presence, with just under 1500 fans following its updates and news.

Analysis

Figure 1 displays the average click-through rate for each campaign before and after the introduction of improved privacy controls. Ads that personalized their content appeared to

Table 2: Summary Statistics

	Mean	Std Dev	Min	Max
Average Impressions	15892.7	63274.2	337	551783
Average Clicks	25.3	53.7	0	374
Average Cost Per Click	0.38	0.096	0.11	0.50
Cost per 1000 views	0.095	0.12	0	0.39
Ad-Reach (000000)	0.095	0.21	0.00098	0.99
Aggregate Click-Through Percentage	0.17	0.23	0	1.37
News Stories Index for Facebook and Privacy	54.48	34.78	11	100
Google Trend Index for Searches about Facebook and Privacy	59.91	27.95	24	100

Campaign-level data. 79 Different Campaigns (78 campaigns based on 39 different targeting variables each with personalized and targeted variants. 1 untargeted campaign)

greatly increase in effectiveness after the introduction of improved privacy controls. This change was highly significant (p-value=0.0047). The effects of targeting ads without personalizing their content before and after the introduction of improved privacy controls were not significantly different (p-value=0.407). There appears to be little change in the effectiveness of the untargeted campaign, though of course with only one campaign it is impossible to assess statistical significance when simply comparing a single before and after period. Analysis of click-through rates at the daily level suggests that there was no statistically significant change in the effectiveness for untargeted ads after the introduction of improved privacy controls.

Figure 2 examines whether there were any differences for the campaigns targeted to undergraduate institutions and celebrities. It is evident that on average the celebrity-focused campaign was more successful at attracting clicks. However, it appears clear that there was a similar incremental jump in the effectiveness of personalized ads after the introduction of improved privacy controls for both kinds of targeting variable.

Figure 1 suggests that the personalization of ads was more effective after Facebook facilitated users' taking control of their personal information. To check the robustness of this result, we also performed regression analysis. This allows us to assess the statistical

Figure 1: Comparison in Click-Through Rates Before and After

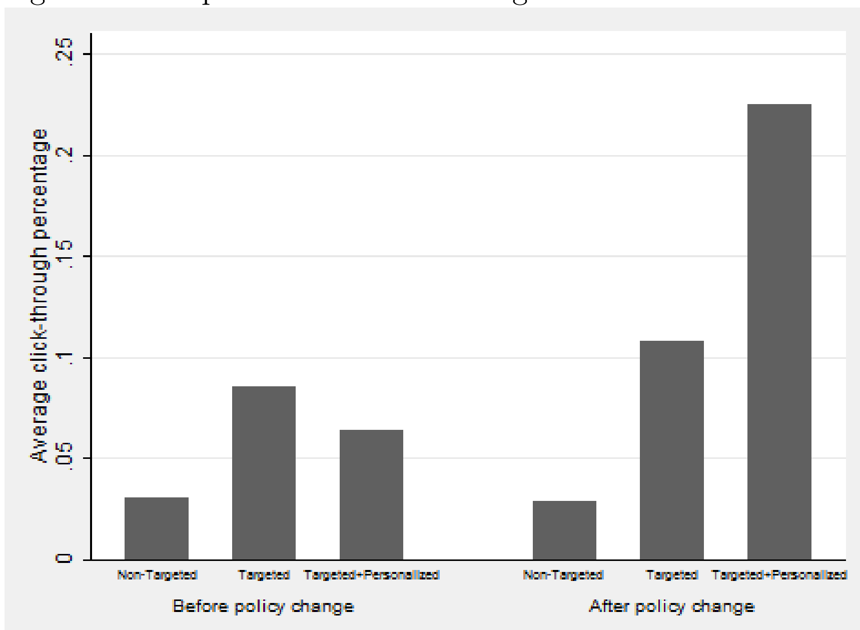

Figure 2: Comparison in Click-Through Rates Before and After

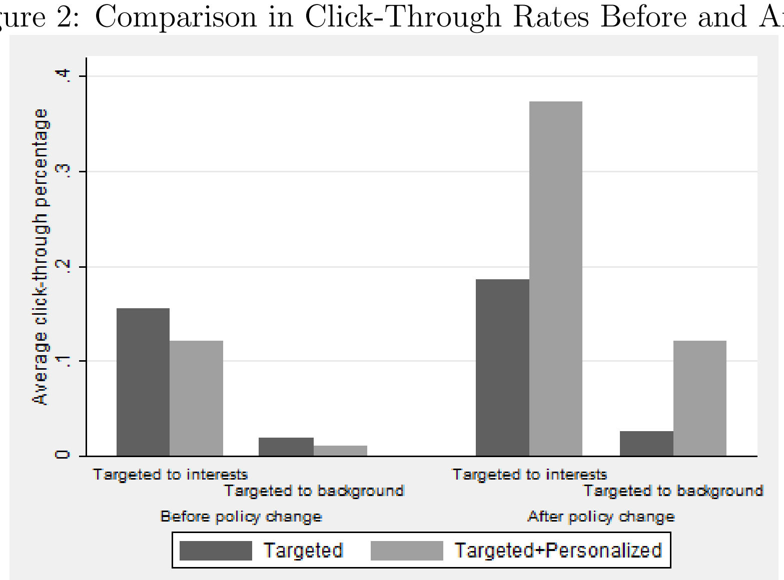

significance of our results in various ways and control for media coverage.

We model the click-through rate $ClickRate_{jt}$ for ad j on day t in the following manner:

$$\begin{aligned} ClickRate_{jt} = \beta Personalized_j \times PostPolicy_t + \alpha Personalized_j + \theta MediaAttention_{jt} \\ + \gamma_k + \delta_t + \epsilon_j \end{aligned} \quad (1)$$

$Personalized_j$ is an indicator variable which is equal to one if the ad contained personalized content matched to the variable on which it was targeted, and zero if there was no personalized content. $PostPolicy_t$ is an indicator variable equal to one if the date was after the privacy-settings policy change took place, and zero otherwise. The coefficient β captures the effect of their interaction. θ captures the effect of various controls we introduce to allow the effectiveness of personalized advertising to vary with media attention. γ_k is a vector of 39 fixed effects for the 20 different undergraduate institutions and each of the 19 celebrities 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. These are collinear with $PostPolicy_t$, which means that $PostPolicy_t$ is dropped from the specification. Because the ads are randomized, δ_t and γ_k should primarily improve efficiency. We estimate the regression using ordinary least squares. 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 we have panel data.

