Social Networks, Personalized Advertising and Privacy Controls

Catherine Tucker*

May 31, 2013

*Catherine Tucker, MIT Sloan School of Management, MIT, Cambridge, MA; 617 252-1499 and NBER; cetucker@MIT.EDU. I thank the Time Warner Cable Research Program on Digital Communications and the Net Institute for financial support. I also thank Alessandro Acquisti, Emilio Calvano, Avi Goldfarb, Cait Lambert, Alex Marthews, Markus Mobius, Martin Peitz, Ken Wilbur and seminar participants at Dartmouth, Harvard, the NBER, the National University of Singapore, New York University, the University of Florida, the University of Mannheim and the University of Munich. 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 on a social networking website. The paper uses data from a randomized field experiment that examined the effectiveness of personalizing ad text with user-posted personal information relative to generic text. The website gave users more control over their personally identifiable information in the middle of the field test. However, the website did not change how advertisers used data to target and personalize ads. Before the policy change, personalized ads did not perform particularly well. However, after this enhancement of perceived control over privacy, users were nearly twice as likely to click on personalized ads. Ads that targeted but did not use personalized ads remained unchanged in effectiveness. The increase in effectiveness was larger for ads that used more unique private information to personalize their message and for target groups who were more likely to use opt-out privacy settings.

Keywords: Privacy, Online Advertising, Social Networks

1 Introduction

Many Internet firms have collected a huge amount of personal data from their users and use this data to allow their advertisers to target and personalize ads. Consumers might see personalized ad content on such sites as more appealing and more connected to their interests (Anand and Shachar, 2009), but they also may see it as 'not only creepy, but off-putting' if they feel that the firm has violated their privacy (Stone, 2010). These privacy concerns may lead to 'reactance' which leads consumers 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).

Internet firms are unsure about whether they should directly address such concerns by strengthening privacy controls. Theoretically, this could minimize the potential for customer reactance and improve the performance of online advertising, because behavioral research shows that consumer perceptions of control reduce reactance (Taylor, 1979). This holds even if the controls are only tangentially related 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 the risk that addressing consumer privacy concerns but continuing to use consumer data to personalize ads may make users less likely to respond to such ads. This ambiguity makes it an empirical question how strengthening privacy controls affects advertising performance.

We explore this question using data from a randomized field experiment conducted by a US-based non-profit organization (NPO) to optimize its advertising campaigns on Facebook, a social networking website. These campaigns were shown to 1.2 million Facebook users. The NPO's aim was to raise awareness of its work improving education for women in East Africa. The NPO randomized whether it explicitly personalized the ad copy to match data from the user's profile. In one condition, consumers saw ads that, for example, explicitly mentioned a celebrity that was featured in their Facebook profile. In other conditions they just saw generic text.

In the middle of the field experiment, in response to mounting media criticism Facebook changed their privacy policies. The policy change introduced an easy-to-use privacy control interface, reduced the amount of information that was automatically required to be displayed, and also gave users new controls over how 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. So, before and after the policy change, advertisers could still choose to target ads towards Oprah Winfrey fans, and personalize the advertising message to mention 'Oprah Winfrey,' in exactly the same manner as they did before. What changed was simply how easy it was for users to control access to their information by regular Facebook users.

The NPO had not anticipated there would be such a change when it launched its field test of its ads. However, the fact that this occurred mid-way through the field experiment is valuable for measuring the effect of a firm responding to privacy concerns by improving privacy controls on advertising effectiveness, while circumventing the usual endogeneity issues.

We analyze five weeks of campaign-level click-through data spanning the introduction of the new privacy controls. This suggests that personalized advertising, surprisingly, was relatively ineffective before Facebook introduced new privacy controls. However, it was nearly twice as effective at attracting users to the NPO'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.

This interpretation rests on the assumption that there were no underlying changes in user behavior or the environment that coincided with the introduction of the new privacy controls but were not directly attributable to the introduction of these controls. One concern is that there was a lot of publicity around the policy change, so we include many different controls for media coverage. We also 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. Last, we show that there was no change in how likely people were to sign up for the NPO's news feed after clicking on the ad, 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 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 also examine how the effect depends on the extent to which consumers use privacy settings on Facebook. This is empirically ambiguous. On the one hand, the kind of consumers who care about privacy and use privacy settings may be more upset if they set high levels of privacy restrictions and still see highly personalized advertising. On the other hand, such consumers who are more aware of privacy matters on Facebook may have experienced the most reactance prior to Facebook addressing privacy complaints with the improved set of privacy policies. They may have also been the most reassured that they could now explicitly prevent Facebook from sharing their click data with third parties. We find that the effect is larger for groups of consumers who used privacy controls to restrict the ability of another ad product to use their data. This also provides evidence that the effect we measure is associated with tastes for privacy controls rather than other external factors.

We replicate our results from the natural experiment with evidence from an online survey that tested consumer reactions to different online ads. These ads displayed either unique or not-at-all-unique private information that the same consumer had supplied earlier, in contexts where respondents either felt they had control over their personal information or that they did 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.

1.1 Contribution

These findings contribute at four levels.

