CATHERINE E. TUCKER*

This article investigates how Internet users' perceptions of control over their personal information affect how likely they are to click on online advertising on a social networking website. The analysis 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 text 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 that were more likely to use opt-out privacy settings.

Keywords: privacy, online advertising, social networks

Online Supplement: http://dx.doi.org/10.1509/jmr.10.0355

Social Networks, Personalized Advertising, and Privacy Controls

Many Internet firms collect a large amount of personal data from their users and use these data to allow advertisers to target and personalize ads. Consumers might perceive personalized ad content on such sites as more appealing and more aligned with their interests (Anand and Shachar 2009; Lambrecht and Tucker 2013), but they also may view it as both creepy and off-putting if they believe that the firm violated their privacy (Stone 2010). These privacy concerns may lead to "reactance," such that consumers resist the ad's appeal (White et al. 2008). Reactance is a motivational state in which consumers resist something they find coercive by behaving in the opposite way to that intended (Brehm 1966, 1989; Clee and Wicklund 1980).

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 reduction in reactance holds even if the controls are only tangentially related to the area in which reactance is evoked (Rothbaum, Weisz, and Snyder 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, the risk is 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 leads to an empirical question as to how strengthening privacy controls affects advertising performance.

*Catherine E. Tucker is Mark Hyman Jr. Career Development Professor and Associate Professor of Management Science, MIT Sloan School of Management, Massachusetts Institute of Technology, and Research Associate, National Bureau of Economic Research (e-mail: cetucker@mit.edu). The author thanks the Time Warner Cable Research Program on Digital Communications and the Net Institute for financial support. This research was also financially supported by NSF CAREER Award No. 6923256. The author also thanks Alessandro Acquisti; Emilio Calvano; Avi Goldfarb; Garrett Johnson; Cait Lambert; Alex Marthews; Markus Mobius; Martin Peitz; Ken Wilbur; and seminar participants at the American Economic Association meetings, Dartmouth, Harvard, the National Bureau of Economic Research, the National University of Singapore, New York University, University of Florida, University of Mannheim, and University of Munich. All errors are hers alone. Ron Shachar served as guest editor for this article.

ISSN: 0022-2437 (print), 1547-7193 (electronic)

I explore this question using data from a randomized field experiment conducted by a U.S.-based nonprofit to optimize its 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 to improve education for women in East Africa. The nonprofit randomized whether the ad copy was explicitly personalized to match data from the user's profile. In one condition, consumers saw ads that, for example, explicitly mentioned a celebrity or an undergraduate institution that was featured in their Facebook profile. In other conditions, they saw just generic text.

In the middle of the field experiment, in response to mounting media criticism, Facebook changed its 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 used to determine which advertising was displayed, targeted, and personalized, because the advertising platform used anonymous data. So, before and after the policy change, advertisers could still choose to target ads to 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 nonprofit had not anticipated such a change when it launched its field test of its ads. However, the fact that the change occurred midway 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.

The analysis spans five weeks of campaign-level click-through data on both sides of the introduction of the new privacy controls. Surprisingly, the findings reveal that personalized advertising was relatively ineffective before Facebook introduced new privacy controls. However, personalized advertising was nearly twice as effective at attracting users to the nonprofit's Facebook page after the shift in Facebook's policy, which gave users more control over their personal information. There was no significant change in the effectiveness of advertising that was shown to the same people but used a generic message during the period. This latter finding is to be expected because such ads do not make clear to consumers whether their private information is being used to target them.

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. There was no significant change in the ads shown, the user composition of Facebook, use of the website, or advertiser behavior during the period studied. There was also no change in how likely people were to sign up for the nonprofit's news feed after clicking on the ad, suggesting that the result is not an artifact of curiosity. Another concern is that there was a great amount of publicity around the policy change, so the analysis deploys many different controls for media coverage.

Building on existing research that documents that reactance to personalized advertising is greatest when the information used is more unique (White et al. 2008), this article then explores whether the positive effect of improved privacy controls was greatest for ads that used more unique information. Although some celebrities in the data (e.g., 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. Personalization was relatively more effective for personalized ads that used unusual information after privacy controls were enhanced. This finding 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.

The analysis then explores whether the measured effect depends on the extent to which consumers use privacy settings on Facebook. This is empirically ambiguous. On the one hand, consumers who care about privacy and use privacy settings may be more upset if they set high levels of privacy restrictions and still receive highly personalized advertising. On the other hand, consumers who are more aware of privacy matters on Facebook may have experienced the most reactance before Facebook addressed 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. Empirical analysis suggests that the measured 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 measured effect is associated with a desire for privacy controls rather than other external factors.

CONTRIBUTION

These findings contribute in four ways. First, to my knowledge, this is the only article to date that uses field data to examine the consequences of advertising outcomes for a firm responding to user privacy concerns by introducing improved privacy controls. Turow et al. (2009) find that 66% of Americans do not want marketers to tailor ads to their interests. In turn, fear of such resistance has led advertisers to limit the use of tailored ads (Lohr 2010). Therefore, the finding that positive effects arise for an advertising platform, in this instance, from addressing users' privacy concerns is useful.

Second, these findings have implications for privacy regulation. Currently, proposed regulations governing online behavioral advertising in the United States, such as "Do Not Track," pertain to 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; Tucker 2012a). In contrast, the results in this article show that in this setting, when a social networking website allowed customers to choose how their personally identifiable information was shared and used, there was no negative effect on advertising performance. This is an important finding for policy makers deciding whether to emphasize user-based controls in privacy regulation both in the United States and elsewhere.

Third, this article builds on research that focuses on more general questions about the role of control in mediating privacy concerns. 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 shown how individual-level control can mediate privacy concerns (Culnan and Armstrong 1999; Fusilier and Hoyer 1980; Malhotra, Kim, and Agarwal 2004).[1] This remains the case even if the control is merely perceptual or over tangential information and access to the focal data remains unchanged (Brandimarte, Acquisti, and Loewenstein 2012; Spiekermann, Grossklags, and Berendt 2001; Xu 2007; Xu et al. 2012). The contribution of the current article to this literature is to demonstrate empirically that perceptual control over privacy can affect responsiveness to advertising in addition to social and personal interactions.

Fourth, this article contributes to the online advertising literature that examines 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, research has labeled social networking websites as problematic venues for advertising because of extremely low click-through rates (Holahan 2007). This article suggests that if such sites are successful at reassuring consumers that they are in control of their privacy, firms can use ad personalization to generate higher click-through rates. Previous studies in marketing on social networking have explored how offline social networks can be used to target (Manchanda, Xie, and Youn 2008), how such social networking sites can use advertising to obtain members (Trusov, Bucklin, and Pauwels 2009), how social networks can be used to target ads (Tucker 2012b), and how makers of applications designed to be used on social networking sites can best advertise their products (Aral and Walker 2011). Beyond social networks, Goldfarb and Tucker (2011a) show that privacy concerns can influence ad effectiveness.