Table 3 presents our results which incrementally build up to the full specification in equation (1). Column (1) is our initial simplified specification. The crucial coefficient of interest is $Personalized \times PostPolicy$. This captures how an individual exposed to a personalized ad responds differently to a personalized ad after Facebook's change in privacy policy, relative to an ad shown to the same people that had generic wording. It suggests a positive and significant increase in the performance of personalized ads relative to merely targeted ads after the introduction of enhanced user privacy controls. The magnitude of our estimates suggest that the click-through rate increased by 0.024, relative to an average baseline click-through rate of 0.007 for personalized ads before the introduction of improved privacy controls. The negative coefficient $Personalized$ which is marginally significant suggests that, prior to the change in privacy settings, personalized ads were less effective than ads that did not use personalized ad copy.

This empirical analysis uses a short time window of 5 weeks. This means that it is unlikely that there is some long-run trend, for example increasing user acceptance of ad personalization or 'habituation' to privacy concerns, that drives the results. To show robustness to an even shorter window, we repeated our estimation data for 10 days from Day 13 to Day 22 (5 days before and 5 days after) around the introduction of improved privacy controls. We use a specification similar to that in Column (2) of Table 3. The results, reported in Column (2) of Table 3, were positive but larger than for the full period.[4]

[4]We ran a falsification check where we split a similar 10-day window in the pre-change period in half, and we found no evidence of any significant change in preferences for personalized advertising.

Table 3: Initial Results

	Initial	10-Day Window	Not 10-Day Window	Controls News	Controls Privacy Search	Full
	(1)	(2)	(3)	(4)	(5)	(6)
Personalized × PostPolicy	0.0236**	0.0554***	0.0149**	0.0243**	0.0204**	0.0162**
	(0.0102)	(0.0208)	(0.00711)	(0.00994)	(0.00956)	(0.00464)
Personalized	-0.0119*	-0.0112	-0.0141***	-0.0772	0.190	0.381
	(0.00627)	(0.0115)	(0.00464)	(0.122)	(0.325)	(0.199)
News Articles				-0.0465**		-0.00101
				(0.0214)		(0.0550)
Personalized Ad × News Articles				0.0129		0.0538
				(0.0248)		(0.0667)
Google Searches					-0.105	-0.0433
					(0.122)	(0.173)
Personalized Ad × Google Searches					-0.0462	-0.152
					(0.0750)	(0.123)
Date Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Targeting Variable Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2730	780	1950	2730	2730	2730
R^2	0.060	0.118	0.044	0.060	0.060	0.061

OLS Estimates. Dependent variable is percentage daily click through rate.
Robust standard errors clustered at ad-level. * $p < 0.10$, ** $p < 0.05$,*** $p < 0.01$
$PostPolicy_t$ is collinear with the date fixed effects and dropped from the specification.

Column (3) of Table 3 reports from a specification that includes all the data not included in this 10-day window (Day 1-Day 12 and Day 23-Day 35). The result is still positive and reasonably large, but is smaller than in Column (2), which is to be expected given the average effect measured in Column (1). One explanation is that the introduction of improved privacy controls was particularly salient in this 10-day window due to the amount of media coverage, meaning that people were more sensitized to personalized advertising. We explore this further in the next three columns when we include controls for news stories and general 'buzz' about the introduction of improved privacy controls.[5]

In Column (4), we add an additional interaction which controls for an index, reported on a scale between 0 and 100, that reflects the number of news stories each day returned by a query on 'Google News' for stories that contained the words 'Facebook' and 'Privacy'. In line with the idea that the results reported in Column (2) were larger because of the media buzz surrounding the introduction of improved privacy controls, the key interaction between $Personalized \times PostPolicy$ is smaller in magnitude, though still statistically and economically significant. Of course, while news stories capture some of the idea of general salience, they do not necessary reflect the extent to which news about Facebook and privacy concerns were being processed and acted on by Facebook users. To explore this, we used an additional control that captures the number of daily searches using the terms 'Facebook and Privacy' on Google as reported by the 'Google Trends' index, which is reported on a scale between 0 and 100. The key interaction $Personalized \times PostPolicy$ in Column (6) is again smaller in magnitude when we control for changes over time using this measure of salience of Facebook.

Column (6) reports results from our full specification which combines both these controls. The point estimate of 0.016 is similar in size to that reported in Column (4) where

[5] The results are also robust if we exclude the period where the service was rolled out and the days spanning announcement and implementation.

we exclude the 10-day window immediately around the introduction of improved privacy controls, suggesting that media attention did inflate the effect we measured in Column (2). It is still an economically significant increase relative to the average baseline click-through rate for personalized ads before the introduction of improved privacy controls of 0.007.

There is a relatively low R^2 across all specifications. This low level of explanatory power is shared by much of the online advertising literature (Reiley and Lewis, 2009; Goldfarb and Tucker, 2011a). One possible explanation is that consumers are skilled at avoiding looking at online advertising when viewing webpages, introducing measurement error (Dreze and Hussherr, 2003) and requiring researchers to assemble large datasets to measure effects precisely.

To try to measure how the introduction of improved privacy controls affected an individual's likelihood of clicking on an ad, we also estimate an individual-level logit model. One advantage of an individual-level model is that we can include the untargeted campaign in our regressions as the baseline. Rather than one observation of a click-through rate of the untargeted campaign which is collinear with the targeting group fixed effects, there are hundreds of thousands of observations of how individuals responded to that campaign.

We model the probability that an individual i clicks on ad j on day t as:

$$\begin{aligned} Clicked_{ijt} = I(&\beta_1 Personalized_j \times PostPolicy_t + \beta_2 Targeted_j \times PostPolicy_t \\ &+\alpha_1 Personalized_j + \alpha_2 Targeted_j + \theta MediaControls_{jt} + \gamma_k + \delta_t + \epsilon_j) \end{aligned} \tag{2}$$

Equation (2) is similar to Equation (1), except for the inclusion of a new indicator variable $Targeted_j$. $Targeted_j$ is an indicator variable for whether the ad was targeted, but had no attempts at personalization. For such ads, it would have been difficult for the consumer to know why they received that ad.

As explained by Ai and Norton (2003), in a logit model, interpretation of interaction terms is not straightforward, as they are a cross-derivative of the expected value of the dependent variable, meaning that neither the sign nor significance of the coefficient necessarily reflects the true marginal effect. This is a particular issue for three-way interactions. To address this, we estimated a logit model and used these logit estimates to calculate marginal effects while taking into account the fact that there were cross-derivatives in the specification. Table 4 reports the results of these logit models and associated marginal effects. Column (2) presents results for marginal effects that are directionally similar to those in Column (1) of Table 3, even though we are now studying behavior at an individual level. Not controlling for media effects, personalized ads performed worse than non-personalized ads before the policy, but performed twice as well after the policy. There was no significant shift in the efficacy of targeted ads before and after the policy. In contrast to the campaign-level results reported in Column 6 of Table 3, the controls for media coverage and search activity surrounding Facebook and privacy appear to have had less of a dampening effect on the interaction between $Personalized_j \times PostPolicy_t$. However, there are of course multiple lower-order interactions, which are not precisely estimated, that have to be considered in order to understand the overall effect.