First, to our knowledge, this is the first paper that uses field data to study the consequences for advertising outcomes of a firm responding to user privacy concerns by introducing improved privacy controls. 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). The finding that there are positive effects for an advertising platform, in this instance, from addressing users' privacy concerns is therefore useful.

Second, our findings have implications for privacy regulation. Currently, proposed regulations such as 'Do Not Track' governing online behavioral advertising in the US are focused around the mechanics of how websites 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 was shared and used, there was no negative effect on advertising performance. This is an important finding for policymakers deciding on whether to emphasize user-based controls in privacy regulation both in the US and elsewhere.

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 can mediate privacy concerns (Fusilier and Hoyer, 1980; Culnan and Armstrong, 1999; Malhotra et al., 2004). This remains the case even if the control is merely perceptual or over tangential information, and access to the focal data remains unchanged (Spiekermann et al., 2001; Xu, 2007; Brandimarte et al., 2012; Xu et al., 2012).

The paper also contributes to the online advertising literature that has studied targeting in data-rich social media sites. This is important because social networking websites now account for one-third of all online display advertising (Marshall, 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 paper 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. Previous studies in marketing about social networking have looked at how offline social networks can be used to target (Manchanda et al., 2008), how such social networking sites can use advertising to obtain members (Trusov et al., 2009), how social networks can be used to target ads (Tucker, 2012), and also how makers of applications designed to be used on social networking sites can best advertise their products (Aral and Walker, 2011). Outside of social networks, Goldfarb and Tucker (2011a) have shown that privacy concerns can influence ad effectiveness.

2 Data

2.1 The Nonprofit Organization (NPO)

The NPO 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 NPO'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 NPO set up a Facebook 'page' that explained its mission and 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 NPO started advertising using Facebook's own advertising platform. Initially, it 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 the NPO to determine whether it could engage further with its potential supporters by both targeting and personalizing ad content.

2.2 Randomized Campaign

The NPO decided to target both graduates from 20 liberal arts colleges with a reputation for supporting female education and Facebook users who had expressed affinity with 19 celebrities and writers who in the past had made statements supporting the education of girls in Africa or African female empowerment in general.[1] Using the Facebook advertising

[1]The NPO is eager to protect the privacy of its supporters, and consequently has asked the authors to not reveal the names of either the actual celebrities or the schools that were used in this advertising campaign. Examples could be Oprah Winfrey, who has set up a girls' school in South Africa, or Serena Williams, who

interface, we also verified that there was very little overlap in fans across these different groups.

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 message, they decided to experiment with two different ad formats. Table 1 summarizes the different conditions used. In the targeted and 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 that had been used to target the ad. 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, our estimates reflect 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 sad expression.

Table 1: Campaign appeals in different conditions

Information used to target ad	College	Interest
Targeted and 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.
Targeted and 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.

The NPO also continued to use as its baseline an untargeted campaign that 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

was a supporter of 'Build African Schools.'

lives through education." All campaigns were restricted to Facebook users who live in the US, and were 18 years and older who were not already fans of the NPO. The NPO 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.

2.3 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 its privacy settings as well as rolling back some previous changes that had made Facebook users' information more public. Studying this change was not the purpose of the randomized field experiment, but it presents a unique opportunity to study how a change in user privacy controls in social networking sites could change consumer responses to advertising, since the NPO tested the ads using the same randomization technique before and after the change in the privacy-control interface.

Facebook introduced this improved privacy interface after being heavily criticized because its privacy settings were very granular and difficult to access. For example, Bilton (2010) pointed out in the national press 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.[2] 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.

[2]As detailed by Table A-2, Facebook had previously acted to reduce the amount of control users had over their data and had attracted negative publicity for doing so.

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

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-1, 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 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 publicly 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.

One consequence of studying a real-life firm shift in privacy policy is that, unlike with a lab experiment, many things were changed at once. Therefore, ultimately we measure the effect of the combination of different privacy measures, as well as how they were reported and received in the press. We are capturing, besides an increased sense of control, other positive effects of Facebook's improved privacy policy, such as higher trust in the firm. Since these improvements in general perceptions may attend any firm's improvement of privacy policies in response to privacy critiques, this still makes our policy estimates of managerial interest. When we move to the lab, however, we can isolate better the effect of improved control.

2.4 Data

We obtained daily data from the NPO 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 NPO's Facebook page. The 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 not the time. It also includes data on the cost to the NPO per click and the imputed cost per thousand impressions. 'Ad-reach' measures the number of Facebook users who were eligible to be shown the ad for any targeted ad campaign. We use this ad-reach data in subsequent regressions to explore the behavioral mechanism. To protect the privacy of the NPO'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.

Table 2 reports the summary statistics. The average number of clicks relative to ad impressions is small, at two-tenths of one percentage point. The maximum click-through rate is 3 percentage points. This average is even smaller when looking at the daily level, since many campaigns received no clicks on a given day, inflating the appearance of low

Table 2: Summary Statistics

	Mean	Std Dev	Min	Max
Average Impressions	15694.0	47807.1	337	376528
Average Click-Through (Percentage Points)	0.023	0.14	0	3.13
Ad-Reach (000000)	0.095	0.21	0.00098	0.99
Average Cost Per Click	0.38	0.095	0.11	0.50
Cost per 1000 views	0.095	0.11	0	0.39
News Article containing 'Facebook' & 'Privacy'	61.1	39.7	10	210
Google Searches	75.7	7.46	62	91
Words Devoted to News	37887	39559	1160	152240

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

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 for the data we use as controls for user exposure to news about Facebook and privacy. The first is data from Factiva about the number of newspapers that had stories that contained the words 'privacy' and 'Facebook' and how many words devoted to the topic. The second is data from Google Trends that gives an index for the number of searches for Facebook and privacy.

