

Why Am I Seeing This Ad? The Effect of Ad Transparency on Ad Effectiveness

TAMI KIM
KATE BARASZ
LESLIE K. JOHN

Given the increasingly specific ways marketers can target ads, consumers and regulators are demanding ad transparency: disclosure of how consumers' personal information was used to generate ads. We investigate how and why ad transparency impacts ad effectiveness. Drawing on literature about offline norms of information sharing, we posit that ad transparency backfires when it exposes marketing practices that violate norms about "information flows"—consumers' beliefs about how their information *should* move between parties. Study 1 inductively shows that consumers deem information flows acceptable (or not) based on whether their personal information was: 1) obtained within versus outside of the website on which the ad appears, and 2) stated by the consumer versus inferred by the firm (the latter of each pair being less acceptable). Studies 2 and 3 show that revealing unacceptable information flows reduces ad effectiveness, which is driven by increasing consumers' relative concern for their privacy over desire for the personalization that such targeting affords. Study 4 shows the moderating role of platform trust: when consumers trust a platform, revealing acceptable information flows increases ad effectiveness. Studies 5a and 5b, conducted in the field with a loyalty program website (i.e., a trusted platform), demonstrate this benefit of transparency.

Keywords: information flows, transparency, targeted advertising, privacy, personalization, ad performance

With marketers' ever-expanding capacity to track online behaviors and display relevant ads, consumers have become increasingly wary of the ways firms gather

Tami Kim is an assistant professor in the Marketing Department at the Darden School of Business, FOB 191B, 100 Darden Blvd. Charlottesville, VA 22903 (kimt@arden.virginia.edu). Kate Barasz is an assistant professor in the Marketing Department at IESE Business School, Av. de Pearson 21 Barcelona, Spain 08034 (kbarasz@iese.edu). Leslie K. John is an associate professor of Business Administration in the Negotiations, Organizations, and Markets Unit at Harvard Business School, Bloomberg Center 462, Soldiers Field Road, Boston, MA 02163 (ljohn@hbs.edu). Please address correspondence to Tami Kim. This research is based on the lead author's dissertation, which was conducted at Harvard Business School. The authors would like to thank the review team and members of the NERD Lab for their invaluable comments. We also gratefully acknowledge Charlotte Blank and her team at Maritz for their field collaboration.

Editor: Eileen Fischer

Associate Editor: Andrea Morales

Advance Access publication May 7, 2018

and use their personal information. For example, the revelation that department stores were surreptitiously tracking customers' in-store movements using cell phone data sparked controversy in the media and beyond (Clifford and Hardy 2013). Target Inc. made headlines for sending pregnancy-related coupons to a teenage customer based on her recent purchases. Although the inference was accurate—the teenager was, in fact, pregnant—the targeting elicited controversy in part because the customer's family had been unaware of her pregnancy (Hill 2012). More broadly, and unbeknownst to many, the sharing of such consumer information amongst firms is extensive: for instance, Facebook purchased data on 70 million US households, enabling the firm to tailor ads based on users' purchases (Wasserman 2012).

Given the invasiveness of such practices, many consumers and regulators are increasingly demanding ad transparency—the disclosure of the ways in which firms collect and use consumer personal data to generate behaviorally

targeted ads—in an effort to empower consumers and better ensure above-board marketing practices (Greiff 2016; Ramirez et al. 2014; Turov et al. 2009). In recent years, many firms have voluntarily instituted such practices. Facebook introduced a feature allowing users to find out why any given ad is being shown to them; for example, a user may be informed that an ad for Acme Home Theaters appears based on what she does on Facebook, such as the pages she has “liked” and the ads she has clicked on (appendix A). Relatedly, a growing number of advertisers have been voluntarily displaying the YourAdChoices icon, a blue symbol indicating that the accompanying ad has been targeted based on the recipient’s characteristics (Alliance 2014); similar to Facebook, consumers can find out why the given ad is displayed to them by clicking the icon. In the same spirit, many websites have begun to explicitly alert visitors when tracking software (e.g., cookies) is in use, to post privacy policies (Culnan 2000), and to display logos (“privacy seals”) signaling that privacy standards have been met (Farmer 2015).

Although some firms have been forthcoming with this information, others have resisted, concerned that drawing attention to potentially off-putting practices may undermine ad effectiveness (Edwards 2009; Learmonth 2009)—that is, the extent to which an ad increases a consumer’s interest in a product. Perhaps reflecting this concern, some firms have implemented weak forms of ad transparency, merely making information available to consumers should they actively seek it out (e.g., by proactively clicking on the “Why am I seeing this ad?” button). How concerned should firms be about making ad practices transparent? How does ad transparency impact ad effectiveness? When might it reduce versus improve ad effectiveness? Might it improve ad effectiveness under certain conditions? In this article, we explore these questions, developing an understanding of 1) how consumers feel about specific advertising practices, and 2) how these perceptions change the way consumers engage with the ad.