Though $PostPolicy$ is not separately identified from the date fixed effects (that is, the trend for untargeted ads), the raw data suggests that click-through rates for untargeted ads doubled, or that for every 10,000 impressions, prior to the introduction of improved privacy controls 3 people clicked, and afterwards 6 people clicked. However, this change is not statistically significant due to the very small baseline probability, especially after controlling for media activity.

Table 4: Individual-level logit results

	Initial	Full		
	(1)	(2)	(3)	(4)
	Logit	Marginal Effects	Logit	Marginal Effects
Personalized Ad × Postpolicy	0.255**	0.00268**	0.235**	0.00290***
	(0.101)	(0.000248)	(0.102)	(0.000792)
Targeted Only Ad × Postpolicy	-0.186*	-0.0000586	-0.275	-0.00123
	(0.100)	(0.000231)	(0.258)	(0.000865)
Targeted Only Ad	0.238***	-0.00110	1.410	0.00628
	(0.0449)	(0.00110)	(1.577)	(0.00558)
Personalized Ad	0.0194	-0.00220**	0.189	-0.00351
	(0.0777)	(0.00110)	(1.456)	(0.00481)
News Articles			0.00376	0.00000285
			(0.00475)	(0.0000128)
Targeted Only × News Articles			-0.00545	-0.0000160
			(0.00481)	(0.0000135)
Personalized × Facebook News Articles			-0.00361	0.00000785
			(0.00487)	(0.0000145)
Google Searches			-0.00451	0.0000154
			(0.0270)	(0.0000746)
Targeted Only × Google Searches			-0.00311	-0.0000532
			(0.0287)	(0.0000898)
Personalized × Google Searches			0.00507	-0.00000604
			(0.0278)	(0.0000844)
Date and Campaign fixed effects	Yes	Yes	Yes	Yes
Observations	1255524	1255524	1255524	1255524

Columns (1) and (3) display Logit Estimates. Columns (2) and (4) display marginal effects. Dependent variable is whether or not the individual clicked. Sample is all occasions a unique Facebook user was exposed to the advertising campaign.
Robust standard errors clustered at ad-level. * $p < 0.10$, ** $p < 0.05$,*** $p < 0.01$.
Postpolicy is not separately identified from the date fixed effects.

Further Robustness Checks

The identification assumption underlying these results is that there was no change in Facebook user and advertiser behavior that drove our results that was not associated with the change in privacy controls. To check robustness to this assumption, we obtained further external data on Facebook user and advertiser behavior.

One potential concern is that our results reflect a change in the numbers of users of Facebook. For example, the negative publicity could have driven more experienced users away, leaving only users who were likely to react to personalized advertising using Facebook. However, the comScore data, based on their panel of two million internet users, in Table A-3 in the appendix suggests that this was not the case, and that instead there was actually an increase in the number of users. Importantly from this study's perspective, there were only small changes in the composition of the user base in June relative to May, and the shifts did not seem to be more dramatic than the shifts seen from April to May. To make sure this was the case, and with the caveat that we only have a limited number of observations, we used the Grubbs (1969) test for outliers for the full year of data. The results of this test did not indicate that observed changes between May and June deviated from the expected normal distribution.

Though observed demographics were reasonably similar, there is always the possibility that the composition of Facebook users changed in an unobserved way and that this influenced the kind of ads that were shown in the period after the introduction of policy controls. For example, there could have been more fans of a celebrity who was famous for directly reaching out to the public and whose fans consequently were more likely to have a taste for personalization using Facebook after the introduction of improved privacy controls. To check for this, we verified empirically that the mix of ads displayed did not change over time. If the composition of ads did change, then this could be a response to the fact that more consumers of that type were going online or, alternatively, that the same number of consumers were spending longer online. Table 5 reports the results. Column (1) suggests that it was not the case that more ads associated with undergraduate institutions were shown after the introduction of improved privacy controls. Column (2) suggests it was not the case that ads which had a larger potential reach, for example ads associated with famous celebrities, were shown more frequently after the introduction of improved privacy controls. Column

(3) combines these two measures and again finds no significant change after the policy in terms of what ads were shown. We also ran a specification which interacted the 39 targeting variables with the *Postpolicy* indicator. None of these interactions were significant.

Table 5: Test of whether there was a change in the types of ads were being shown before and after the introduction of improved privacy controls

	(1)	(2)	(3)	(4)
PostPolicy	-77.71	-108.6	-17.32	26.90
	(117.9)	(236.1)	(70.60)	(163.6)
PostPolicy × School Indicator		-60.14		71.91
		(241.3)		(170.9)
PostPolicy × Ad Reach			633.0	710.1
			(923.0)	(908.0)
Targeting Variable Fixed Effects	Yes	Yes	Yes	Yes
Observations	2730	2730	2730	2730
R^2	0.050	0.050	0.051	0.051

OLS Estimates. Dependent variable is number of times each ad is shown.
Robust standard errors clustered at ad-level. * $p < 0.10$, ** $p < 0.05$,*** $p < 0.01$
Ad-Reach$_k$ and *School* are collinear with the targeting variable fixed effects and also dropped from the specification.

It is also possible that, rather than a change in user composition, what we are measuring is actually driven by a dramatic change in how people use Facebook. For example, an alternative explanation of our results could be that after the introduction of improved privacy controls, people were more likely to spend time on Facebook and consequently more likely to eventually click on a personalized ad, perhaps because they mistook it for non-commercial content. Table A-4 in the appendix presents data from Compete, Inc., about how users' browsing behavior on Facebook changed over 2010. There does not appear to be any large or dramatic change in users' browsing behavior in the period we study, compared to the natural fluctuations that are apparent for the rest of the year. Though checking for outliers in such a short time series faces obvious problems, the Grubbs (1969) test for outliers indicated that the post-policy period did not deviate from the expected normal distribution.