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 U.S. residents and also engaging their enthusiasm and support for its programs. To do so, the nonprofit set up a Facebook page to explain its mission and allowed people to view 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 began advertising using Facebook's own advertising platform. Initially, it ran an untargeted ad campaign that displayed an ad in April 2010 to all users of Facebook who lived in the United States and were 18 years of age and older. This campaign experienced a low click-through rate and attracted fewer than five new "fans" to the website. The disappointing results of this campaign led the nonprofit to determine whether it could engage further with its potential supporters by both targeting and personalizing ad content.

Randomized Campaign

The nonprofit 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. In an effort to protect the privacy of its supporters, the nonprofit has asked for the names of both the actual celebrities and the schools used in this advertising campaign to remain anonymous. Examples could be Oprah Winfrey, who has set up a girls' school in South Africa, or Serena Williams, who was a supporter of the U.S.-registered nonprofit Build African Schools. Data from the Facebook advertising interface suggest that there was little overlap in fans across these different groups.

To establish whether Facebook user data should be used merely to target ads or should also be used to personalize the content of the advertising message, the nonprofit decided to experiment with two 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 nonpersonalized case, the ad was similar in content but did not explicitly mention the undergraduate institution or the celebrity's name used to target the ad. In both cases, the baseline or "nonpersonalized" message was not completely generic but instead alluded to a broad user characteristic. Therefore, the estimates reflect the incremental benefit of personalized ad content that has specific and concrete personal information compared with ad content that uses non-

[1]Some evidence shows, however, that such individual controls impose costs on the systems considering adopting privacy controls (Campbell, Goldfarb, and Tucker 2015; Miller and Tucker 2009).

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 nonpersonalized	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.

specific and nonconcrete information. In each case, the ad was accompanied by the same picture of a girl who had been helped by the program. In line with the work of Small and Verrochi (2009), this girl had a sad expression.

The nonprofit also continued to use as its baseline an untargeted campaign that reached out to all adult U.S. 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." All campaigns were restricted to Facebook users who live in the United States and were 18 years of age and older who were not already fans of the nonprofit. The nonprofit set a daily maximum spending cap on advertising campaigns. It also agreed to pay at most $.50 for each click produced by the different advertising campaigns.

The Introduction of Improved Privacy Controls

A unique and potentially valuable aspect of 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 chief executive officer 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. Evaluating this change was not the purpose of the randomized field experiment, but it presents a unique opportunity to examine how a change in user privacy controls in social networking sites could change consumer responses to advertising because the nonprofit 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) points 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. As Table A1 in Appendix A details, Facebook had previously acted to reduce the amount of control users had over their data and had attracted negative publicity for doing so. In December 2009, ten privacy groups filed a complaint with the Federal Trade Commission over changes to Facebook's privacy policy,[2] which included default settings that made users' status updates available potentially to all Internet users, as well as made users' friend lists publicly available.

Facebook's change in privacy interface had three major components. First, all privacy settings were aggregated into one simple control. Users no longer needed to deal with 170 granular options. As Figure B1 in Appendix B depicts, 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 for users to opt out with a single click from third-party applications accessing their personal information. In general, these changes were received favorably. For example, Chris Conley (2010), the chairman of the American Civil Liberties Union, 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 on Facebook were targeted or whether advertisers could use user information to personalize ads. Display advertising was treated separately because, at the time of this study, Facebook stated, "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" (see http://www.zdnet.com/blog/btl/facebook-settings-a-tale-of-two-defaults/44095). To reassure advertisers that the change would not adversely affect them, Facebook sent out an e-mail to them saying that "this change will not affect your advertising campaigns." (The full letter appears in Figure B2 in Appendix B.) This means that though users were given control over how much information was being shared publicly and the extent to which they could be tracked by third parties, the actual mechanism by which the ads tested were targeted and served did not change.

A consequence of examining a real-life firm shift in privacy policy is that, unlike with a lab experiment, many things were changed all at once. Therefore, ultimately, the measured effect is a combination of different privacy measures, as well as how they were reported and received in the press. In addition to an increased sense of control, the estimate captures other positive effects of Facebook's improved privacy policy, such as higher trust in the firm. Because these improvements in general perceptions of the firm may occur in response to any firm improving its privacy policies because of public criticism, the estimates of this article are still of managerial interest.

Data

The nonprofit shared daily data from Facebook on how well each of the ads performed for the duration of the experiment. For each of the 79 ad campaigns, daily data were collected on the number of times they were shown and the number of clicks they received. In total, these ads were shown to 1.2 million users and received 1,995 clicks. The nonprofit was not informed by Facebook whether the ads appeared on mobile devices or in news feeds, but it believes that the majority of ads appeared on the right-hand side of the Facebook screen. When users clicked on the ad, they were taken to the nonprofit's Facebook page. The data spanned 2.5 weeks on both sides of the introduction of privacy controls on May 28, 2010.

The data included the number of unique impressions (i.e., the number of users to whom the ad was shown) and the number of clicks each ad received. Each click came from a unique user. The data contained information on the total unique clicks per day but did not convey any more detailed information about the timing of clicks. The data also included information on the cost to the nonprofit per click and the imputed cost per thousand impressions. Finally, the data contained an "Ad Reach" variable, which measures the number of Facebook users who were eligible

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

to be shown the ad for any targeted ad campaign. This Ad Reach variable allows for exploration of the behavioral mechanism in subsequent regressions. To protect the privacy of the nonprofit's supporters, the data did not contain information about the backgrounds or identities of those who clicked "like" 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 evaluating the daily level, as many campaigns received no clicks on a given day, inflating the appearance of low click-through rates. However, this is similar to rates reported by other advertisers for Facebook ads. For example, Webtrends (2011) reports that average click-through rates were .051% in 2010.

Table 2 also reports summary statistics for the data used as controls for user exposure to news about Facebook and privacy. The first set of data came from Factiva on the number of newspapers with stories that contained the words "privacy" and "Facebook" and how many words each article had devoted to the topic. The data also indicated whether the news article appeared in a renowned newspaper. The second set of data came from Google Trends, which provided an index for the number of searches for "Facebook" and "privacy." This index lies between 0 and 100 and captures the relative volume of searches over time for particular search terms people used on Google's search engine in the United States.

ANALYSIS

Figure 1 displays the average click-through rates for each campaign before and after the introduction of improved privacy controls. Before the policy change, the personalized ads were underperforming relative to their generic counterparts. This is surprising, given the expectation that displaying personalized ad text would increase the ads' relevance and, consequently, their appeal. However, after the policy change, the more expected pattern prevailed, in which ads with personalized content were relatively more effective than generically worded but targeted ads or untargeted ads. This change was highly significant ($p = .0047$). The effects of targeted ads without personalized content before and after the introduction of improved privacy controls were not significantly different ($p = .407$). There appears to be little change in the effectiveness of the untargeted campaign, though with only one campaign, it is impossible to assess statistical significance when 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 of untargeted ads after the introduction of improved privacy controls.