We study these questions by testing the impact of actively and conspicuously revealing information about ad practices to consumers (e.g., by presenting an ad and its ad practice alongside each other). Such conspicuous disclosure is uncommon in today’s marketplace; digital advertisements are not usually accompanied by information on how they were generated, and when they are, this information is typically inconspicuous, merely made available for the motivated consumer to find. However, our experiments are designed to understand how consumers react to ad transparency, conditional on exposure. The topic of when and why consumers seek out such information, including potential individual differences in this propensity, is an interesting and important topic but one that is beyond the scope of the current article. (We further discuss consumers’ demand for ad transparency in the General Discussion,

including what it implies for the generalizability of our findings.)

We situate our account in research on offline norms of information sharing and disclosure and, using a combination of induction and deduction, develop a conceptual account of how and why ad transparency affects ad effectiveness. Specifically, our account demonstrates that ad transparency backfires when it exposes marketing practices that violate norms of “information flows”—consumer beliefs of how their information *ought* to flow between parties. Moreover, we address why this occurs: revealing information flows that consumers deem unacceptable reduces ad effectiveness by increasing consumers’ relative concern for their privacy over their interest in the increased personalization that such targeting affords. Finally, we demonstrate that platform trust enhances the upside of ad transparency: when trust in a given platform is high, ads that are displayed on this platform and reveal underlying acceptable information flows outperform nontransparent ones.

CONCEPTUAL DEVELOPMENT

In contrast to traditional targeted advertising, such as placing ads in specific markets or broadcasting in targeted time slots, online ad platforms allow advertisers to target consumers with greater efficiency, specificity, and accuracy. For example, the modern marketer can gather and integrate individual-level data to facilitate behaviorally targeted ads—ads based on the recipient consumer’s online behavior. Such developments have added many new and effective tools to the marketer’s arsenal, including behavioral retargeting (displaying ads for products consumers recently viewed; Lambrecht and Tucker 2013), content-based targeting (displaying ads based on what consumers read; Zhang and Katona 2012), and keyword-based targeting (displaying ads based on the terms consumers entered into search engines; Desai, Shin, and Staelin 2014; Sayedi, Jerath, and Srinivasan 2014; Yang, Lu, and Lu 2014), all of which have improved marketers’ capacity to reach and persuade consumers.

However, better-targeted ads necessarily require access to consumers’ personal data. This presents its own set of challenges with respect to how data are collected, which data are collected, and whether consumers would consent to their data being used this way. Most consumers are aware only in the general sense that companies collect and use their data to display ads, and are almost never privy to the specifics. For example, fewer than 20% of consumers realize they share their communication history, IP addresses, and web-surfing history when using a standard web browser (Morey, Forbath, and Schoop 2015), erroneously believing that only explicitly divulged information is harvested. Furthermore, despite the fact that companies

routinely sell and package consumer data in aggregated formats, most consumers do not understand how their personal information is consolidated across platforms (Kroft 2013).

Therefore, while improved targeting techniques have been advantageous for firms, the result may be more mixed for consumers (White 2004). On the one hand, well-targeted ads are objectively more personalized; thus, they should by definition be more relevant and interesting to consumers. They can also facilitate the potential discovery of new products uniquely suited to consumers' needs, wants, and interests (Charts 2014; Tam and Ho 2006). Transparency, in turn, could make such targeting practices more salient and further increase consumers' awareness of the degree to which ads are tailored. Indeed, informing consumers that an ad has been targeted based on their behavior can increase the perceived person-product fit, increasing ad effectiveness (Summers, Smith, and Reczek 2016). However, within the context of consumers' general lack of awareness about how their personal data are used, ad transparency could also make salient something that is not typically top of mind: the fact that marketers are collecting and using their personal information, which could raise privacy concerns and potentially decrease ad effectiveness.

We reconcile these competing predictions by developing a framework that, at its core, invokes consumer beliefs about which personal information flows—that is, ways in which marketers collect and use their personal information—are acceptable versus unacceptable. We begin by developing a conceptual account to predict which personal information flows consumers are likely to find acceptable versus unacceptable. In the empirical portion of the article, we first assess whether these predictions emerge in the data; we then proceed to test how disclosure of these information flows (i.e., ad transparency) affects ad effectiveness. In so doing, we help to account for when and why awareness of behavioral targeting enhances versus detracts from ad effectiveness.

A Theory of Offline-to-Online Norm Transference

In the offline world, acceptable versus unacceptable flows of information are generally well established and agreed upon (Grice, Cole, and Morgan 1975; Nissenbaum 2011). A wealth of research shows that people inherently understand, and generally adhere to, implicit rules of communication (Fishbein 1979; Grice et al. 1975). For instance, people are adept at complying to social rules of disclosure (Altman and Taylor 1973; Caltabiano and Smithson 1983; Derlega and Chaikin 1977; Huang et al. 2017) and penalize noncompliance (Archer and Berg 1978; Rubin 1975). And, with respect to the sharing of personal

information in particular, there are norms or “rules” that span across contexts.