Another concern is that the results could reflect a change in the composition of advertisers. For example, perhaps other advertisers pulled out of Facebook as a result of the negative publicity concerning the privacy interface, meaning that perhaps there were fewer advertisers competing to personalize advertising, which made the personalized ads relatively more attractive. Though we cannot check for evidence of this directly, we are able to provide some suggestive evidence against this counter-explanation by looking at the pricing data for the ads. If there had been a drop-off in advertisers, we would expect also to see a decrease in the price paid in the ad auction, as the price should theoretically be a function of the number of bidders (McAfee and McMillan, 1987). However, the small drop in cost per click of 1.5 cents (3%) after the introduction of improved privacy controls was not statistically significant (p-value=0.59).

Another concern is that though there could be an increase in the proportion of clicks for an ad, this increase might not have been helpful for the marketing aims of the nonprofit. For example, an alternative explanation for our results is that after the introduction of improved privacy controls, consumers became more likely to click on an ad that appeared too intrusive, in order to find out what data the advertiser had or because they were curious as to how they obtained their data, rather than it being the case they were more likely to respond positively to the advertising appeal.

To investigate this possibility we obtained confidential data from the nonprofit, based on weekly update emails from Facebook that recorded how many people had become their 'fan' on Facebook, that is, subscribed to their newsfeed. Prior to the introduction of improved privacy controls there was a 0.97 correlation on average at the weekly level. After the introduction of improved privacy controls there was a 0.96 correlation. There was no statistically significant difference between these two correlations, suggesting that it was not the case that after the introduction of improved privacy controls people were more likely to click on the ad even if they had no interest in the work of the nonprofit.

Economic Significance

These robustness checks reassure that the change in clicks we measure is associated with the introduction of improved privacy controls, but it is not clear whether the change was economically significant. This is always an important question, but is particularly important here because click-through rates are so low. Comscore data suggests that users of Facebook.com had 297 billion display ad impressions in the third quarter of 2010. The estimates presented in Column (6) of Table 3, suggests that the pre-policy click-through rates of 0.07 increased by 0.16 percentage points for personalized advertising after the introduction of improved privacy controls. We do not know, of course, how many ads on Facebook are indeed personalized and that our results consequently imply were affected by the introduction of improved privacy controls. However, even if a small percentage such as 10 percent were personalized, then our estimates would suggest that there were 47.5 million more clicks on such ads after the introduction of improved privacy controls. If all advertisers were also paying around 50 cents for these clicks, then this would translate into an increase in quarterly revenue of $23.8 million. This is a striking boost for any advertising platform. Of course, these estimates should be interpreted cautiously, as they are based on many unverifiable assumptions about the distribution of Facebook advertiser behavior. They do, however, highlight that the magnitude of these effects are likely to be large, even though click-through rates remain low, for advertising platforms as enormous as Facebook.

Mechanism: Rarity of User Information

We now turn to explore the behavioral mechanism that underlies our results. Edwards et al. (2002); White et al. (2008) have shown that personalized ads can lead to a process of 'reactance' (Brehm, 1966), where consumers deliberately resist ads that they perceive as intrusive. A potential explanation for the efficacy of privacy controls that we find in our natural experiment is that they reduced consumers' level of reactance to personalized advertising.

To provide evidence for this proposed mechanism, we exploit the fact that earlier studies have shown that reactance to personalized advertising is larger for ads that use more unique information about the consumer (White et al., 2008). For example, if an ad was personalized around the fact that a Facebook user liked 'cooking,' then Facebook has 3,167,540 users who say on their profiles that they like cooking. The use of such information might be felt to be less intrusive and consequently less likely to invoke reactance. However, if an ad was personalized around the fact that a user liked the Korean delicacy kimchi, then there are only 6,180 Facebook users who say that they do like kimchi; knowing that such a preference is relatively rare might make the user more concerned they were being tracked by the advertiser in a privacy-violating manner, increasing intrusiveness and consequently provoking reactance.

To explore this in our empirical setting we use additional data on how many users were in the target group for that particular campaign. We modify our equation (1) for the click-through rate $ClickRate_{jt}$ for ad j on day t in the following manner:

$$\begin{aligned} ClickRate_{jt} = \beta_1 Personalized_j \times PostPolicy_t \times AdReach_k + \\ \beta_2 Personalized_j \times PostPolicy_t + \beta_3 PostPolicy_t \times AdReach_k + \\ \alpha_1 Personalized_j + \alpha_2 Personalized_j \times AdReach_k \\ \theta MediaControls_{jt} + \gamma_k + \delta_t + \epsilon_j \end{aligned} \tag{3}$$

$AdReach_k$ and $Postpolicy_t$ are collinear with the date and campaign fixed effects so are dropped from the equation.

Table 6 uses equation (3) to investigate how our effects were moderated by how large or small the reach of the ad was - how many people, potentially, the ad could be shown to. Column (1) of Table 6 reports how the efficacy of personalized ads relative to ads that were targeted to users' interests before and after the introduction of improved privacy controls was affected by ad-reach for our initial specification. The negative coefficient on *Post-Policy* × *Personalized* × *Ad-Reach* suggests that the positive effect is smaller for ads that had a larger ad-reach than those that had a smaller ad-reach. In other words, personalization was relatively more successful after the introduction of privacy controls for celebrities who had smaller fan bases or schools with smaller numbers of graduates on Facebook, as can be seen from the larger point estimate for *Post-Policy* $\times$ *Personalized* relative to Table 3, Column (1).

Column (2) of Table 6 repeats this exercise for ads that used the shorter ten-day window (Day 13-Day 22). Again, the results appear robust, providing evidence against an interpretation that an unobserved time-trend unrelated to the change in policy drove the results. Column (3) of Table 6 shows that the change observed is similar in our main specification that includes media controls. Ad-Reach is denominated in millions of users. Therefore, roughly extrapolating from the linear functional form, our estimates suggest that for ads for the campaigns in our sample that have target audiences of greater than 243,000, the effect of the policy was canceled out. However, for the median campaign, which had 7,560 people in the target market, the introduction of privacy controls actually raised the click-through percentage by 0.03, relative to a mean of 0.02.

Further Evidence from an Experimental Setting

These results suggest that the shift towards giving users control over their personal information had the largest effect for personalized advertising that attempted to use more unusual pieces of information. This provides suggestive evidence that the change in privacy policy reduced reactance to personalized advertising and that that is why there were more clicks after the policy. However, nothing in the empirical data allows us to actually measure 'reactance' directly. Therefore, we turned to an artificial lab-like setting to gather direct evidence on the effect of privacy controls and the uniqueness of personal information on reactance.