3 Analysis

Figure 1 displays the average click-through rate for each campaign before and after the introduction of improved privacy controls. Before the policy change, the personalized ads were under-performing relative to their generic counterparts. This is surprising, given the expectation that displaying personalized ad text would increase their relevance and consequently their appeal. However, after the policy change the more expected pattern prevails where ads with personalized content were relatively more effective than generically worded but

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

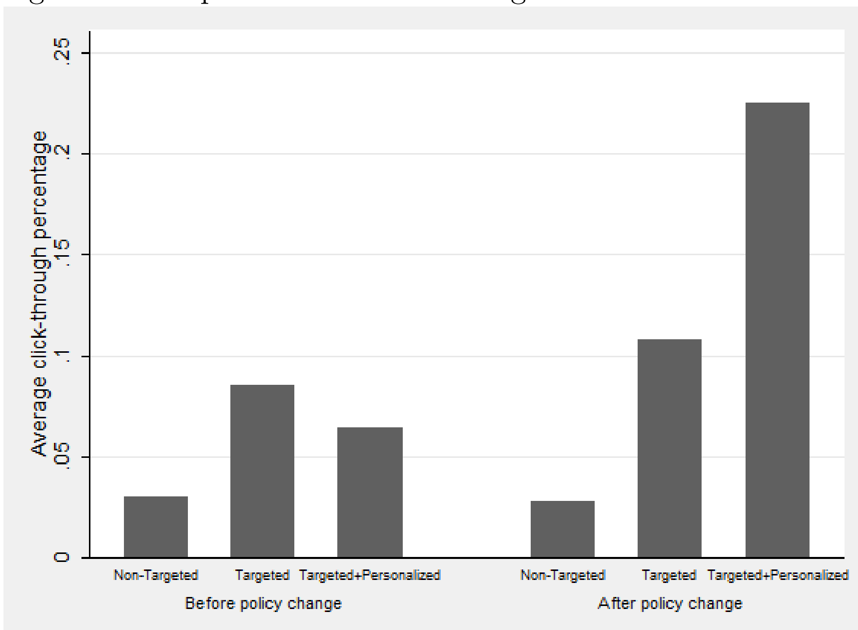

targeted ads or untargeted ads. This change was highly significant (p-value=0.0047). The effects of targeting ads without personalizing their content before and after the introduction of improved privacy controls were not significantly different (p-value=0.407). There appears to be little change in the 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 Facebook users with affinities to different schools and different celebrities.

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

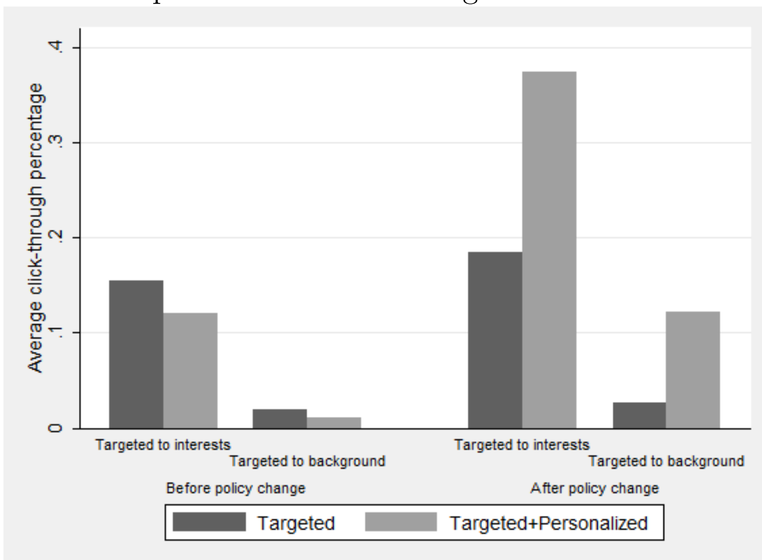

Figure 1 suggests that the personalization of display ads was more effective after Facebook facilitated users' taking control of their personal information. To check the robustness of this result, we also performed regression analysis. This allows us to assess the statistical significance of our results in various ways and to control for media coverage.

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

$$\begin{aligned} ClickRate_{jt} = {} & \beta Personalized_j \times PostPolicy_t + \alpha Personalized_j \\ & + \theta_1 MediaAttention_t \times Personalized_j \\ & + \gamma_k + \delta_t + \epsilon_j \end{aligned} \tag{1}$$

$Personalized_j$ is an indicator variable equal to one if the ad contained personalized content that reflected the data on celebrity affinity or undergraduate school used to target, 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 as are the vector of controls for the direct effect of the media $MediaAttention$. Because the ads are randomized, δ_t and γ_k should primarily improve efficiency. We estimate the regression using ordinary least squares. Following the strategy 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 an 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 0.023 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.[4]

[4]In earlier versions of this paper, we also showed the robustness of our results to a logit model. Since this was achieved by simply converting the aggregate-level data, the results were similar.