For instance, consider the following two examples of information flow norms in the offline world. Example 1: an individual may find it acceptable to share some personal information (e.g., that she is trying to lose weight) with two different friends, but would likely find it unacceptable for one friend to take it upon himself to tell the other. In both cases, the two friends ultimately know the same information; however, the flow pattern of this information is acceptable in the first case but decidedly not in the second case. Example 2: an individual may directly inform a friend about some personal information (e.g., weight loss efforts), but would likely find it unacceptable for that friend to make an overt inference (e.g., speculating about said weight loss efforts), even if the inference is accurate. In both cases the friend ultimately has the same knowledge; however, the first flow pattern (information directly divulged by the source to the friend) is acceptable, while the other (an overt inference—i.e., the lack of a direct link from source to the friend) is not. In short, in offline contexts, there are strong norms with respect to appropriate flows of personal information, and violations of these norms can feel invasive (Nissenbaum 2004, 2011).

Although offline and online worlds differ greatly in the way personal information is gathered and used, we posit that when determining whether a given ad practice is acceptable, consumers turn to the offline world as a guidepost, where these norms—the ways in which individuals should behave in sharing personal information—are well established and well understood. Indeed, philosopher Helen Nissenbaum suggests that in orienting to the often bewildering online world, consumers look for “the contours of familiar social activities and structures” (Nissenbaum 2011). Consistent with this theorizing, Moon (2000) empirically demonstrated that the tendency to reciprocate others' disclosures—a well-established norm in offline contexts—also applies in online human-computer interactions.

These ideas, coupled with the fact that consumers have idiosyncratic, preconceived notions of how firms ought to handle their personal data (Malhotra, Kim, and Agarwal 2004; Milne and Gordon 1993; Smith, Milberg, and Burke 1996), suggest that consumer response to ad transparency will depend on consumers' perceptions of which information flow patterns should—or should not—be used to generate the ads. Specifically, we argue that consumer response to transparency messages will manifest in how they choose to engage with the targeted ad: all else equal, consumers are more likely to engage with an ad when it is accompanied by an acceptable, norm-adherent transparency message (i.e., conveying that the consumer data behind it were obtained via an acceptable information flow) than an ad with an unacceptable, norm-violating message. It follows, then, that the impact of ad transparency on ad

effectiveness depends on the perceived acceptability of the information flow it exposes. Specifically, we predict:

H1a: Ad transparency that reveals unacceptable information flows will reduce ad effectiveness relative to ad transparency that reveals acceptable flows.

How will ad transparency fare relative to no ad transparency at all? Relative to norm-adherent behavior, norm-violating behavior—and negative events more generally—tend to trigger particularly strong reactions (Baumeister et al. 2001; Fehr and Fischbacher 2004; Fehr and Gächter 2002; Morewedge 2009; Rozin and Royzman 2001). Therefore, we hypothesize that revealing unacceptable information flows (i.e., norm-violating activities) will have a particularly strong effect on consumer behavior: it will reduce ad effectiveness relative to no transparency. Following this logic that “bad is stronger than good” (Baumeister et al. 2001), revealing acceptable information flows (i.e., norm-adherent activities) may not, on its own, increase effectiveness relative to no transparency. In other words, people are likely to punish bad behavior, but may not equivalently reward good behavior; therefore, the simple act of disclosing acceptable information may not necessarily lead to a positive, non-status quo response in consumers. Indeed, as we later delineate, additional factors might be necessary to unlock the capacity for norm-adherent ad transparency to boost ad effectiveness relative to no transparency. We predict:

H1b: Ad transparency that reveals unacceptable information flows will reduce ad effectiveness relative to no ad transparency.

However, one natural question arises: Which information flows do consumers perceive to be acceptable versus unacceptable? The answer to this question, especially as it pertains to behaviorally targeted ads, is not yet known. Therefore, as a prerequisite to testing hypotheses 1a and 1b, we use an inductive approach to identify what norms of information flows—as they pertain to behaviorally targeted ads—consumers believe companies ought to follow. Guided by our theory of offline-to-online norm transference, we predict that these domains will mirror offline injunctive norms and be predictive of the effect of ad transparency on ad effectiveness. For instance, as in the aforementioned example 1, consumers may expect companies to abstain from sharing their personal information with other companies; and as in example 2, consumers may expect companies to use only what consumers have stated about themselves and to refrain from making inferences about them.

Privacy versus Personalization

After we determine what constitutes acceptable versus unacceptable information flows, a related second question arises: What underlies the relationship between the acceptability of the information flows revealed by ad

transparency and the resultant effectiveness of the ad? From the firm’s perspective, producing effective targeted advertisements entails the careful balance of consumer interest in ad personalization (i.e., how much the ad fits one’s needs and interests) and consumer information privacy concerns (i.e., concern over the safety and control over one’s personal information; Malhotra et al. 2004; Milne and Gordon 1993; Phelps, Nowak, and Ferrell 2000; Smith et al. 1996). Previous research suggests that consumers assess the attractiveness of an advertised product by looking to its fit with, or personalization to, their needs and wants (Aguirre et al. 2015; Montgomery and Smith 2009; Tam and Ho 2006). Targeted ads—which should represent a better fit with and personalization to consumer needs and wants—may drive higher ad effectiveness.