We recruited 240 survey-takers from Mechanical Turk to take part in an online survey. The study has a 2×2 design, which varied celebrity fame (Non-Unique Data, Unique Data)

Table 6: Stratification

	(1) Main	(2) 10-Day Window	(3) Controls
Post-Policy × Personalized × Ad-Reach	-0.0852**	-0.199**	-0.0852***
	(0.0421)	(0.0966)	(0.0313)
Personalized	-0.0153**	-0.0228	0.377
	(0.00670)	(0.0188)	(0.364)
Post-Policy × Personalized	0.0317**	0.0662***	0.0243**
	(0.0147)	(0.0251)	(0.0114)
Personalized × Ad-Reach	0.0354	0.125	0.0354*
	(0.0214)	(0.0850)	(0.0211)
Post-Policy × Ad-Reach	0.0150	-0.00677	0.0150
	(0.0350)	(0.0497)	(0.0204)
Personalized Ad × Number Facebook News Articles			0.0538
			(0.0356)
News Articles			0.00149
			(0.0269)
Google Facebook Privacy Searches			0.0472
			(0.0632)
Personalized Ad × Google Facebook Privacy Searches			-0.152
			(0.108)
Date Fixed Effects	Yes	Yes	Yes
Targeting Variable Fixed Effects	Yes	Yes	Yes
Observations	2730	780	2730
R^2	0.062	0.129	0.063

OLS Estimates. Dependent variable is percentage daily click through rate.
Robust standard errors clustered at ad-level. * $p < 0.10$, ** $p < 0.05$,*** $p < 0.01$
$PostPolicy_t$ is collinear with the date fixed effects and dropped from the specification. $Ad\text{-}Reach_k$ is collinear with the targeting variable fixed effects and also dropped from the specification.

and level of privacy control (No Privacy Controls, Privacy Controls). We asked people to imagine that they regularly used a social networking website. They were also told that when they were browsing the website they saw an ad for a product they were interested in and the text of the ad mentions the name of one of the people that they were a fan of. The survey respondents were then randomly allocated to two conditions. In the 'Non-Unique Data' condition, they were told that 'This person is internationally famous and their name is known by everyone.' In the 'Unique Data' condition they were told that 'This person is very obscure and hardly anyone would recognize their name.' We also varied users' perception of privacy control. In the 'Privacy Controls' condition, they were told that the 'Social networking website had recently introduced privacy controls that gave them control over their personal information.' In the 'No Privacy Controls' condition, they were told that the social networking website did not have adequate controls governing their personal information. This manipulation appeared to be effective. In a preliminary manipulation check, respondents reported that on a 7-point scale they were more likely to feel in control of their privacy in the privacy control condition (5.3 vs 3.95, F=7.01, p-value$<$0.01).

We then asked respondents seven questions designed to gauge their level of 'reactance' to the ad and the situation. These questions were based on the scales developed by Edwards et al. (2002); White et al. (2008) and Lamberton (2011), which in turn were based on the scale developed by Hong and Faedda (1996). This scale covers the extent to which the ad was considered to be interfering, intrusive, forced, unwelcome, discomforting, curtailing of freedom and manipulative, measured on a 7-point scale (α=0.89). Column (1) of Table 7 reports the results. In line with the work of White et al. (2008), the mention of a rare celebrity increases reactance significantly. However, the introduction of Privacy Controls for users in the Rare Celebrity condition reduces reactance significantly. There is no significant main effect of 'Privacy Controls' for respondents in the 'Common' celebrity condition where there was less reactance, which accords with the results reported in Table 6.

We also asked respondents five questions about how likely they were to positively respond to the ad. This scale covered the extent to which they were likely they were to take notice of the ad, remember the ad, purchase the product, and think about the product on a seven-point scale (α=0.87). Column (2) reports the results. As expected, the results reverse themselves from Column (1). Respondents report that they are less likely to react favorably to an ad using unique data in the absence of privacy controls. However, in the presence of privacy controls they are actually more likely to react favorably an ad with unique data than for a non-unique data. Column (3) analyzes click-through intent. The results are very similar to Column (3), though slightly less significant. However, the main finding of the natural experiment is replicated. That is, after the introduction of privacy controls, respondents are more likely to click on an ad that uses unusual personal information.

Table 7: Lab Experiment Results

	(1) Reactance	(2) Favorable Ad Response	(3) Click Intent
Unique Data × Privacy Controls	-0.730**	1.212***	0.994**
	(0.303)	(0.369)	(0.414)
Unique Data	1.022***	-0.755***	-0.563*
	(0.214)	(0.281)	(0.311)
Privacy Controls	-0.0435	0.0493	0.193
	(0.212)	(0.256)	(0.301)
Constant	3.377***	4.997***	3.714***
	(0.151)	(0.191)	(0.223)

OLS Estimates. 240 Respondents. Dependent variable is a seven-point scale.
* $p < 0.10$, ** $p < 0.05$,*** $p < 0.01$

There are obvious limitations about the generalizability of the results of any experiment in an artificial setting, but there are also some obvious advantages to having replicated the effect in a controlled experimental environment. First, we are able to explicitly measure reactance and how it is ameliorated by privacy controls and in turn how this interacts with how 'personal' the personal information used in a personalized ad is. Second, we ask questions about the purchase of a generic product, suggesting that the earlier results are not limited to the nonprofit sector. Last, and crucially, because we use a randomized between-subjects design, we are able to rule out alternative explanations for the results in Table 6 that involve endogeneity or selection into the liking of either unusual or very popular celebrities.

Implications

This paper is the first to explore the consequences for an advertising-supported website of giving users more control over how their personal information is shared. The paper uses data from a randomized experiment conducted by a nonprofit that was designed to explore the relative merits of targeting ads, and ads that used user information to personalize content of the ad. During the field experiment, the social networking site on which the experiment was being conducted unexpectedly announced that it would change the interface through which users controlled their privacy settings. These changes, which were publicly applauded by consumer advocates, gave users greater control over what personally identifiable information was shared and whether third parties could track their movements. They did not however, affect the anonymous use of information by advertisers to target their information.

Empirical analysis suggests that after this change in policy, Facebook users were roughly twice as likely to react positively to personalized ad content and click on personalized ads. There was generally no economically significant change in their reactions to untargeted or merely targeted ads. This suggests that publicly giving users control over their private information can benefit advertising-supported media and advertisers on such sites. This has important consequences for the current policy debate which has generally viewed the introduction of privacy controls as being harmful to advertising outcomes (Goldfarb and Tucker, 2011b).