In Column (2), we add an additional interaction which controls for how many news stories there were that day 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 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. Column (3) reports results from our main specification which, combines both of these controls as summarized by equation (1). Though the estimated effect for $Personalized \times PostPolicy$ is smaller at 0.0174, it is still a significant increase relative to the average baseline click-through rate for personalized ads before the introduction of improved privacy controls of 0.007.

Table 3: Initial Results

	No Controls	Controls News	Combined	Lags	Lags 2	Specialized vs. Main	Word-Weighted
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Personalized × PostPolicy	0.0236**	0.0246**	0.0174**	0.0191***	0.0212***	0.0161***	0.0168***
	(0.0102)	(0.0100)	(0.00358)	(0.00218)	(0.00181)	(0.00144)	(0.00140)
Personalized	-0.0119*	-0.106	0.339	0.309	0.297	0.106	0.133
	(0.00627)	(0.106)	(0.187)	(0.151)	(0.139)	(0.176)	(0.232)
Personalized Ad × News Articles		0.0185	0.0510	0.0644	0.0660		
		(0.0214)	(0.0531)	(0.0746)	(0.0752)		
Personalized Ad × Google Searches			-0.139	-0.139	-0.134	-0.0200	-0.0260
			(0.0976)	(0.0988)	(0.0952)	(0.0460)	(0.0570)
Personalized Ad × News Articles (Lag)				-0.0102	0.0158		
				(0.0152)	(0.0194)		
Personalized Ad × News Articles (Lag 2)					-0.0305***		
					(0.00518)		
Personalized Ad × Major Newspapers						-0.00278	
						(0.00353)	
Personalized Ad × Specialized News						-0.0124	
						(0.00768)	
Personalized Ad × Major Newspapers (Lag)						0.00420	
						(0.00457)	
Personalized Ad × Specialized News (Lag)						-0.00312*	
						(0.00130)	
Personalized Ad× Major News (words)							-0.00153
							(0.00153)
Personalized Ad × Specialized News (words)							-0.00238
							(0.00195)
Date Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Targeted Group Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2730	2730	2730	2730	2730	2730	2730
R^2	0.060	0.060	0.061	0.061	0.061	0.061	0.060

OLS Estimates. Dependent variable is percentage daily click through rate for each of 79 campaigns.
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.

For the remainder of the table we investigate alternative ways of controlling for media coverage. Column (4) reports the results of a specification that includes lags for media stories appearing and Column (5) reports results of a specification that has a double lag for these media stories. This is motivated by work by Sinaceur et al. (2005), who examine how there are often delays in affective response to policy news - in their case, the advent of mad cow disease. In both cases these lags enhance the precision and increase the size of the estimates.

Column (6) reports a specification where we use data from Factiva that allows us to distinguish between news stories that were reported in the major press and the specialized press and allow the effect to vary with these different types of media. Column (7) reports a specification where we use a word-weighted measure for these controls, to reflect the extent of coverage rather than just the presence of coverage. Both of these alternative ways of specifying media presence produce a similarly sized estimate to that of our main specification in Column (3). In general, the results of Column (4)-(7) are reassuring that the measured effect of the introduction of privacy controls is robust to different ways of controlling for the extensive media coverage.

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. In addition, we only measure average effects of the policy - there may have been many Facebook users for whom the policy had essentially no effect as they had no intention of ever clicking on ads.

This empirical analysis uses a short time window of 5 weeks. This means it is less likely 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. The results for a specification with no controls, reported in Column (1) of Table 4, were positive but larger than for the full period.[5] 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 in Column (2) where we include controls from our main specification in Column (3) of Table 3 for news stories and general 'buzz' about the introduction of improved privacy controls.[6] This holds when we extend the media controls as we did in Columns (4)-(7) of Table 3. In Column (3) and (4) of Table 4 we examine the patterns when we exclude this immediate 10-day window around the policy change. Similarly, in Columns (5) and (6) we examine the patterns in a window that excludes the immediate 20 days around the policy change. The results are reasonably similar, though the news controls generally appear to increase the precision of our estimates for these windows outside the immediate time of the policy change.

Of course with any research that relies on a natural experiment for identification, there are open questions about what is precisely the local average treatment effect that is being estimated. In Table 3 we control as far as possible for media attention. The reason we do this is because of concern for the direct effect of media - if the media is telling users that Facebook is bad and intrusive, they may be less likely to click on personalized ads regardless of whether or not Facebook has privacy controls in place. The media controls we have in place (in particular the lags) should help control for this. What the media controls cannot do, however, is control for the fact that as a result of the media publicity all Facebook users

[5]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.