But marketers can go too far; personalization can backfire. Of particular relevance to the present inquiry, highly personalized ads (compared to less personalized ads) can rouse consumer privacy concerns (Aguirre et al. 2015). Such concerns can overshadow interest in purchasing the product itself, and more immediately, dampen the likelihood of engaging with advertisements for it—even for a product that personalized targeting would predict to be a good person-product fit (White et al. 2008). Similarly, we propose that ad transparency affects ad effectiveness by shifting consumers’ relative concerns for privacy versus their interest in personalization. Specifically, a message that communicates a norm-violating practice, thus making salient that “acceptable” information flows have been breached, will activate privacy concerns that ultimately eclipse the benefits of personalization; even the most personalized, perfectly targeted advertisement will flop if the consumer is more focused on the (un)acceptability of how the targeting was done in the first place. Such privacy-related backlash has been documented in other domains as well. Privacy concerns prompt behavioral inhibition, akin to how a prevention focus triggers defensive behavior (Higgins 1997), making consumers less willing to make purchases (Tsai et al. 2011), open commercial emails (White et al. 2008), and divulge personal information (Culnan 2000; John, Acquisti, and Loewenstein 2011; Phelps et al. 2000; Singer, Hippler, and Schwarz 1992). In the realm of behavioral advertising, we suggest that this inhibition manifests as ad avoidance: refraining from engaging with the ad or the product it promotes.

H2: The reduction in ad effectiveness resulting from revealing unacceptable information flows will be driven by consumers’ relative concern for privacy over their desire for the personalization that such targeting affords.

Unlocking the Upside of Ad Transparency

Our hypotheses thus far have focused on the downside of ad transparency. Specifically, we predict that

unacceptable ad transparency backfires (by reducing ad effectiveness), but that acceptable ad transparency may not in and of itself increase ad effectiveness relative to no transparency. However, we also wanted to investigate the following question: When might ad transparency increase ad effectiveness? We suggest that trust in the platform on which the ad is displayed (e.g., Facebook) moderates the effect of ad transparency on ad effectiveness (our predictions so far pertain to the contexts in which consumers' feelings about the ad platform are relatively neutral—perceived by consumers as neither particularly trustworthy nor untrustworthy). When acceptable information flows are revealed on trusted platforms, we predict that ad transparency may increase ad effectiveness. This prediction is predicated on the idea that the trustworthiness of the ad platform has spillover effects on consumers' reactions to the advertisements it displays. In support of this claim, consumers' assessments of trustworthiness in the digital space tend to be generalized (Pavlou, Liang, and Xue 2006; Stewart 2003); if a user trusts the Facebook brand, her trust applies more expansively to the Facebook website and its features, including ads and the transparency messages that accompany them. Therefore, we expect ad transparency to have a different impact on ad effectiveness as a function of consumers' individual levels of trust in the ad platform.

High consumer trust is associated with feelings of safety and reliability, and generates a willingness to engage (Chaudhuri and Holbrook 2001; Doney and Cannon 1997). Thus, when consumers trust an ad platform and norm-adhering ad transparency practices are displayed, their response—and, in turn, the effect on ad effectiveness—will be most favorable. However, when consumers trust an ad platform and norm-violating practices are displayed, their trust is violated; in this case, consistent with our previous predictions, reactions to norm-violating transparency messages—and the subsequent effect on ad effectiveness—will be negative.

Low consumer trust is associated with skepticism and a belief that firms are out to maximize their own self-interest (Bleier and Eisenbeiss 2015; Wang and Benbasat 2005). In these cases, we suggest that ad transparency—of any kind—is unlikely to overcome low trust; even acceptable targeting practices will be met with the skepticism accompanying distrust. Consistent with this idea, previous work has shown that suspicion of ulterior motives leads consumers to discount firms' benevolent acts, such as charitable giving (Foreh and Grier 2003; Lin-Healy and Small 2013; Newman and Cain 2014; Yoon, Gurhan-Canli, and Schwarz 2006). See figure 1.

In sum, with respect to the relationship between trust, ad transparency, and ad effectiveness, we predict:

H3a: When consumers trust the ad platform, ad transparency that adheres to information flow norms will increase ad effectiveness.

FIGURE 1

UNLOCKING THE UPSIDE OF TRANSPARENCY

Consumer's trust in platform	High	+ Boosts ad effectiveness	– Betrayed trust
	Low	– Met with skepticism	– Met with skepticism
		Acceptable	Unacceptable
		Information flow revealed	

H3b: When consumers distrust the ad platform, ad transparency will decrease ad effectiveness regardless of whether it adheres to information flow norms.