There are obvious limitations to this research that are worth mentioning. First, the randomized experiment was conducted by a nonprofit with an appealing cause. Consumers may be more or less ready to ascribe less pernicious motives to a nonprofit than to a for-profit company when they observe their advertising. Second, this randomized experiment was conducted at a time when privacy concerns were particularly sensitive and salient in consumers' eyes. Though we do use controls for the extent of the publicity surrounding privacy and Facebook, it is not clear how the results would change when the introduction of controls is not so heavily publicized by the media. Third, we do not know how long the positive effects we measured after the introduction of privacy controls for personalized advertising persisted. Fourth, the type of privacy control introduced by Facebook that we study was just one of a myriad of potential ways that social networks or other advertising-supported websites could have used to give control over their privacy settings to their users.

It would be interesting for future research to see whether an explicit 'opt-in' approach to sharing information or changes in privacy policies that explicitly addressed advertising could produce equally striking results. Notwithstanding these limitations, this paper does provide initial evidence of how addressing privacy concerns of consumers is important for online advertising venues.

References

Acquisti, A. and R. Gross (2006). Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *6th Workshop on Privacy Enhancing Technologies*, pp. 36–58.

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

Anand, B. and R. Shachar (2009, September). Targeted advertising as a signal. *Quantitative Marketing and Economics 7*(3), 237–266.

Aral, S. and D. Walker (2010). Creating social contagion through viral product design: Theory and evidence from a randomized field experiment. *Mimeo, NYU*.

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.

Bilton, N. (2010, May 12). Price of Facebook Privacy? Start Clicking. *New York Times*.

Brandimarte, L., A. Acquisti, and G. Loewenstein (2010). Misplaced confidences: Privacy and the control paradox. *Mimeo, Carnegie Mellon University*.

Brehm, J. W. (1966). *A theory of psychological reactance.* Academic Press, New York.

Brehm, J. W. (1989). Psychological reactance: Theory and applications. *Advances in Consumer Research 16*, 72–75. eds. Thomas K. Srull, Provo, UT.

Caverlee, J. and S. Webb (2008). A Large-Scale Study of MySpace: Observations and Implications for Online Social Networks. In *Proceedings from the 2nd International Conference on Weblogs and Social Media (AAAI).*

Clee, M. A. and R. A. Wicklund (1980). Consumer behavior and psychological reactance. *The Journal of Consumer Research 6*(4), pp. 389–405.

Cormier, J. (2010, November 10). Facebook sees 23 percent of online display ad impressions. *Dallas Examiner*.

Dreze, X. and F.-X. Hussherr (2003). Internet advertising: Is anybody watching? *Journal of Interactive Marketing 17*(4), 8–23.

Edwards, S. M., H. Li, and J. Hyun Lee (2002). Forced exposure and psychological reactance: Antecedents and consequences of the perceived intrusiveness of pop-up ads. *Journal of Advertising 31*, 83–95.

Golder, S. A., D. M. Wilkinson, and B. A. Huberman (2007). Rhythms of social interaction: Messaging within a massive online network. In *Proceedings of the Third International Conference on Communities and Technologies.*

Goldfarb, A. and C. Tucker (2011a). Online display advertising: Targeting and obtrusiveness. Forthcoming, Marketing Science.

Goldfarb, A. and C. E. Tucker (2011b). Privacy regulation and online advertising. *Management Science 57*(1), 57–71.

Grubbs, F. E. (1969). Procedures for detecting outlying observations in samples. *Technometrics 11*, 1–21.

Holahan, C. (2007, November). So Many Ads, So Few Clicks: Can more targeted pitches on Facebook and other sites reverse the shrinking response to online ads? *Business Week*.

Hong, S.-M. and S. Faedda (1996). Refinement of the Hong psychological reactance scale. *Educational and Psychological Measurement 56*(1), 173–182.

Kerr, R. (2011, January 18). Facebook ad revenue to top $4 billion in 2011. *Vator News*.

Kihlstrom, R. E. and M. H. Riordan (1984, June). Advertising as a signal. *Journal of Political Economy 92*(3), 427–50.

Lamberton, C. (2011). A spoonfull of choice: Using allocation to increase consumer satisfaction with compulsory payments. *Mimeo, University of Pittsburgh*.

Lipsman, A. (2010, January 2010). 2009: Another Strong Year for Facebook. *Comscore Voices*.

Lohr, S. (2010, April 30). Privacy concerns limit online ads, study says. *New York Times*.

Manchanda, P., Y. Xie, and N. Youn (2008). The role of targeted communication and contagion in new product adoption. *Marketing Science 27*(6), 961–976.

McAfee, R. P. and J. McMillan (1987). Auctions and bidding. *Journal of Economic Literature 25*(2), 699–738.

Miller, A. (1971). *The assault on privacy: Computers, data banks, and dossiers*. University of Michigan Press.

Nielsen (2010, January 14, 2010). Top U.S. web brands and site usage: December 2009. *Nielsen Wire*.

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.

Rothbaum, F., J. R. Weisz, and S. S. Snyder (1982). Changing the world and changing the self: A two-process model of perceived control. *Journal of Personality and Social Psychology 42*(1), 5–37.

Small, D. and N. Verrochi (2009). The face of need: Facial emotion expression on charity advertisements. *Journal of Marketing Research 46*, 777–787.

Spiekermann, S., J. Grossklags, and B. Berendt (2001). E-privacy in 2nd generation E-commerce: privacy preferences versus actual behavior. In *EC '01: Proceedings of the 3rd ACM conference on Electronic Commerce*, New York, NY, USA, pp. 38–47. ACM.

Stone, B. (2010, March 3). Ads Posted on Facebook Strike Some as Off-Key. *New York Times*.

Taylor, S. E. (1979). Hospital patient behavior: Reactance, helplessness, or control? *Journal of Social Issues 35*(1), 156–184.

Thompson, S. C., A. Sobolew-Shubin, M. E. Galbraith, L. Schwankovsky, and D. Cruzen (1993). Maintaining perceptions of control: finding perceived control in low-control circumstances. *Journal of Personality and Social Psychology 64*(2), 293–304.

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.

Turow, J., J. King, C. J. Hoofnagle, A. Bleakley, and M. Hennessy (2009). Americans Reject Tailored Advertising and Three Activities that Enable It. *Mimeo, Berkeley*.

White, T., D. Zahay, H. Thorbjornsen, and S. Shavitt (2008, March). Getting too personal: Reactance to highly personalized email solicitations. *Marketing Letters 19*(1), 39–50.

Xu, H. (2007). The effects of self-construal and perceived control on privacy concerns. In *28th Annual International Conference on Information Systems (ICIS 2007),Montreal, Canada.*

A Data Appendix

Table A-1: Top 10 U.S. Online Display Ad Publishers Q3 2010

Website	Display Ad Impressions (MM)	Share of Display Ad Impressions
Facebook.com	297,046	23.1
Yahoo! Sites	140,949	11.0
Microsoft Sites	64,009	5.0
Fox Interactive Media	48,252	3.8
Google Sites	35,043	2.7
AOL, Inc.	32,330	2.5
Turner Network	21,268	1.7
Glam Media	13,274	1.0
eBay	8,421	0.7
ESPN	8,261	0.6

Source: comScore Ad Metrix. Display ads include static and rich media ads.