Figure 1

COMPARISON OF CLICK-THROUGH RATES BEFORE AND AFTER POLICY CHANGE

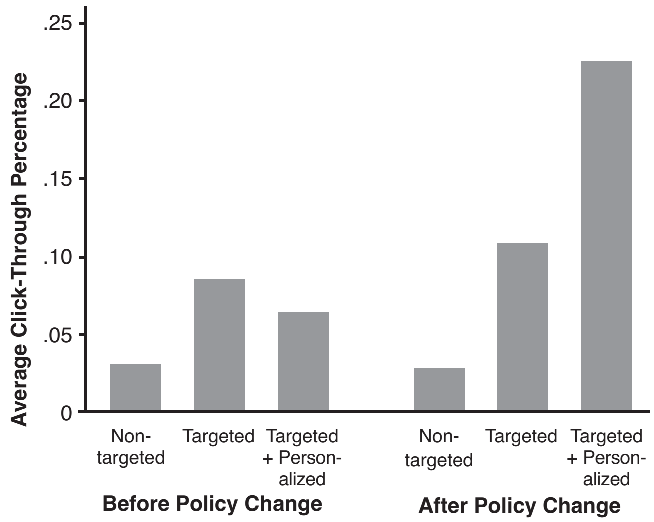

Figure 2 examines whether there were any differences in the campaigns targeted to undergraduate institutions and celebrities. On average, the celebrity-focused campaign was more successful at attracting clicks. However, there was a similar incremental jump in the effectiveness of personalized ads after the introduction of improved privacy controls for both categories of Facebook users: those with affinities for various schools and those with affinities for various celebrities.

Figure 1 suggests that the personalization of display ads was more effective after Facebook allowed users to take control of their personal information. Regression analysis, however, allows an assessment of both statistical significance and the use of more controls. In the regression, the

Table 2

SUMMARY STATISTICS

	M	*SD*	*Min*	*Max*
Average impressions	15,694	47,807	337	376,528
Average click-through (percentage points)	.0	.1	.0	3.1
Ad reach (000)	95	210	9,800	990,000
Average cost per click	.4	.1	.1	.5
Cost per 1,000 views	.1	.1	.0	.4
News articles containing words "privacy" and "Facebook"	61.1	39.7	10	210
Major newspaper articles containing words "privacy" and "Facebook"	3.2	1.5	1	7
Google searches	75.7	7.5	62	91
Words devoted to news	37,887	39,559	1,160	152,240

Notes: Campaign-level data; 79 campaigns (78 campaigns based on 39 different target groups each with personalized and targeted variants; 1 untargeted campaign). Table A2 in Appendix A clarifies the sources of each of these pieces of data.

Figure 2

COMPARISON OF CLICK-THROUGH RATES BEFORE AND AFTER BY TARGETING METHOD

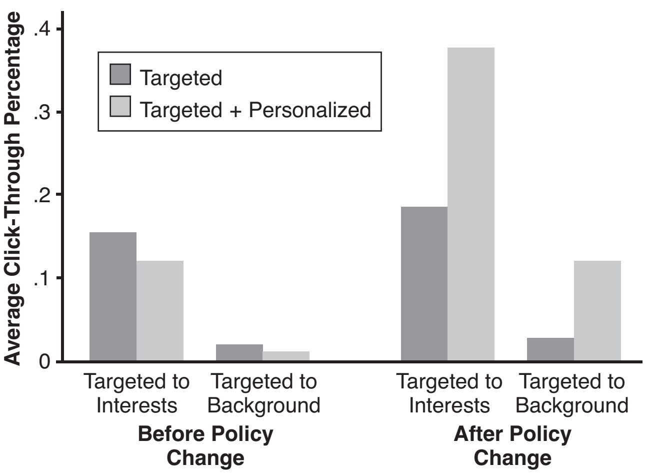

click-through rate $ClickRate_{jt}$ for ad j on day t targeted at group k is a function of

$$(1)\quad ClickRate_{jt} = \beta Personalized_j \times PostPolicy_t + \alpha Personalized_j + \theta_1 MediaAttention_t \times Personalized_j + \gamma_k + \delta_t + \varepsilon_j.$$

$Personalized_j$ is an indicator variable that equals 1 if the ad contained personalized content that reflected the data on celebrity affinity or undergraduate school used to target and 0 if there was no personalized content. $PostPolicy_t$ is an indicator variable that equals 1 if the date was after the privacy-settings policy change took place and 0 if otherwise. The coefficient β captures the effect of their interaction, θ captures the effect of various controls that allow the effectiveness of personalized advertising to vary with media attention, and γ_k is a vector of 39 fixed effects for the 20 undergraduate institutions and each of the 19 celebrities targeted. These all control for the underlying systematic differences in how likely people in that target segment were to respond to this charity. The specification includes a vector of date dummies δ_t. These dummies 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 media attention (MediaAttention). Because the ads are randomized, δ_t and γ_k should primarily improve efficiency. The specification is estimated with ordinary least squares. In line with Bertrand, Duflo, and Mullainathan's (2004) research, standard errors are clustered at the ad-campaign level to avoid artificially understating the size of the standard errors as a result of the use of repeated panel data.

Table 3 presents empirical estimates that incrementally build up to the full specification in Equation 1. Column 1 is an initial simplified specification. The crucial coefficient of interest is Personalized × PostPolicy, which captures how a Facebook user exposed to a personalized ad responds to this ad after Facebook's change in privacy policy relative to an ad that had generic wording. This coefficient suggests a positive and significant increase in the performance of personalized ads compared with merely targeted ads after the introduction of enhanced user privacy controls. The magnitude of the estimates suggests that the click-through rate increased by .024, compared with an average baseline click-through rate of .023 for personalized ads before the introduction of improved privacy controls. The negative coefficient of Personalized, which is marginally significant, suggests that before the change in privacy settings, personalized ads were less effective than ads that did not use personalized ad copy.[3]

Column 2 of Table 3 adds another interaction that controls for how many news stories appeared that day that contained the words "privacy" and "Facebook." Although news stories capture some general salience, they do not necessarily reflect the extent to which Facebook users processed and acted on news about privacy concerns on Facebook. Therefore, the analysis expands to include an additional control that captures the number of daily searches with the terms "privacy" and "Facebook" by the Google Trends index, reported on a scale from 0 to 100. In line with the results reported in Column 1, which were larger because of the media buzz surrounding the introduction of improved privacy controls, the key interaction of Personalized × PostPolicy is smaller in magnitude, though still statistically significant. Column 3 reports the results from the main specification, which combines both these controls as summarized in Equation 1. Although the estimated effect for Personalized × PostPolicy is smaller (.0174), it is still a significant increase compared with the average baseline click-through rate for personalized ads before the introduction of improved privacy controls (.007).

The other results reported in Table 3 investigate alternative ways of controlling for media coverage. Column 4 reports the results of a specification that allows for an additional day-delay lag in the effect of media stories appearing, and Column 5 reports the results of a specification that allows a lag for the effect of stories from both the previous day and the day preceding it. The motivation for these lags comes from Sinaceur, Heath, and Cole (2005), who examine delays in affective responses 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. It is noticeable that the effect of the second lag is significant and negative in Column 5, suggesting that there was a general negative effect on clicking behavior from Facebook media coverage two days before the user saw the ad. Speculatively, one explanation may be that multiday coverage augments the effect of critical media.