OVERVIEW OF STUDIES

As a prerequisite to testing hypotheses 1a and 1b, we first generated a list of ways firms obtain and use consumers' personal information in generating targeted ads, and then measured consumer acceptance of these practices (study 1). A factor analysis identified two key dimensions of consumer acceptability of information flows. Specifically, acceptability varies according to whether the personal information used to generate the ad was 1) obtained within versus outside of the website on which the ad appears; and 2) explicitly stated by the consumer versus inferred by the firm. Next, we conducted confirmatory experiments. Studies 2 and 3, lab experiments that simulate behavioral targeting, tested whether the dimensions identified in study 1 indeed predict ad effectiveness (hypotheses 1a and 1b). Studies 2 and 3 additionally tested the mediating role of consumers' relative concern for privacy over their interest in personalization (hypothesis 2). The final set of studies speaks to the upside of ad transparency (hypotheses 3a and 3b). Specifically, study 4 shows that the effect of ad transparency on ad effectiveness is moderated by trust: disclosures of consumer-accepted information flows will be particularly effective when consumers trust the ad platform. Studies 5a and 5b, conducted within a context in which trust is present (a loyalty program website), demonstrate the upside of transparency in the field.

Guided by marketing practice and previous research, we captured our dependent measure—ad effectiveness—in

different, but converging, ways across experiments: self-reported purchase interest (Summers, Smith, and Reczek 2016) in studies 2–4 and clicking on the ad (Bleier and Eisenbeiss 2015; Tucker 2014) in studies 5a and 5b.

STUDY 1: IDENTIFYING DIMENSIONS

Study 1 was an inductive study. We compiled a list of practices that firms use to generate behaviorally targeted ads and measured how acceptable consumers found them to be.

Procedure

Item Generation. We compiled a list of the ways in which the two largest online advertising platforms, Google and Facebook, allow advertisers to employ users' personal data to generate ads. We supplemented this list by asking a group of Facebook users ($N = 79$, 54% male; $M_{\text{age}} = 22.9$, $SD = 4.15$) to cull the ad transparency text provided in the first ad displayed in their Facebook newsfeed (see appendix B for exact procedure). Finally, we added further practices of which we were aware, based on industry experiences. This effort resulted in 30 items (appendix C).

Item Acceptability. Participants ($N = 149$, 52.3% male; $M_{\text{age}} = 35.13$, $SD = 10.14$) were told that "Facebook generates personalized advertisements for their users using various methods" and were presented with the 30 ad practices in random order. For each item, participants rated the extent to which they agreed with the statement: "Facebook should show me advertisements based on [practice]" (1 = "Strongly disagree"; 7 = "Completely agree"). In this and all other studies, we report all independent and dependent variables and data exclusions. We recruited participants by following the minimum threshold of 100 participants per cell (though we used smaller prespecified thresholds for pretests). We did not analyze data until we finished collecting the prespecified number of participants. Stimuli and data (with the exception of the field data collected in studies 5a and 5b) are available at osf.io/4fjqr.

Results

To prepare for factor analyses, we began by performing the Kaiser–Meyer–Olkin test of sampling adequacy and Bartlett's test of sphericity on the 30 items. The Kaiser–Meyer–Olkin measure of sampling adequacy was within an acceptable range (total matrix sampling adequacy = .90), and Bartlett's test of sphericity was significant ($p < .001$), suggesting that our data were appropriate for factor analyses. We then conducted a principal axis factor analysis with Varimax rotation, suppressing coefficients whose absolute values were below .5. This analysis suggested five factors, accounting for 74.3% of the total variance. One of these factors included six items that did not reveal a clear pattern,

and one item did not fit any of the factors. Following the approach used in past research (Aaker 1997), we excluded these items and conducted a principal axis factor analysis again. This revealed a four-factor structure accounting for 75.2% of the total variance. The loadings of all items on each factor were .5 or higher (absolute value).

Cronbach's alphas for each factor were .86 or higher (see appendix D for factor loadings). We named the four factors: information obtained within-website, information obtained cross-website, attributes stated by the consumer, and attributes inferred by the firm. The four factors thus represent the four ends of two dimensions—whether the personal information used to generate the ad was 1) obtained within versus outside of the website on which the ad appears; and 2) explicitly provided by the consumer versus inferred by the firm.

Participants thought that firms should not use information obtained cross-website ($M_{\text{cross}} = 3.05$, $SD = 1.80$) relative to that obtained within-website ($M_{\text{within}} = 4.58$, $SD = 1.52$; $t(148) = -10.61$, $p < .001$). Participants also thought that firms should not make inferences about their attributes ($M_{\text{inferred}} = 3.11$, $SD = 1.70$) relative to relying on ones that the consumer had stated ($M_{\text{stated}} = 4.08$, $SD = 1.88$; $t(148) = -7.30$, $p < .001$).