Table A-2: Timeline for Facebook Growth, Privacy and Advertising

Date	Event
February 2004	Facebook launched from Harvard dorm room.
November 2007	Facebook launches 'Facebook ads'. Advertising pilot involving 'beacons' (small 1x1 pixel web bugs) allows Facebook to track users' movements over other websites for purposes of targeting.
December 2007	Facebook makes Beacon an opt-out service after negative publicity.
September 2009	Beacon ad targeting program shut down amid class-action suit.
November 2009	Facebook changes its default settings to publicly reveal more of its users' information that had previously only been available to Facebook users. This information could now be tracked by third-party search engines.
December 9 2009	Privacy settings are entirely removed from certain categories of users' information. These categories include the user's name, profile photo, list of friends and pages they were a fan of, gender, geographic region, and networks the user was connected to. They are instead labeled as publicly available to everyone, and can only be partially controlled by limiting search privacy settings. Founder Mark Zuckerberg's photos are apparently inadvertently made public by the change in settings.
December 17 2009	Coalition of privacy groups led by the Electronic Frontier Foundation files a complaint with Federal Trade Commission over changes to privacy settings
April 2010	Facebook users' General Information becomes publicly exposed whenever they connect to certain applications or websites such as the online review site Yelp. General Information includes users' name and their friends' names, profile pictures, gender, user IDs, connections, and any content shared using the Everyone privacy setting.
May 12 2010	New York Times publishes article entitled 'Facebook Privacy: A Bewildering Tangle of Options' (Bilton, 2010). This ignites a firestorm of negative press about Facebook and privacy.
Monday May 24 2010	Facebook founder Mark Zuckerberg announces in an editorial in the Washington Post that Facebook will institute new privacy settings
Wednesday May 26 2010	Facebook unveils new privacy settings in press event
Thursday May 27 2010	Facebook starts rollout of privacy settings. New York Times publishes 'A Guide to Facebook's New Privacy Settings'.
Saturday May 29 2010	First reports of new privacy setting controls being seen by users

Additional Sources: Facebook's official public timeline; 'Facebook's Eroding Privacy Policy: A Timeline': Electronic Frontier Foundation April 2010.

Table A-3: There were only small changes in Facebook user composition

Proportion of Group	April 2010	May 2010	June 2010
Age <17	10.4	10.6	11.4
Age 18-24	19.2	19.4	18.6
Age 25-34	20.8	20.7	20.8
Age 35-44	20.4	19.9	19.9
Age 45-54	16.7	16.5	16.5
Age 55-64	8	8.1	8.1
Age 65+	4.6	4.8	4.7
Income <$15k	10.1	10.3	9.7
Income $15-24k	6.2	6.1	5.9
Income $25-39k	12.5	12.7	13.5
Income $40-59k	22.1	22	24.2
Income $60-74k	10.9	11.3	9.6
Income $75-99k	16.8	16.3	15.3
Income $100k+	21.5	21.2	21.8
Male	47.2	47.1	48.2
Female	52.8	52.9	51.8
Total Unique Visitors	121 Million	130 Million	141 Million

Source: Comscore Marketer Database

Table A-4: There was little change in how Facebook users used the website

Date	Average Stay	Visits / Person	Pages / Visit
Dec-09	21:29	22.27	29.46
Jan-10	23:06	22.15	33.52
Feb-10	22:14	21.08	35.33
Mar-10	21:30	23.4	29:00
Apr-10	21:54	23.27	25.45
May-10	22:39	24.9	27.27
Jun-10	21:50	24.37	24.78
Jul-10	22:28	24.61	28.64
Aug-10	22:28	26.86	30.33
Sep-10	22:25	26.12	27.49
Oct-10	24:30	26.52	24.64
Nov-10	24:56	26.55	23.86
Dec-10	25:48	26.46	24.24