Column 6 of Table 3 reports a specification that uses data from Factiva, to distinguish news stories reported in the major press from those reported in the specialized press, and allows the effect to vary with these different types of media. Column 7 reports a specification that uses a word-weighted measure for these controls, to reflect the extent of coverage rather than just the presence of coverage. Specifically, this measure weights the effect of each article by how

[3]Previous versions of this article demonstrated robustness to a logit model. Because this was achieved by simply converting the aggregate-level data to the individual level, the results were similar.

Table 3

AFTER THE INTRODUCTION OF IMPROVED PRIVACY CONTROLS, PERSONALIZED ADVERTISING BECAME MORE EFFECTIVE

	No Controls (1)		*Controls News (2)*		*Combined (3)*		*Lags (4)*		*Lags 2 (5)*		*Specialized vs. Main (6)*		*Word Weighted (7)*	
Personalized × PostPolicy	.024**	(.010)	.025**	(.010)	.017**	(.004)	.019***	(.002)	.021***	(.002)	.016***	(.001)	.017***	(.001)
Personalized	–.012*	(.006)	–.106	(.106)	.339	(.187)	.309	(.151)	.297	(.139)	.106	(.176)	.133	(.232)
Personalized ad × news articles			.019	(.021)	.051	(.053)	.064	(.075)	.066	(.075)				
Personalized ad × Google searches					–.139	(.098)	–.139	(.099)	–.134	(.095)	–.020	(.046)	–.026	(.057)
Personalized ad × news articles (lag)							–.010	(.015)	.016	(.019)				
Personalized ad × news articles (lag 2)									–.031***	(.005)				
Personalized ad × major newspapers											–.003	(.004)		
Personalized ad × specialized news											–.012	(.008)		
Personalized ad × major newspapers (lag)											.004	(.005)		
Personalized ad × specialized news (lag)											–.003*	(.001)		
Personalized ad × major news (words)													–.002	(.002)
Personalized ad × specialized news (words)													–.002	(.002)
Date fixed effects	Yes		Yes		Yes		Yes		Yes		Yes		Yes	
Targeted group fixed effects	Yes		Yes		Yes		Yes		Yes		Yes		Yes	
Observations	2,730		2,730		2,730		2,730		2,730		2,730		2,730	
R^2	.060		.060		.061		.061		.061		.061		.060	

*$p < .10$.
**$p < .05$.
***$p < .01$.
Notes: Ordinary least squares estimates. Dependent variable is the percentage daily click-through rate for each of 79 campaigns. Robust standard errors are clustered at the ad level. $PostPolicy_t$ is collinear with the date fixed effects and is dropped from the specification.

many words it contains. This allows for a more profound effect of a longer article that took up more pages than a brief or paragraph-long article. Both alternative ways of specifying media presence produce a similarly sized estimate to that of the main specification in Column 3. In general, the results of Columns 4–7 help ensure 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-square across all specifications. This low level of explanatory power is shared by much of the online advertising literature (Goldfarb and Tucker 2011a; Reiley and Lewis 2009). A possible explanation is that consumers are skilled at avoiding online advertising when viewing web pages, which introduces measurement error (Drèze and Hussherr 2003) and requires researchers to assemble large data sets to measure effects precisely. In addition, this is simply the average effects of the policy—the policy may have had essentially no effect on many Facebook users because they had no intention of ever clicking on ads.

This empirical analysis uses a short time window of five weeks. Therefore, it is less likely that some long-term trend—for example, increasing user acceptance of ad personalization or "habituation" to privacy concerns—is driving the results. To show robustness to an even shorter window, Column 1 of Table 4 reports estimates exclusively 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, also reported in Column 1, were positive but larger than for the full period. Examination of similar 10-day windows in the prepolicy change period revealed no evidence of any significant change in preferences for personalized advertising. 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. Column 2 explores this by including controls from the main specification in Column 3 of Table 3 for news stories and general "buzz" surrounding the introduction of improved privacy controls. The results are also robust to excluding the period in which the service was rolled out and the days spanning announcement and implementation. Columns 3 and 4 of Table 4 examine the effect of excluding the immediate 10-day window around the policy change. Similarly, Columns 5 and 6 examine the effect of using 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 the estimates for these windows outside the immediate time of the policy change. Another noticeable pattern is that the point estimates are largest for the 10-day window immediately surrounding the period. One explanation could be that users were more likely to recall or remember the change in privacy settings close to the time the policy change occurred. However, this speculation is based on the relative size of point estimates—a combined regression suggests that there is no statistically significant difference in the size of the estimated point estimates.

With any research that relies on a natural experiment for identification, there are open questions about precisely what local average treatment effect is being estimated. The estimates include controls for the direct effect of media; if the media is telling users that Facebook is bad and intrusive, users may be less likely to click on personalized ads, regardless of whether Facebook has privacy controls in place. The media controls (in particular, the lags) should help control for this. What the media controls cannot do, however, is control for Facebook users' awareness of the change in policy as a result of the media publicity. This means that the correct interpretation of what is measured is that these estimates apply to a large firm that is under media scrutiny for its privacy policies and consequently can be fairly confident that its users found out about changes in its privacy policy. These estimates may be less applicable to small firms that cannot rely on the media to announce changes in their privacy policies for them.

Further Robustness Checks

The assumption that allows causality to be ascribed to the estimates is that there was no change in Facebook user and advertiser behavior or in the external environment that could provide an alternative explanation of the shift in advertising response that was not associated with the

Table 4
DIFFERENT TIME WINDOWS

	10-Day				*Excluding 10-Day*				*Excluding 20-Day*			
	1		*2*		*3*		*4*		*5*		*6*	
Personalized × PostPolicy	.055**	(.021)	.051*	(.023)	.015*	(.007)	.016*	(.008)	.024*	(.011)	.021**	(.007)
Personalized	–.011	(.012)	.011	(.040)	–.014**	(.005)	.286	(.312)	–.005*	(.002)	.025	(.036)
Date fixed effects	Yes		Yes		Yes		Yes		Yes		Yes	
Targeting variable fixed effects	Yes		Yes		Yes		Yes		Yes		Yes	
News + search controls	No		Yes		No		Yes		No		Yes	
Observations	780		780		1,872		1,872		1,326		1,326	
R^2	.118		.119		.044		.043		.052		.045	

*p < .05.
**p < .01.
Notes: Ordinary least squares estimates. Dependent variable is the percentage daily click-through rate for each of 79 campaigns. Robust standard errors are clustered at the ad level. $PostPolicy_t$ is collinear with the date fixed effects and is dropped from the specification.

change in privacy controls.[4] This section uses further external data to provide evidence for this assumption.

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 nonprofit. For example, an alternative explanation for the results is that after the introduction of improved privacy controls, consumers were more likely to click on ads that appeared too intrusive, to find out what data the advertisers had or how advertisers obtained their data, rather than responding positively to the advertising message.

The nonprofit shared weekly update e-mails from Facebook that recorded how many people had become a “fan” on Facebook and subscribed to its news feed, to help rule out this alternative explanation. In the two weeks before the introduction of improved privacy controls, there was a .97 correlation between the number of fans and the number of clicks. After the introduction of improved privacy controls, there was a .96 correlation between the number of fans and the number of clicks. There was no statistically significant difference between these two correlations, suggesting that after the introduction of improved privacy controls, people were not more likely to click on the ad even if they had no interest in the nonprofit. In general, the nonprofit considers the campaign a success, especially given the relatively small cost of the trial (less than $1,000). In its most recent fundraising campaign, approximately 6% of revenues from new donors came directly from its Facebook page.