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

Table 4: Looking at different time windows

	10 Day		Excluding 10 Day		Excluding 20 Day	
	(1)	(2)	(3)	(4)	(5)	(6)
Personalized × PostPolicy	0.0554***	0.0454**	0.0149**	0.0197***	0.0235**	0.0205***
	(0.0208)	(0.0204)	(0.00711)	(0.00721)	(0.0105)	(0.00723)
Personalized	-0.0112	0.0189	-0.0141***	-0.0329	-0.00528**	0.0248
	(0.0115)	(0.0355)	(0.00464)	(0.0266)	(0.00248)	(0.0355)
Date Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Target Group Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
News+Search Controls	No	Yes	No	Yes	No	Yes
Observations	780	780	1872	1872	1326	1326
R^2	0.118	0.019	0.044	0.043	0.052	0.045

OLS Estimates. Dependent variable is percentage daily click-through rate for each of 79 campaigns. 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.

were aware of the change in the policy. This means that the correct interpretation here is that these estimates apply to a large firm that is under media scrutiny for its privacy policies and consequently can rely on the fact that its users found out about changes in its privacy policy. These estimates may be less applicable to small firms who cannot rely on the media to highlight changes in their privacy policy and controls for them.

3.1 Further Robustness Checks

The assumption that allows us to ascribe causality to our estimates is that there was no change in Facebook user and advertiser behavior and the external environment 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 the NPO, and on Facebook user and advertiser behavior.

The NPO 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. An obvious 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 NPO. 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 message.

To investigate this possibility, we obtained confidential data from the NPO, based on weekly update emails from Facebook that recorded how many people had become their 'fan' on Facebook and subscribed to their newsfeed. In the two weeks prior to the introduction of improved privacy controls, there was a 0.97 correlation between the number of fans and number of clicks. After the introduction of improved privacy controls, there was a 0.96 correlation between the number of fans and number of clicks. 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 NPO.

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. According to the figures on the advertising 'reach' made available to advertisers by Facebook, the potential audience did not decrease for any of the ads following the policy change, instead showing only a very small increase. In Table A-3 in the appendix, comScore data based on their panel of two million internet users suggests that this was not the case, and that there was actually an increase in the number of users. 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 privacy 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 A-5 in the appendix reports the results and shows that there was no change in which ads were shown based on their observable characteristics, though there may of course have been unobserved changes based on their unobserved characteristics.

It is also possible that, rather than a change in user composition, what we are measuring was 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. Again 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).

3.2 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 why addressing privacy concerns by improving privacy controls was associated with improved advertising performance is that it 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 were 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 were personalized around the fact that a user liked the Korean delicacy kimchi, then there are only 6,180 Facebook users who say that they 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 targeted at group k in the following manner:

$$
\begin{aligned}
ClickRate_{jt} = \beta_1 Personalized_j \times PostPolicy_t \times AdReach_k + \\
\beta_2 Personalized_j \times PostPolicy_t + \beta_3 PostPolicy_t \times AdReach_k + \\
\alpha_1 Personalized_j + \alpha_2 Personalized_j \times AdReach_k \\
\theta MediaAttention_t \times Personalized_j + \gamma_k + \delta_t + \epsilon_j
\end{aligned} \tag{2}
$$

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

Table 5 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 5 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* $\times$ *Personalized* $\times$ *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).

Columns (2)-(7) of Table 5 echo our earlier analysis from Table 3 by adding incremental controls for different types of media attention. The result remains robust to these different 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.

Table 5: Mechanism: Role of Rarity of Information

	No Controls	Controls	Lags	Lags 2	Specialized vs. Main	Word-Weighted
	(1)	(2)	(3)	(4)	(5)	(6)
Post-Policy × Personalized × Ad-Reach	-0.0852**	-0.0852***	-0.0852***	-0.0852***	-0.0852***	-0.0852***
	(0.0421)	(0.0312)	(0.0313)	(0.0312)	(0.0313)	(0.0312)
Post-Policy × Personalized	0.0317**	0.0262**	0.0273***	0.0293***	0.0242**	0.0250**
	(0.0147)	(0.0112)	(0.00996)	(0.0102)	(0.0107)	(0.0119)
Personalized	-0.0153**	0.294	0.306	0.293	0.102	0.129
	(0.00670)	(0.363)	(0.268)	(0.267)	(0.250)	(0.266)
Personalized × Ad-Reach	0.0354	0.0354*	0.0354*	0.0354*	0.0354*	0.0354*
	(0.0214)	(0.0209)	(0.0210)	(0.0210)	(0.0211)	(0.0210)
Post-Policy × Ad-Reach	0.0150	0.0150	0.0150	0.0150	0.0150	0.0150
	(0.0350)	(0.0204)	(0.0204)	(0.0203)	(0.0204)	(0.0204)
Personalized Ad × Number Facebook News Articles		0.0467				
		(0.0355)				
Personalized Ad × Google Facebook Privacy Searches		-0.124	-0.139**	-0.134**	-0.0200	-0.0260
		(0.108)	(0.0620)	(0.0616)	(0.0588)	(0.0607)
Personalized Ad × News Articles			0.0644*	0.0660*		
			(0.0342)	(0.0342)		
Personalized Ad × News Articles (Lag)			-0.0102	0.0158		
			(0.0206)	(0.0237)		
Personalized Ad × News Articles (Lag 2)				-0.0305*		
				(0.0176)		
Personalized Ad × Major Newspapers					-0.00278	
					(0.00955)	
Personalized Ad × Specialized News					-0.0124*	
					(0.00663)	
Personalized Ad × Major Newspapers (Lag)					0.00420	
					(0.00781)	
Personalized Ad × Specialized News (Lag)					-0.00312	
					(0.00918)	
Personalized Ad× Major News (words)						-0.00153
						(0.00282)
Personalized Ad × Specialized News (words)						-0.00238
						(0.00187)
Date Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Targeted Group Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2730	2730	2730	2730	2730	2730
R^2	0.062	0.062	0.062	0.063	0.062	0.062

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 fixed effects for the group targeted and also dropped from the specification.