Discussion

In sum, study 1 revealed two dimensions that predict most of the variance in how consumers perceive information flow norms; specifically, perceived acceptability varies depending on whether the personal information used to generate the ad was 1) obtained within versus outside of the website on which the ad appears; and 2) explicitly stated by the consumer versus inferred by the firm. These dimensions align with offline norms of information sharing. Indeed, the first dimension recollects the information flow norms described in example 1: generating an ad based on consumer information obtained from a different website is akin to talking behind someone's back. Similarly, the second dimension recollects the information flow norms described in example 2: making an overt inference about someone can be taboo. Furthermore, these inductive insights are broadly consistent with Smith, Milberg and Burke (1996), who identified unauthorized secondary use of personal information as one of the key drivers of privacy concerns in organizations.

STUDY 2: REVEALING WITHIN- VERSUS CROSS-WEBSITE INFORMATION FLOWS

Study 2 tested the effectiveness of revealing within- versus cross-website information flows. We predicted that revealing cross-website information flows would decrease ad effectiveness relative to the within-website flows and also relative to no transparency (hypotheses 1a and 1b). Study 2 also assessed whether the effect of ad transparency on ad

effectiveness is mediated by the differential activation of consumers' concerns for privacy over their interest in personalization (hypothesis 2). Participants were shown an ad for a bookstore. Between-subjects, participants were either given no information on why they were seeing the ad (baseline), or one of two transparency messages (i.e., that the ad had been generated based on information obtained either within-platform or cross-platform). To increase the credibility of this manipulation and to simulate real behavioral targeting, we had all participants first engage in an online movie browsing task.

Procedure

Participants ($N = 449$, 48% male; $M_{\text{age}} = 35.9$, $SD = 11.7$) were recruited from Amazon Mechanical Turk (MTurk) and assigned to one of three conditions: a baseline condition or one of two transparency conditions (within-website or cross-website). Participants were informed that the study consisted of two parts.

Part 1. Participants were informed that they would first be asked to browse movies. Between-subjects, we manipulated whether they browsed movies within the website on which they would subsequently receive an advertisement, or on a different website. Within-website participants encountered everything in a single browser window, while cross-website participants encountered two distinct browser windows. Specifically, participants in the within-website [cross-website] condition were told:

You will be browsing movies within [outside of] this survey platform. Please click on any movies that catch your eye to learn more about them, and feel free to check out as many movies as you would like. We will give you a minute to do this task, but you can spend more time if you'd like. When you are done browsing, please proceed to the next page [When you are done browsing, please return back to this page and proceed to the next page].

In the within-website condition, participants saw a list of 10 movies on the same page as the above instructions (e.g., *Logan*, *La La Land*; see appendix E). When they clicked a title, they saw a short description for the movie. In the cross-website condition, participants clicked on a link that took them to an external site. Specifically, upon the participant clicking the link, a new browser page opened, displaying the same web page as that for the within-website condition (the only difference being that for the cross-website participants, this web page was clearly external to the survey platform as opposed to embedded within the survey; appendix E).

To keep participants' experiences consistent, we had participants in the baseline condition also engage in this movie browsing task, with half randomized to browse movies within-website, and half randomized to browse movies cross-website. In our results, we collapse across

this control factor. In addition, to ensure participant engagement in the browsing task (i.e., part 1), we told all participants up front that they may be asked to answer questions about the browsing experience.

Part 2. Participants were informed that part 2 entailed evaluating an ad. Specifically, they read, "We will be showing you an ad that is targeted for you. Targeted advertising is a type of advertising that allows companies to reach consumers based on various traits." Participants waited for 5 seconds, during which we were ostensibly generating an ad for them, and then saw an ad about a fictitious online bookstore, UBooks.com. The transparency conditions included a message next to the ad. Those in the within-platform condition read, "You are seeing this ad based on the products you clicked on while browsing our website (i.e., within this survey platform)" and those in the cross-website condition read, "You are seeing this ad based on the products you clicked on while browsing a third-party website (i.e., outside of this survey platform)." In the baseline condition, there was no message (see appendix F for stimuli). The language used for the two transparency conditions in this and subsequent studies was adapted from the actual text that firms use when listing their ad practices (see appendix G for an example).

Ad Effectiveness. Participants indicated the extent to which they agreed with the statements: "I am interested in visiting the website for UBooks.com" and "I am interested in buying products from UBooks.com" on a seven-point scale (1 = "Strongly disagree" to 7 = "Strongly agree"). We averaged these items to create a composite score of ad effectiveness ($\alpha = .92$).

Privacy over Personalization. To assess our proposed mediator, we asked participants: "In order to provide more personalized recommendations for you, marketers need to gather more information about you. In other words, when receiving an advertisement, there is a tradeoff between maintaining your privacy and enjoying the benefits of greater personalization. Upon seeing the above ad, which factor is more important to you when evaluating a targeted ad?" on a 10-point scale (1 = "Privacy is more important to me" to 10 = "Personalization is more important to me"). We reverse-coded this item such that higher scores represented higher relative concern for privacy.

Pretests

We conducted two pretests to ensure that ad practices disclosed in the within-website and cross-website conditions 1) were indeed perceived as acceptable and unacceptable, respectively; and 2) did not differ in perceived accuracy.