A potential concern is that the results reflect a change in the number of Facebook users. 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 Facebook advertisers, the potential audience did not decrease for any of the ads following the policy change, instead showing a small increase. As Table A3 in Appendix A shows, comScore data based on its panel of two million Internet users suggest 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 shift did not seem to be more dramatic than that from April to May. The Grubbs (1969) test allows testing for outliers in the full year of data, with the caveat that there is a limited number of observations. The results of this test did not indicate that observed changes between May and June deviated from the expected normal distribution.

Although observed demographics were reasonably similar, there is always the possibility that the composition of Facebook users changed in an unobserved way and that this change influenced which ads 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 therefore whose fans were more likely to accept personalization using Facebook after the introduction of improved privacy controls. Empirically, the mix of ads displayed did not change over time. If the composition of ads did change, this could be the reason more consumers of that type were going online or, alternatively, why the same number of consumers were spending longer online. Table A4 in Appendix A shows that there was no change in terms of which ads were shown based on their observable characteristics, though there may have been unobserved changes based on their unobserved characteristics.

It is also possible that, rather than a change in user composition, the measured effect was driven by a dramatic change in how people use Facebook. For example, an alternative explanation of the results is 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 noncommercial content. Table A5 in Appendix A presents data from Compete Inc. about how users’ browsing behavior on Facebook changed in 2010. The data show that no large or dramatic changes in users’ browsing behavior occurred in the period studied, compared with the natural fluctuations that are apparent for the rest of the year. Again, the Grubbs (1969) test for outliers indicated that the postpolicy 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 withdrew from Facebook as a result of the negative publicity surrounding the privacy interface, meaning that perhaps fewer advertisers were competing to personalize advertising, which made the personalized ads relatively more attractive. Though not a direct test, there is evidence against this interpretation in the pricing data for the ads. If there had been a drop-off in advertisers, this should have translated to a decrease in the price paid in the ad auction because the price is theoretically 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 is not statistically significant ($p = .59$).

Mechanism: Rarity of User Information

Edwards, Li, and Lee (2002) and White et al. (2008) show that personalized ads can lead to a process of reactance (Brehm 1966), in which consumers deliberately resist ads 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 the proposed mechanism, this analysis exploits the findings from prior studies that reactance to personalized advertising is greater for ads that use more unique information about the consumer (White et al. 2008). For example, Facebook has roughly 3,167,540 users whose profiles indicate that they like cooking. Therefore, if an ad were personalized around a user’s love of cooking, which can be fairly common, he or she might consider the use of such information less intrusive and consequently be less likely to react. However, if an ad were personalized around the user’s love of the Korean delicacy kimchi, which

[4]The Web Appendix replicates the natural experiment by using data 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 consumers had supplied previously, in contexts in which they believed that either they had control over their personal information or they did not. The results from this experiment confirm the previous findings and, by explicitly measuring stated reactance, provide support for a behavioral mechanism in which reactance is reduced for highly personal advertising if consumers perceive having control over their privacy.

is liked by only 6,180 Facebook users, he or she might be more concerned about the advertiser's use of private information, increasing perceptions of intrusiveness and consequently provoking reactance.

To explore this issue empirically, the analysis employed additional data on how many users were in the target group for that particular campaign. Equation 1 is expanded so that the click-through rate ClickRate_{jt} for ad j on day t targeted at group k is now a function of

$$(2)\quad \text{ClickRate}_{jt} = \beta_1\text{Personalized}_j \times \text{PostPolicy}_t \times \text{AdReach}_k + \beta_2\text{Personalized}_j \times \text{PostPolicy}_t + \beta_3\text{PostPolicy}_t \times \text{AdReach}_k + \alpha_1\text{Personalized}_j + \alpha_2\text{Personalized}_j \times \text{AdReach}_k + \theta\text{MediaAttention}_t \times \text{Personalized}_j + \gamma_k + \delta_t + \varepsilon_j.$$

AdReach_k, MediaAttention_t, and PostPolicy_t are collinear with the date and campaign fixed effects and thus are dropped from the equation.

Table 5 uses Equation 3 to investigate how the size of the ad reach (large or small) moderated the estimates—that is, the number of people to whom the ad could potentially be shown. Column 1 of Table 5 reports the moderating effect of the ad reach for the initial specification on the efficacy of personalized ads relative to simply targeted ads before and after the introduction of improved privacy controls. The negative coefficient on PostPolicy × Personalized × AdReach 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 the larger point estimate for PostPolicy × Personalized relative to Table 3, Column 1, shows.

Columns 2–6 of Table 5 echo the analysis in Columns 3–7 of Table 3 by adding incremental controls for different types of media attention. The result remains robust to these different controls. Ad reach includes millions of users. Therefore, roughly extrapolating from the linear functional form, the estimates suggest that for ads that had target audiences of greater than 243,000, the positive effect of the policy was largely 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 .03.

The Role of Privacy Settings Usage

So far, the empirical analysis provides reasonably compelling evidence that the increase in click-throughs for personalized ads is driven by the rarity of information used, which in turn is linked to privacy concerns. However, this association is only indirectly linked to privacy controls. Although Facebook did not share information about whether users indeed changed their privacy settings using the newly introduced controls, the nonprofit provided data on average usage of alternative privacy controls by different groups of Facebook users, allowing for an examination of whether the effect was strongest for users who were most likely to use privacy controls.

However, it is not clear whether this effect will be the greatest for users who care about privacy controls. It may be that such consumers would be more upset if they set highly protective privacy settings and still saw highly personalized advertising. Conversely, consumers who are more aware of privacy matters on Facebook may have experienced the most reactance before Facebook addressed the privacy complaints by improving privacy policies. They may also have been the most reassured that they could now explicitly prevent Facebook from sharing their click data with third parties.

Specifically, this new analysis exploits Facebook's introduction of a new product called "social advertising" after the events examined in this article took place (for a detailed description of this new form of advertising, see Tucker 2012b). Social advertising contains ads that feature the names of friends who have "liked" a web page. Facebook users can use the new privacy controls to prevent their names from being featured in ads to their friends. The percentage of people who opt out is shared with the advertising firm. Therefore, two years after the original data were collected, the nonprofit briefly (at my request) ran new social ad campaigns targeted to each of the 39 original target groups. This enabled measurement of the proportion of ads shown to the target Facebook users that had a friend's name and image removed.

Two assumptions underlie this measure. First, because this ad product and data did not exist at the time of the natural experiment, this procedure assumes that the desire for privacy controls remains static in the target groups over time or at least that if there was a change, all target groups changed at the same rate. This assumption is based on empirical evidence that suggests that although the absolute level of privacy desires can change, the rate of change across demographic groups does not differ (Goldfarb and Tucker 2012).

The second assumption is that a desire for privacy controls among the target group's friends also reflects the desire for privacy controls of the target group itself. The advertising data are constructed such that a firm can only observe usage of privacy controls by the friends of any target group it selects, not the target group itself. Therefore, an implicit assumption is that there is a correlation across friends in terms of the desire for privacy controls. Prior research has documented such homophily or correlations across friends in terms of privacy desires in field data from online social networks (Acquisti and Gross 2006; Lewis, Kaufman, and Christakis 2008).