3.3 The Role of Privacy Settings Usage

These results are reasonably compelling that the result is driven by the rarity of information used which in turn should be linked to privacy concerns. However, it is still is only indirectly linked to privacy controls. Though Facebook did not share information about whether users did indeed change their privacy settings using the newly introduced controls, we are able to use data on average usage of alternative privacy controls by different groups of Facebook users to examine whether the effect was indeed strongest for the kind of users who are most likely to use privacy controls.

It is not clear whether the effect will be largest for users who care about privacy controls. It may be that the kind of consumers who care about privacy would be most upset if they set highly protective privacy settings and still see highly personalized advertising. On the other hand, such consumers who are more aware of privacy matters on Facebook may have experienced the most reactance prior to Facebook addressing privacy complaints with the improved set of privacy policies. They may have also been the most reassured that they could now explicitly prevent Facebook from sharing their click data with third parties.

Specifically, we exploited the fact that Facebook has introduced a new product called 'social advertising'. These are ads which feature the name of friends who have 'liked' a webpage.[7] Facebook users can use the Facebook privacy controls to prevent their name being featured in these ads to their friends who may be targeted by a firm. This percentage of 'opt-outs' is shared with the advertising firm. Therefore, two years after the original data was collected, the non-profit briefly (at our request) ran new social ad campaigns that were targeted at each of the 39 original target groups. This allowed us to collect data for each of these target groups on what proportion of ads shown to the target Facebook users had the friend's name and image removed.

There are two obvious assumptions that we make when using this measure. The first is

[7] This new form of advertising is described in detail in Tucker (2012).

that because this ad product and data did not exist at the time of the natural experiment, we are assuming that a taste for using privacy controls remains static in our target groups over time, or at least that if there was a change, that all target groups changed at the same rate. This is based on empirical evidence that suggests that though the absolute level of privacy taste can change, the rate of change across demographic groups does not differ (Goldfarb and Tucker, 2012).

The second assumption is that a taste for privacy controls among the friends of people who are in the target group reflects the taste for privacy controls of the members of the target group themselves. The advertising data is constructed so that a firm can only ever observe usage of privacy controls by the friends of any target group it selects, not the target group itself. Therefore, we are implicitly assuming that there is a correlation across friends in terms of tastes for the use of privacy controls. Such homophily or correlations across friends in tastes for privacy has been documented in field data from online social networks by Acquisti and Gross (2006); Lewis et al. (2008). We recognize, though, that this is a data limitation.

On average, 14.9% of the ad impressions were affected by Facebook users opting-out of showing their name and photo. There was considerable variation in this statistic across the 39 target groups, varying from 56% to 0%.

We modify our equation (1) for the click-through rate $ClickRate_{jt}$ for ad j on day t targeted at group k in the following manner:

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

The results are reported in Table 6. The estimates for $Personalized_j \times PostPolicy_t \times PrivacyControlsUse_k$ are reasonably consistent across the different specifications that allow for different media controls. In each case they suggest that as more people in the group targeted use privacy controls, the larger is the effect of introducing privacy controls.

The fact that the effect size is mediated by the usage of privacy controls across the different groups targeted by the firms is direct evidence that the increase in personalized advertising effectiveness that we observed in Figure 1 was driven by customers' affinity for privacy controls, rather than by media coverage or other explanations of the effect.

Table 6: Mechanism: Role of Privacy Setting Usage

	No Controls	Controls	Lags	Lags 2	Specialized vs. Main	Word-Weighted
	(1)	(2)	(3)	(4)	(5)	(6)
Personalized × Postpolicy × PrivacyAwareScore	0.0942**	0.0942**	0.0942**	0.0942**	0.0942**	0.0942**
	(0.0419)	(0.0419)	(0.0418)	(0.0419)	(0.0419)	(0.0419)
Personalized × PostPolicy	0.0375***	0.0314***	0.0331***	0.0351***	0.0300**	0.0308**
	(0.0126)	(0.0113)	(0.0115)	(0.0117)	(0.0117)	(0.0128)
Personalized	-0.0119*	0.339	0.309	0.297	0.106	0.133
	(0.00627)	(0.266)	(0.268)	(0.267)	(0.250)	(0.266)
Postpolicy × PrivacyAwareScore	0.00394	0.00394	0.00394	0.00394	0.00394	0.00394
	(0.0398)	(0.0396)	(0.0396)	(0.0396)	(0.0396)	(0.0397)
Personalized Ad × News Articles		0.0510***	0.0644*	0.0660*		
		(0.0191)	(0.0341)	(0.0341)		
Personalized Ad × Google Searches		-0.139**	-0.139**	-0.134**	-0.0200	-0.0260
		(0.0619)	(0.0620)	(0.0615)	(0.0588)	(0.0606)
Personalized Ad × News Articles (Lag)			-0.0102	0.0158		
			(0.0206)	(0.0236)		
Personalized Ad × News Articles (Lag 2)				-0.0305*		
				(0.0175)		
Personalized Ad × Major Newspapers					-0.00278	
					(0.00954)	
Personalized Ad × Specialized News					-0.0124*	
					(0.00662)	
Personalized Ad × Major Newspapers (Lag)					0.00420	
					(0.00781)	
Personalized Ad × Specialized News (Lag)					-0.00312	
					(0.00918)	
Personalized Ad× Major News (words)						-0.00153
						(0.00282)
Personalized Ad × Specialized News (words)						-0.00238
						(0.00186)
Date Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Targeted Group Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2730	2730	2730	2730	2730	2730
R^2	0.061	0.062	0.062	0.062	0.062	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. $PrivacyControlUse$ is collinear with the fixed effects for the group targeted and also dropped from the specification.