Source: Compete, Inc

Figure A-1: Facebook: Screenshot of Ad Targeting Interface

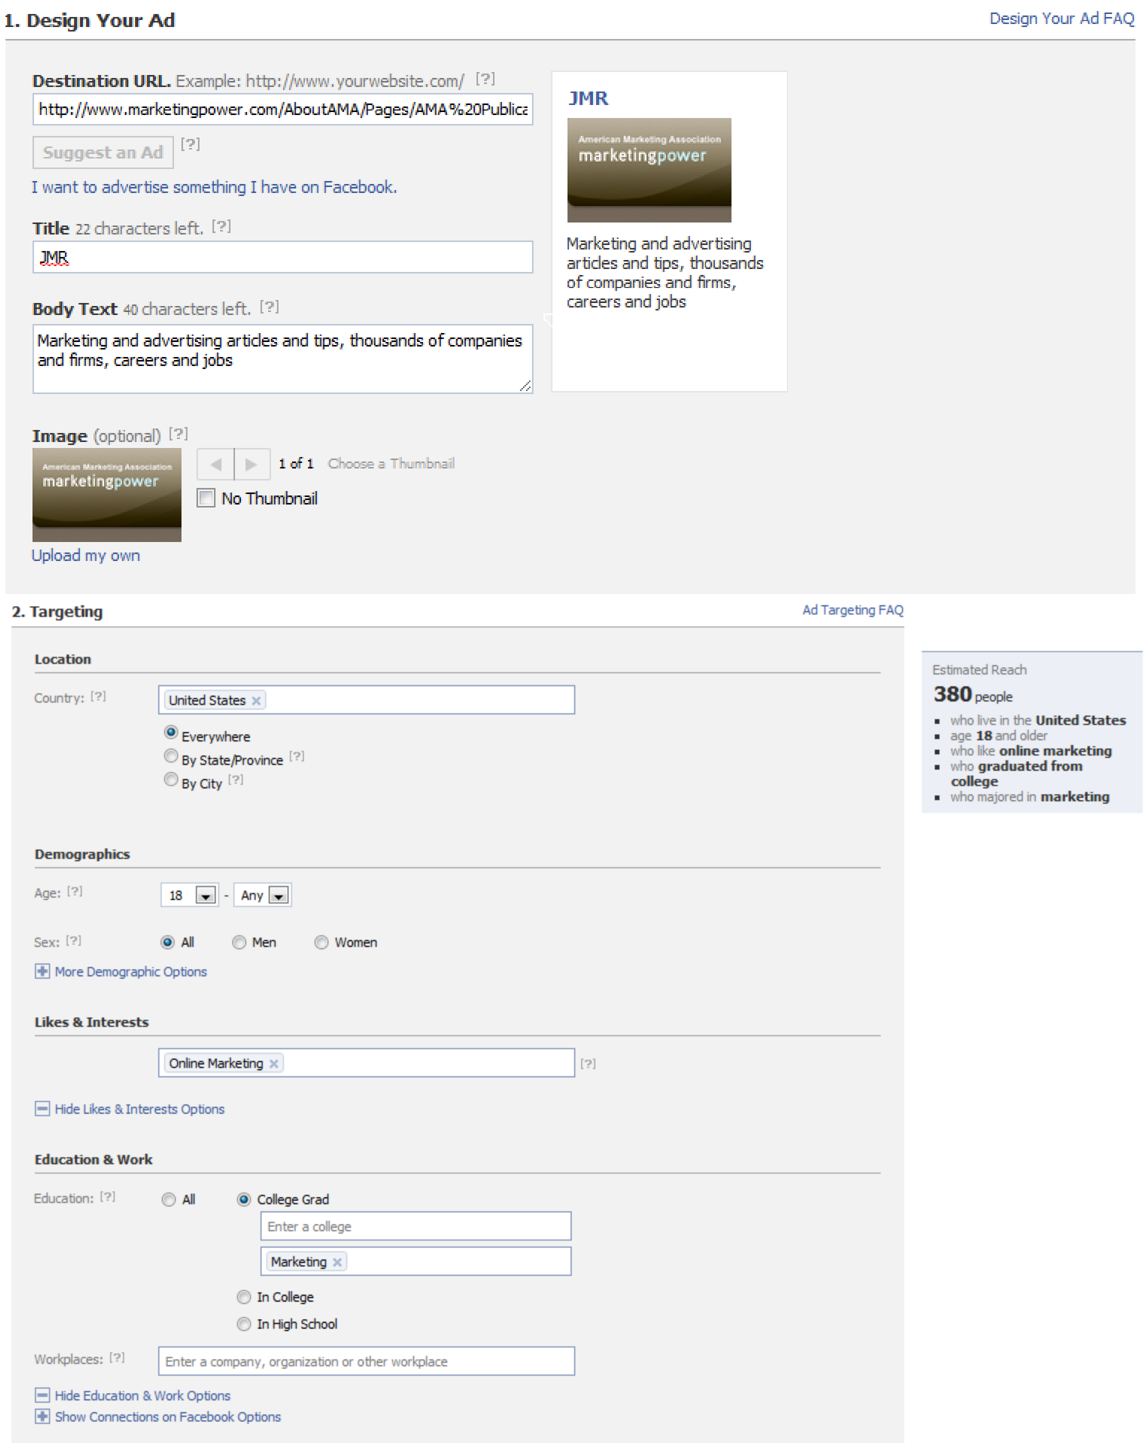

Source: Mock-up ad campaign for marketing journal created by authors This is a screenshot which shows the Facebook interface used to design and target an ad. To preserve the anonymity of the nonprofit, it shows a mock ad for a marketing journal. On the right-hand side, there is a sample ad which is similar in format to the ad used in the tests, and gives an accurate representation of the relative size of text and photo in the actual ad. The lower panel shows how an advertiser would theoretically target people who are interested in online marketing and who also had a college degree in marketing.

The change in privacy controls

Figure A-2: Facebook: Screenshots of Privacy Options before and after the introduction of privacy controls

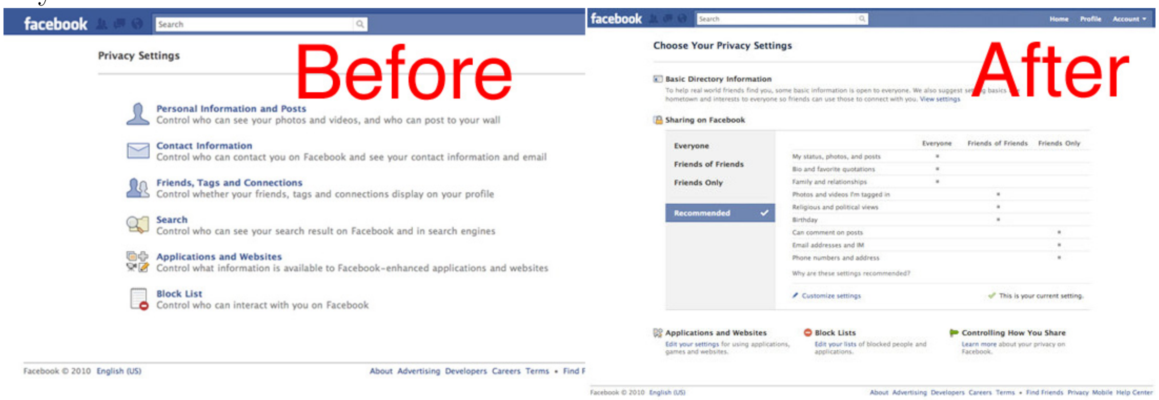

Source: Gawker Media

Exhibit A: Facebook's Notification to Advertisers: May 26, 2010

Facebook will roll out changes today that will make it easier for our users to understand and control their privacy settings. As this change will have an impact on our users, we wanted to let you, a valued advertising partner, know about it. Please note that this change will not affect your advertising campaigns and there is no action required on your part.

Facebook is a company that moves quickly, constantly innovating and launching new products to improve the user experience. The feedback we heard from users was that in our efforts to innovate, some of our privacy settings had become confusing.

We believe in listening to our users and taking their feedback into account whenever possible. We think the following changes address these concerns by providing users with more control over their privacy settings and making them more simple to use.

Starting today, Facebook will:

* Provide an easy-to-use "master" control that enables users to set who can see the content they share through Facebook. This enables users to choose, with just one click, the overall privacy level they're comfortable with for the content they share on Facebook. Of course, users can still use all of the granular controls we've always offered, if they wish.

* Significantly reduce the amount of information that must be visible to everyone on Facebook. Facebook will no longer require that users' friends and connections are visible to everyone. Only Name, Profile Picture, Networks and Gender must be publicly available. Users can opt to make all other connections private.

* Make it simple to control whether other applications and websites access any user information. While a majority of our users love Facebook apps and Facebook-enhanced websites, some may prefer not to share their information outside of Facebook. Users can now opt out with just one click.

I encourage you to take a moment to read our CEO Mark Zuckerberg's blog post and check out the new Facebook Privacy Page.

Thanks, The Facebook Ads Team