On average, 14.9% of ad impressions were affected by Facebook users opting out of having their names and photos used. This statistic varied considerably across the 39 target groups, from 56% to 0%.

Including this measure of privacy control use modifies Equation 1 for the click-through rate ClickRate_{jt} for ad j on day t targeted at group k as follows:

$$(3)\quad \text{ClickRate}_{jt} = \beta_1\text{Personalized}_j \times \text{PostPolicy}_t \times \text{PrivacyControlsUse}_k + \beta_2\text{Personalized}_j \times \text{PostPolicy}_t + \beta_3\text{PostPolicy}_t \times \text{PrivacyControlsUse}_k + \alpha_1\text{Personalized}_j + \alpha_2\text{Personalized}_j \times \text{PrivacyControlsUse}_k + \theta\text{MediaAttention}_t \times \text{Personalized}_j + \gamma_k + \delta_t + \varepsilon_j.$$

Table 5
MECHANISM: ROLE OF RARITY OF INFORMATION

	No Controls (1)		*Controls (2)*		*Lags (3)*		*Lags 2 (4)*		*Specialized vs. Main (5)*		*Word Weighted (6)*	
PostPolicy × Personalized × Ad Reach	–.085**	(.042)	–.085***	(.031)	–.085***	(.031)	–.085***	(.031)	–.085***	(.031)	–.085***	(.031)
PostPolicy × Personalized	.032**	(.015)	.026**	(.011)	.027***	(.010)	.029***	(.010)	.024**	(.011)	.025**	(.012)
Personalized	–.015**	(.007)	.294	(.363)	.306	(.268)	.293	(.267)	.102	(.250)	.129	(.266)
Personalized × Ad Reach	.035	(.021)	.035*	(.021)	.035*	(.021)	.035*	(.021)	.035*	(.021)	.035*	(.021)
PostPolicy × Ad Reach	.015	(.035)	.015	(.020)	.015	(.020)	.015	(.020)	.015	(.020)	.015	(.020)
Personalized ad × number Facebook news articles			.047	(.036)								
Personalized ad × Google Facebook privacy searches			–.124	(.108)	–.139**	(.062)	–.134**	(.062)	–.020	(.059)	–.026	(.061)
Personalized ad × news articles					.064*	(.034)	.066*	(.034)				
Personalized ad × news articles (lag)					–.010	(.021)	.016	(.024)				
Personalized ad × news articles (lag 2)							–.031*	(.018)				
Personalized ad × major newspapers									–.003	(.010)		
Personalized ad × specialized news									–.012*	(.007)		
Personalized ad × major newspapers (lag)									.004	(.008)		
Personalized ad × specialized news (lag)									–.003	(.009)		
Personalized ad × major news (words)											–.002	(.003)
Personalized ad × specialized news (words)											–.002	(.002)
Date fixed effects	Yes		Yes		Yes		Yes		Yes		Yes	
Targeted group fixed effects	Yes		Yes		Yes		Yes		Yes		Yes	
Observations	2,730		2,730		2,730		2,730		2,730		2,730	
R^2	.062		.062		.062		.063		.062		.062	

*$p < .10$.
**$p < .05$.
***$p < .01$.
Notes: Ordinary least squares estimates. Dependent variable is percentage daily click-through rate for each of 79 campaigns. Robust standard errors are clustered at the ad level. $PostPolicy_t$ is collinear with the date fixed effects and is dropped from the specification. $AdReach_k$ is collinear with the fixed effects for the group targeted and is also dropped from the specification.

The results appear in Table 6. Note that these results pool data collected two years after the date of the original data with the original data to construct the interaction variable $\text{PrivacyControlsUse}_k$. The estimates for $\text{Personalized}_j \times \text{PostPolicy}_t \times \text{PrivacyControlsUse}_k$ are reasonably consistent across the different specifications that allow for different media controls. In each case, they suggest that as more people in the targeted group use privacy controls, the greater is the effect of introducing privacy controls. That the usage of privacy controls across the different targeted groups mediates the effect size provides further evidence that the increase in personalized advertising effectiveness observed in Figure 1 was driven by customers' affinity for privacy controls, rather than by media coverage or other explanations of the effect.

IMPLICATIONS

This article explores the consequences when data-rich websites that receive revenue through advertising, such as social networks, address privacy concerns by offering users more control of their privacy. The article uses data from a randomized experiment conducted by a nonprofit 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 conducted unexpectedly announced that it would change the interface through which users controlled their privacy settings. These changes, which were publicly lauded by consumer advocates, gave users greater control over what personally identifiable information was shared and whether third parties could track their movements. However, after the policy change, advertisers could still use the same personal data to target and personalize advertising messages they had used previously.

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

Empirical analysis suggests that after the policy change, 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 of their private information can benefit advertising-supported media and advertisers on social networking sites. This implication has important consequences for the current policy debate, which views the introduction of privacy controls as harmful to advertising outcomes (Goldfarb and Tucker 2011b).

This research has several limitations. First, the randomized experiment was conducted by a nonprofit with an appealing cause. Consumers may ascribe less pernicious motives to a nonprofit than to a for-profit company that uses personalized advertising. Second, this randomized experiment was conducted at a time when privacy concerns were particularly sensitive and salient in consumers' eyes. Although the analysis deployed controls for the extent of the publicity surrounding privacy and Facebook, it is not clear how the results might change if the introduction of controls is not so heavily publicized in the media. Third, the analysis focused on a particular platform with a specific business model and community of users, and therefore, the results cannot be extrapolated more generally to other websites or situations without further analysis. Fourth, it is not clear how long the measured positive effects persisted after the introduction of privacy controls for personalized advertising. Fifth, it is difficult to project general economic effects of such policy changes without more information on the penetration of personalized advertising on platforms such as Facebook. Last, the type of privacy control introduced by Facebook was just one of multiple ways that social networks or other advertising-supported websites can give control of privacy settings to their users. Notwithstanding these limitations, this article provides initial evidence that addressing consumers' privacy concerns is important for online advertising venues.

REFERENCES

Acquisti, Alessandro and Ralph Gross (2006), "Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook," in *Privacy Enhancing Technologies*, George Danezis and Philippe Golle, eds. Berlin: Springer Verlag, 36–58.

Anand, Bharat and Ron Shachar (2009), "Targeted Advertising as a Signal," *Quantitative Marketing and Economics*, 7 (3), 237–66.

Aral, Sinan and Dylan Walker (2011), "Creating Social Contagion Through Viral Product Design: A Randomized Trial of Peer Influence in Networks," *Management Science*, 57 (9), 1623–39.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan (2004), "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics*, 119 (1), 249–75.

Bilton, Nick (2010), "Price of Facebook Privacy? Start Clicking," *The New York Times*, (May 12), [available at http://www.nytimes.com/2010/05/13/technology/personaltech/13basics.html?_r=0].

Brandimarte, Laura, Alessandro Acquisti, and George Loewenstein (2012), "Misplaced Confidences: Privacy and the Control Paradox," *Social Psychological and Personality Science*, 4 (3), 340–47.