3.4 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 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 used in advertising on reactance.

We recruited 178 survey-takers from Amazon's Mechanical Turk to take part in an online survey.[8] The survey takers were asked about their educational history. They were then taken to a simulated social networking website page where they saw an ad displayed that offered them a discount and were asked about their responses to the ad. The study has a 2×2 design, which varied how uniquely identifying the content of the ad was (Unique Information, Generic Information) and the level of privacy control that survey takers were told existed on the website. (No Privacy Controls, Privacy Controls).

In the 'Non-Unique Data' condition, users were offered a discount on the basis of the state of the high school they attended. In the 'Unique Data' condition, they were offered a discount on the basis of actual name of the high school they attended. We also varied users' perception of privacy control. In the 'Privacy Controls' condition, they were told that 'The website has been praised for the extent of control it gives its users over their personal information. To restrict access to personal information you need to use their easy-to-understand privacy settings. In the 'No Privacy Controls' condition, they were told that 'The website has been criticized for the lack of control it gives its users over their personal information. To restrict

[8] We excluded 21 survey takers who failed to input their high school or state name correctly in the pre-survey. We also excluded one subject who was homeschooled. The results are similar if not quite as precise when we include these subjects.

access to personal information you need to use their hard to understand system of privacy-controls.' 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 (4.43 vs 3.23, F(2,176)=4.48, p-value=0.0000).

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 (2013), 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 unique and personally identifying information 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 non-unique information condition where there was less reactance, which accords with the results reported in Table 5. We also use controls for age, education and use of Facebook as reported for survey takers, though the results are robust to their exclusion, as would be expected in a randomized design. Interestingly, none of these are significant, which is evidence against alternative explanations for the results of our natural experiment that are based on changes in user demographics.

We also asked respondents questions about how likely they were to respond positively to the ad. We asked them whether they were likely to click on the ad, visit the store website and use the discount coupon. Column (2) reports the results for click intent. 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 to an ad with unique data than to an ad using non-unique data. Column (3) analyzes whether they are likely to visit the store website. The results are very similar to Column (3), though slightly less significant. Column (4) shows that the results echo (though less precisely) for the measure about whether or not the person was likely to use the discount coupon. In general, 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. Also of interest is that in Column (2) the effect of privacy controls is negative for click intent. However, this significant estimate does not carry over for the other dependent variables. A potential explanation is that there were fewer people who contemplated clicking on the ad to investigate the ad due to privacy concerns when respondents were told that the privacy controls on the website had been praised.

Table 7: Lab Experiment Results

	(1)	(2)	(3)	(4)
	Reactance	Click Intent	Visit Store	Use Discount
Unique Information × Privacy Controls	-0.797**	1.664***	1.170**	0.993*
	(0.380)	(0.569)	(0.592)	(0.538)
Unique Information	0.948***	-0.889**	-0.703*	-0.757**
	(0.267)	(0.399)	(0.417)	(0.378)
Privacy Controls	0.267	-0.957**	-0.621	-0.637
	(0.271)	(0.407)	(0.425)	(0.385)
High School Grad	-0.0992	0.0963	-0.0868	-0.277
	(0.372)	(0.562)	(0.583)	(0.532)
College Grad	-0.186	0.464	0.120	0.101
	(0.362)	(0.546)	(0.568)	(0.517)
Post Grad	-0.429	0.512	0.423	0.381
	(0.424)	(0.642)	(0.666)	(0.607)
Facebook User	0.599*	-0.433	-0.673	-0.715
	(0.317)	(0.468)	(0.486)	(0.443)
Age	-0.00372	0.0143	0.00406	0.00305
	(0.0107)	(0.0157)	(0.0163)	(0.0148)
Constant	2.838***	3.047***	3.957***	4.033***
	(0.585)	(0.877)	(0.910)	(0.830)
R^2	0.099	0.064	0.038	0.054

OLS Estimates. 178 Respondents. Dependent variable is a seven-point scale as shown for reactance in columns (1)-(3) and various intent measures in columns (4)-(6).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

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 an actual 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 5 that involve endogeneity or selection.