Pretest 1. To examine whether ad practices disclosed in the within-website and cross-website conditions were indeed perceived as acceptable and unacceptable,

respectively, we recruited a separate group of participants ($N=92$, 35.9% male; $M_{\text{age}}=37.0$, $SD=11.5$) and informed them that “Targeted advertising is a type of advertising that allows companies to reach consumers based on various traits.” Participants were randomly assigned to one of two conditions: within-website and cross-website. We asked participants in the within-website condition to suppose that they saw an ad that was “based on the products you clicked on while browsing our website (i.e., within this survey platform)”; those in the cross-website condition were asked to suppose that they saw an ad that was “based on the products you clicked on while browsing a third-party website (i.e., outside of this survey platform).” Participants rated how acceptable they found the revealed advertising practice on a seven-point scale (1 = “Not at all acceptable” to 7 = “Very acceptable”). As intended, those in the within-website condition rated their ad practice as significantly more acceptable than those in the cross-website condition ($M_{\text{within}}=4.92$, $SD=1.73$; $M_{\text{cross}}=3.07$, $SD=1.70$; $t(90)=5.16$, $p<.001$).

Pretest 2. To ensure that perceived accuracy did not differ across the within-website and cross-website conditions, we recruited participants ($N=82$, 41% male; $M_{\text{age}}=36.1$, $SD=13.0$) and gave them the same information about targeted advertising as above. Participants were randomly assigned to one of two conditions (i.e., within-website, cross-website), and saw the same scenario as above. Then, they rated how much they thought the ad would fit their needs and wants (1 = “Not at all” to 7 = “Very much”). There was no statistically significant difference between conditions ($M_{\text{within}}=4.55$, $SD=1.60$; $M_{\text{cross}}=4.53$, $SD=1.15$; $t(80)=.07$, $p=.94$).

Results

Ad Effectiveness. A one-way ANOVA revealed a significant main effect of transparency on ad effectiveness ($F(2, 446)=14.05$; $p<.001$). As predicted, ad effectiveness was reduced in the cross-website condition ($M=2.83$, $SD=1.77$) relative to both the within-website ($M=3.56$, $SD=1.73$; $t(298)=-3.67$, $p<.001$) and baseline conditions ($M=3.87$, $SD=1.72$; $t(293)=5.14$, $p<.001$). The latter two conditions did not differ ($t(301)=1.54$, $p=.13$).

Privacy over Personalization. Results of the process measure were similar ($F(2, 446)=3.12$; $p=.045$). Specifically, privacy concerns were higher in the cross-website condition ($M=7.47$, $SD=2.43$) relative to both the within-website ($M=6.75$, $SD=2.49$; $t(298)=-2.51$, $p=.01$) and baseline conditions ($M=6.97$, $SD=2.64$; $t(293)=1.67$, $p=.09$). The latter two conditions did not differ ($t(301)=-.75$, $p=.46$).

Mediation. The differential impact of the within- versus cross-website conditions on ad effectiveness was mediated by participants’ relative concern for their privacy over

their interest in personalization: a 5,000 sample bootstrap analysis using PROCESS model 4 (Hayes 2013) indicated a significant indirect effect ($b=-.17$, $SE=.08$; 95% confidence interval: $[-.23, -.02]$).

Discussion

Study 2 showed that ad transparency reduced ad effectiveness when it revealed cross-website tracking—an information flow that consumers deem unacceptable, as identified by our inductive study 1. Consistent with hypothesis 2, ad transparency that revealed unacceptable information flows heightened concern for privacy over interest in personalization, reducing ad effectiveness.

STUDY 3: REVEALING STATED VERSUS INFERRED INFORMATION FLOWS

Study 3 tested the effectiveness of revealing that an ad is based on stated versus inferred attributes. We predicted that ads based on inferred information would decrease ad effectiveness relative to stated attribute transparency and relative to no transparency (hypotheses 1a and 1b). Like study 2, study 3 also predicted privacy over personalization concerns to mediate the relationship between ad transparency and ad effectiveness. Participants were shown an ad for an online art gallery. Between-subjects, participants were given either no information on why they were seeing the ad (baseline), or one of two transparency messages (i.e., that the ad had been generated based on information either stated by the consumer or inferred by the firm). To increase the credibility of this manipulation and to simulate real behavioral targeting, we had all participants first complete an online shopper profile.

Procedure

Participants ($N=348$, 45% male; $M_{\text{age}}=37.66$, $SD=12.45$) were recruited from MTurk and shown one of three different versions of an ad varying in transparency: no transparency (baseline condition), or one of two transparent ads—one revealing it had been targeted based on stated attributes, and one revealing it had been targeted based on inferred attributes.

First, participants were told: “We are a group of researchers in marketing. Today, we will be showing you targeted advertisements. Targeted advertising is a type of advertising that allows companies to reach consumers based on various traits.” Next, participants were asked to build an online shopper profile by filling out a form. This form listed three demographic questions that were randomly chosen from the following: gender, age, highest level of education completed, current relationship status, and student status. On the next screen, participants were further informed: “In order to provide targeted online advertisements for you, marketers can rely on the

information that you've given them voluntarily (i.e., what you've stated about you), or infer other types of information about you based on things like your Internet connection and things that you click on (i.e., what they infer about you)."