Brehm, Jack W. (1966), *A Theory of Psychological Reactance*. New York: Academic Press.

——— (1989), "Psychological Reactance: Theory and Applications," in *Advances in Consumer Research*, Vol. 16, Thomas K. Srull, ed. Provo, UT: Association for Consumer Research, 72–75.

Campbell, James D., Avi Goldfarb, and Catherine Tucker (2015), "Privacy Regulation and Market Structure," *Journal of Economics & Management Strategy*, forthcoming.

Clee, Mona A. and Robert A. Wicklund (1980), "Consumer Behavior and Psychological Reactance," *Journal of Consumer Research*, 6 (4), 389–405.

Conley, Chris (2010), "Facebook Addresses Several Privacy Problems," Blog of Rights, (May 26), [available at https://www.aclu.org/blog/content/facebook-addresses-several-privacy-problems].

Culnan, Mary and Pamela Armstrong (1999), "Information Privacy Concerns, Procedural Fairness, and Interpersonal Trust: An Empirical Investigation," *Organization Science*, 10 (1), 104–115.

Drèze, Xavier and Fracois-Xavier Hussherr (2003), "Internet Advertising: Is Anybody Watching?" *Journal of Interactive Marketing*, 17 (4), 8–23.

Edwards, Steven M., Hairong Li, and Joo Hyun Lee (2002), "Forced Exposure and Psychological Reactance: Antecedents

Table 6
MECHANISM: ROLE OF RARITY OF INFORMATION

	No Controls (1)		*Controls (2)*		*Lags (3)*		*Lags 2 (4)*		*Specialized vs. Main (5)*		*Word Weighted (6)*	
Personalized × PostPolicy × PrivacyControlUse	.094**	(.042)	.094**	(.0420)	.094**	(.042)	.094*	(.042)	.094**	(.042)	.094**	(.042)
Personalized × PostPolicy	.038***	(.013)	.031**	(.011)	.033***	(.012)	.035***	(.012)	.030**	(.012)	.031*	(.013)
Personalized	–.012*	(.006)	.339	(.266)	.309	(.268)	.297	(.267)	.106	(.250)	.133	(.266)
PostPolicy × PrivacyControlUse	.004	(.040)	.004	(.040)	.004	(.040)	.004	(.040)	.004	(.040)	.004	(.040)
Personalized ad × news articles			.051***	(.019)	.064*	(.034)	.066*	(.034)				
Personalized ad × Google searches			–.139**	(.062)	–.139**	(.062)	–.134**	(.061)	–.020	(.059)	–.026	(.061)
Personalized ad × news articles (lag)					–.010	(.021)	.016	(.024)				
Personalized ad × news articles (lag 2)							–.031*	(.018)				
Personalized ad × major newspapers									–.003	(.010)		
Personalized ad × specialized news									–.012*	(.007)		
Personalized ad × major newspapers (lag)									.004	(.008)		
Personalized ad × specialized news (lag)									–.003	(.010)		
Personalized ad× major news (words)											–.002	(.003)
Personalized ad × specialized news (words)											–.002	(.002)
Date fixed effects	Yes		Yes		Yes		Yes		Yes		Yes	
Targeted group fixed effects	Yes		Yes		Yes		Yes		Yes		Yes	
Observations	2,730		2,730		2,730		2,730		2,730		2,730	
R^2	.061		.062		.062		.062		.062		.061	

**p* < .10.
***p* < .05.
****p* < .01.
Notes: Ordinary least squares estimates. Dependent variable is percentage daily click-through rate for each of 79 campaigns. Robust standard errors are clustered at the ad level. $PostPolicy_t$ is collinear with the date fixed effects and is dropped from the specification. PrivacyControlUse is collinear with the fixed effects for the group targeted and is also dropped from the specification.

and Consequences of the Perceived Intrusiveness of Pop-Up Ads," *Journal of Advertising*, 31 (3), 83–95.

Fournier, Susan and Jill Avery (2011), "The Uninvited Brand," *Business Horizons*, 54, 193–207.

Fusilier, Marcelline and Wayne Hoyer (1980), "Variables Affecting Perceptions of Invasion of Privacy in a Personnel Selection Situation," *Journal of Applied Psychology*, 65 (5), 623–26.

Goldfarb, Avi and Catherine Tucker (2011a), "Online Display Advertising: Targeting and Obtrusiveness," *Marketing Science*, 30 (May), 389–404.

——— and ——— (2011b), "Privacy Regulation and Online Advertising," *Management Science*, 57 (1), 57–71.

——— and ——— (2012), "Privacy and Innovation," *Innovation Policy and the Economy*, 12 (1), 65–90.

Grubbs, Frank E. (1969), "Procedures for Detecting Outlying Observations in Samples," *Technometrics*, 11 (1), 1–21.

Holahan, Catherine (2007), "So Many Ads, So Few Clicks," *BusinessWeek*, (November 11), [available at http://www.businessweek.com/stories/2007-11-11/so-many-ads-so-few-clicks].

Lambrecht, Anja and Catherine Tucker (2013), "When Does Retargeting Work? Information Specificity in Online Advertising," *Journal of Marketing Research*, 50 (September), 561–76.

Lewis, Kevin, Jason Kaufman, and Nicholas 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, Steve (2010), "Privacy Concerns Limit Online Ads, Study Says," *The New York Times*, (April 30), [available at http://bits.blogs.nytimes.com/2010/04/30/privacy-concerns-limit-online-ads-study-says/].

Malhotra, Naresh K., Sung S. Kim, and James Agarwal (2004), "Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model," *Information Systems Research*, 15 (4), 336–55.

Manchanda, Puneet, Ying Xie, and Nara Youn (2008), "The Role of Targeted Communication and Contagion in New Product Adoption," *Marketing Science*, 27 (6), 961–76.

Marshall, Jack (2011), "Facebook Served a Third of Display Impressions in Q1," ClickZ, (May 10), [available at http://www.clickz.com/clickz/news/2069381/facebook-served-display-impressions-q1].

McAfee, R. Preston and John McMillan (1987), "Auctions and Bidding," *Journal of Economic Literature*, 25 (2), 699–738.

Miller, Amalia R. and Catherine Tucker (2009), "Privacy Protection and Technology Adoption: The Case of Electronic Medical Records," *Management Science*, 55 (7), 1077–1093.

Miller, Arthur R. (1971), *The Assault on Privacy: Computers, Data Banks, and Dossiers*. Ann Arbor: University of Michigan Press.

Reiley, David and Randall Lewis (2009), "Retail Advertising Works! Measuring the Effects of Advertising on Sales via a Controlled Experiment on Yahoo!" working paper, Yahoo! Research.

Rothbaum, Fred, John R. Weisz, and Samuel 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, Marwan, Chip Heath, and Steve Cole (2005), "Emotional and Deliberative Reactions to a Public Crisis: Mad Cow Disease in France," *Psychological Science*, 16 (3), 247–54.

Small, Deborah and Nicole Verrochi (2009), "The Face of Need: Facial Emotion Expression on Charity Advertisements," *Journal of Marketing Research*, 46 (December), 777–87.