4 Implications

This paper explores the consequences for data-rich websites, such as social networks, that try to support themselves through advertising, of addressing privacy concerns by offering users more control over their privacy. The paper uses data from a randomized experiment conducted by an NPO that was designed to explore the relative merits of targeted ads with generic text, and ads that used user information to personalize the 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. However, advertisers could still use the same personal data to target and personalize advertising messages with before and after the policy change.

Recent research (Fournier and Avery, 2011), has emphasized that to succeed in the new world of social media brands must relinquish control. This research parallels such findings by emphasizing that to succeed, web platforms need to give control over privacy settings to their users if they want to use user data to enhance their offerings.

Empirical analysis suggests that after this change in policy, Facebook users were nearly 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 social-networking sites. This has important consequences for the current policy debate, which views the introduction of privacy controls as being harmful to advertising outcomes (Goldfarb and Tucker, 2011b).

There are limitations to this research. First, the randomized experiment was conducted by an NPO with an appealing cause. Consumers may be ready to ascribe less pernicious motives to an NPO 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 are studying a particular platform with a specific business model and community of users and the results cannot be extrapolated more generally to other websites or situations without further analysis. Fourth, we do not know how long the positive effects we measured after the introduction of privacy controls for personalized advertising persisted. Last, 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.

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

Barefoot, D. and J. Szabo (2008, April). Facebook ad click-through rates are really pitiful.

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.

Brandimarte, L., A. Acquisti, and G. Loewenstein (2012). Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science*.

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.

Culnan, M. and P. Armstrong (1999, Jan-Feb). Information privacy concerns, procedural fairness, and interpersonal trust: An empirical investigation. *Organization Science 10*(1), 104–115.

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.

Fournier, S. and J. Avery (2011). The uninvited brand. *Business Horizons 54*, 193–207.

Fusilier, M. and W. Hoyer (1980). Variables affecting perceptions of invasion of privacy in a personnel selection situation. *Journal of Applied Psychology 65*(5), 623–626.

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, May). Online display advertising: Targeting and obtrusiveness. *Marketing Science 30*, 389–404.

Goldfarb, A. and C. Tucker (2012). Shifts in privacy concerns. *American Economic Review 102*(3), 349–53.

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?

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

Lamberton, C. (2013). A spoonful of choice: How allocation increases satisfaction with tax payments. *Journal of Public Policy and Marketing*, Forthcoming.

Lewis, K., J. Kaufman, and N. Christakis (2008). The taste for privacy: An analysis of college student privacy settings in an online social network. *Journal of Computer-Mediated Communication 14*(1), 79–100.

Lohr, S. (2010, April 30). Privacy concerns limit online ads, study says. Technical report.

Malhotra, N. K., S. S. Kim, and J. Agarwal (2004). Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research 15*(4), 336–355.

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.

Marshall, J. (2011, May 10). Facebook served a third of display impressions in Q1.

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.

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.

Sinaceur, M., C. Heath, and S. Cole (2005, March). Emotional and deliberative reactions to a public crisis: Mad cow disease in France. *Psychological science 16*(3), 247–54.

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.

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.

Tucker, C. (2012). Social Advertising. *Mimeo, MIT*.

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.*

Xu, H., H.-H. Teo, B. C. Y. Tan, and R. Agarwal (2012). Effects of individual self-protection, industry self-regulation, and government regulation on privacy concerns: A study of location-based services. *Information Systems Research*.

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

Table A-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)
Targeted Group 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
$Ad\text{-}Reach_k$ and *School* are collinear with the fixed effects for the group targeted and also dropped from the specification.
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 that 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 groups targeted with the *Postpolicy* indicator. None of these interactions were significant.

B The change in privacy controls

Figure A-1: Facebook: Screenshots of Privacy Options before and after the introduction of privacy controls

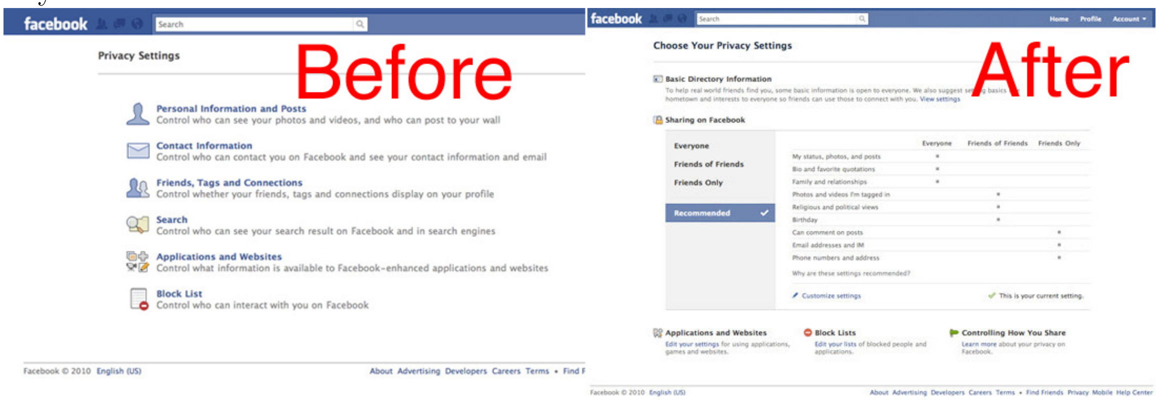

Source: Gawker Media

A-1 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