Spiekermann, Sarah, Jens Grossklags, and Bettina 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: ACM, 38–47.

Stone, Brad (2010), "Ads Posted on Facebook Strike Some as Off-Key," *The New York Times*, (March 3), [available at http://www.nytimes.com/2010/03/04/technology/04facebook.html?_r=0].

Taylor, Shelley E. (1979), "Hospital Patient Behavior: Reactance, Helplessness, or Control?" *Journal of Social Issues*, 35 (1), 156–84.

Thompson, Suzanne C., Alexandria Sobolew-Shubin, Michael E. Galbraith, Lenore Schwankovsky, and Dana Cruzen (1993), "Maintaining Perceptions of Control: Finding Perceived Control in Low-Control Circumstances," *Journal of Personality and Social Psychology*, 64 (2), 293–304.

Trusov, Michael, Randolph E. Bucklin, and Koen Pauwels (2009), "Effects of Word-of-Mouth Versus Traditional Marketing: Findings from an Internet Social Networking Site," *Journal of Marketing*, 73 (September), 90–102.

Tucker, Catherine (2012a), "The Economics of Advertising and Privacy," *International Journal of Industrial Organization*, 30 (3), 326–329.

——— (2012b), "Social Advertising," working paper, Sloan School of Management, Massachusetts Institute of Technology.

Turow, Joseph, Jennifer King, Chris J. Hoofnagle, Amy Bleakley, and Michael Hennessy (2009), "Americans Reject Tailored Advertising and Three Activities that Enable It," working paper, School of Communication, University of Pennsylvania.

Webtrends (2011), "Facebook Advertising Performance Benchmarks & Insights," (January 31), [available at https://docs.google.com/viewer?url=http%3A%2F%2Ff.cl.ly%2Fitems%2F2m1y0K2A062x0e2k442l%2Ffacebook-advertising-performance.pdf].

White, Tiffany, Debra Zahay, Helge Thorbjornsen, and Sharon Shavitt (2008), "Getting Too Personal: Reactance to Highly Personalized Email Solicitations," *Marketing Letters*, 19 (1), 39–50.

Xu, Heng (2007), "The Effects of Self-Construal and Perceived Control on Privacy Concerns," in *28th Annual International Conference on Information Systems (ICIS 2007),Montreal, Canada*.

———, Hock-Hai Teo, Bernard C.Y. Tan, and Ritu 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*, 23 (4), 1342–63.

APPENDIX A: DATA APPENDIX

Table A1

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 1 × 1 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 to which the user was connected. Instead, they are 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 the 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' names and their friends' names, profile pictures, gender, user IDs, connections, and any content shared using the Everyone privacy setting.
May 12, 2010	*The New York Times* publishes an article that ignites a firestorm of negative press about Facebook and privacy (see Bilton 2010).
May 24, 2010	Facebook founder Mark Zuckerberg announces in an editorial in the *Washington Post* that Facebook will institute new privacy settings.
May, 26 2010	Facebook unveils new privacy settings in press event.
May 27, 2010	Facebook starts rollout of privacy settings. *The New York Times* publishes "A Guide to Facebook's New Privacy Settings."
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 A2

SOURCES OF DATA USED IN THE STUDY

Data Variable	*Data Source*
Impressions	Facebook campaign data provided by the nonprofit
Clicks	Facebook campaign data provided by the nonprofit
Ad reach (000,000)	Facebook campaign data provided by the nonprofit
Average cost per click	Facebook campaign data provided by the nonprofit
Cost per 1,000 views	Facebook campaign data provided by the nonprofit
Daily volume of news articles mentioning terms "privacy" and "Facebook"	Factiva database: http://global.factiva.com
Daily volume of words devoted to news	Factiva database: http://global.factiva.com
Google searches using terms "privacy" and "Facebook"	Google Trends: www.google.com/trends
Facebook users in that interest group using privacy controls for social advertising	Facebook campaign data provided by the nonprofit two years after the initial field test

Table A3

SMALL CHANGES IN FACEBOOK USER COMPOSITION

Proportion of Group	*April 2010*	*May 2010*	*June 2010*
Age <17 years	10.4	10.6	11.4
Age 18–24 years	19.2	19.4	18.6
Age 25–34 years	20.8	20.7	20.8
Age 35–44 years	20.4	19.9	19.9
Age 45–54 years	16.7	16.5	16.5
Age 55–64 years	8.0	8.1	8.1
Age 65+ years	4.6	4.8	4.7
Income <$15k	10.1	10.3	9.7
Income $15k–24k	6.2	6.1	5.9
Income $25k–39k	12.5	12.7	13.5
Income $40k–59k	22.1	22.0	24.2
Income $60k–74k	10.9	11.3	9.6
Income $75k–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 A4

TEST OF WHETHER THERE WAS A CHANGE IN THE TYPES OF ADS BEING SHOWN BEFORE AND AFTER THE INTRODUCTION OF IMPROVED PRIVACY CONTROLS

	(1)	*(2)*	*(3)*	*(4)*
PostPolicy	–77.71 (117.9)	–108.6 (236.1)	–17.32 (70.60)	26.90 (163.6)
PostPolicy × School indicator		–60.14 (241.3)		71.91 (170.9)
PostPolicy × Ad Reach			633.0 (923.0)	710.1 (908.0)
Targeted group fixed effects	Yes	Yes	Yes	Yes
Observations	2,730	2,730	2,730	2,730
R^2	.050	.050	.051	.051

Notes: Ordinary least squares estimates. The dependent variable is number of times each ad is shown. Robust standard errors are clustered at the ad level. $AdReach_k$ and School are collinear with the fixed effects for the group targeted and are dropped from the specification. Column 2 suggests that it was not the case that more ads associated with undergraduate institutions were shown after the introduction of improved privacy controls. Column 3 suggests that it was not the case that ads that had a larger potential reach (e.g., ads associated with famous celebrities) were shown more frequently after the introduction of improved privacy controls. Column 4 combines these two measures and again indicates no significant change after the policy in terms of what ads were shown. A specification of the interaction between the 39 groups targeted and the PostPolicy indicator produced no significant effects.

Table A5

LITTLE CHANGE IN HOW FACEBOOK USERS USED THE WEBSITE

Date	*Average Stay*	*Visits/Person*	*Pages/Visit*
December 2009	21:29	22.27	29.46
January 2010	23:06	22.15	33.52
February 2010	22:14	21.08	35.33
March 2010	21:30	23.40	29.00
April 2010	21:54	23.27	25.45
May 2010	22:39	24.90	27.27
June 2010	21:50	24.37	24.78
July 2010	22:28	24.61	28.64
August 2010	22:28	26.86	30.33
September 2010	22:25	26.12	27.49
October 2010	24:30	26.52	24.64
November 2010	24:56	26.55	23.86
December 2010	25:48	26.46	24.24

Source: Compete Inc.

APPENDIX B: THE CHANGE IN PRIVACY CONTROLS

Figure B1

SCREENSHOTS OF FACEBOOK'S PRIVACY OPTIONS BEFORE AND AFTER THE INTRODUCTION OF PRIVACY CONTROLS

Source: Gawker Media.

Figure B2

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 byproviding 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