After completing an unrelated task, participants were presented with one of three different versions of an ad for an online art gallery store.

Manipulation. The transparency conditions included a message next to the ad: "You are seeing this ad based on your information that you stated about you" (stated attribute condition); "You are seeing this ad based on your information that we inferred about you" (inferred attribute condition). In the baseline condition, there was no message (see appendix H for stimuli).

Ad Effectiveness. Participants indicated the extent to which they agreed with the following statements on a seven-point scale: "I am interested in visiting the website for UGallery.com" and "I am interested in buying products from UGallery.com" (1 = "Strongly disagree" to 7 = "Strongly agree"). We averaged these items to create a composite ad effectiveness score ($\alpha = .93$).

Privacy over Personalization. We measured the relative activation of privacy concerns versus interest in personalization using the same item as in study 2.

Pretests

As in study 2, we conducted two pretests to ensure that ad practices disclosed in the stated attribute and inferred attribute conditions 1) were indeed perceived as acceptable and unacceptable, respectively; and 2) did not differ in perceived accuracy. Both pretests used the same designs as in study 2's pretests, with the exception that participants read about stated or inferred information flows.

Pretest 1. To examine whether ad practices disclosed in the stated attribute and inferred attribute conditions were indeed perceived as acceptable and unacceptable, respectively, we randomly assigned participants ($N = 70$, 45.7% male; $M_{\text{age}} = 38.7$, $SD = 13.3$) to one of two conditions (i.e., stated, inferred). As intended, those in the stated condition rated their ad practice as significantly more acceptable than those in the inferred condition ($M_{\text{stated}} = 4.55$, $SD = 1.57$; $M_{\text{inferred}} = 3.54$, $SD = 1.78$; $t(68) = 2.49$, $p = .02$).

Pretest 2. To ensure that perceived accuracy did not differ across the stated and inferred conditions, we randomly assigned participants ($N = 114$, 40.4% male; $M_{\text{age}} = 34.7$, $SD = 11.6$) to one of two conditions (i.e., stated, inferred). There was no statistically significant difference across conditions ($M_{\text{stated}} = 4.86$, $SD = 1.30$; $M_{\text{inferred}} = 4.54$, $SD = 1.45$; $t(112) = 1.26$, $p = .21$).

Results

Ad Effectiveness. A one-way ANOVA revealed a significant main effect of transparency on ad effectiveness ($F(2, 345) = 4.12$; $p = .02$). As predicted, ad effectiveness was reduced in the inferred attribute condition ($M = 2.52$, $SD = 1.35$) relative to both the stated attribute ($M = 3.10$, $SD = 1.82$; $t(232) = 2.77$, $p = .01$) and baseline conditions ($M = 2.96$, $SD = 1.62$; $t(230) = 2.22$, $p = .03$). The latter two conditions did not differ ($t(228) = -.63$, $p = .53$).

Privacy over Personalization. Results were similar for privacy ($F(2, 345) = 2.83$; $p = .06$). Specifically, privacy concerns were higher in the inferred attribute condition ($M = 3.46$, $SD = 2.26$) relative to both the stated attribute ($M = 4.15$, $SD = 2.70$; $t(232) = 2.12$, $p = .04$) and baseline conditions ($M = 4.07$, $SD = 2.30$; $t(230) = 2.05$, $p = .04$). The latter two conditions did not differ ($t(228) = -.23$, $p = .82$).

Mediation. The differential impact of inferred versus stated conditions on ad effectiveness was mediated by participants' relative concern for their privacy over their interest in personalization. A 5,000 sample bootstrap analysis using PROCESS model 4 (Hayes 2013) suggested a significant indirect effect ($b = -.17$, $SE = .08$; 95% confidence interval: $[-.35, -.01]$).

Discussion

Study 3 showed that ad transparency reduced ad effectiveness when it revealed that the ad was based on consumer attributes inferred by the firm—an information flow that consumers deem unacceptable, as identified by our inductive study 1. Consistent with hypothesis 2, ad transparency that revealed unacceptable information flows increased concern for privacy over interest in personalization, thus reducing ad effectiveness.

Finally, note that in studies 2 and 3, ads that revealed acceptable information flows performed just as well as those in the baseline condition. To examine whether and when ad transparency could boost ad effectiveness, above and beyond the baseline control, we conducted the next set of studies on ad platforms for which consumers have preestablished (dis)trust. Thus, these studies test whether trust can unlock the benefits of disclosing (vs. not disclosing) acceptable information flows.

STUDY 4: THE MODERATING ROLE OF TRUST

Study 4 tested hypotheses 3a and 3b by examining the moderating role of individual differences in ad platform trust. To do so, we measured consumer trust in Facebook, and examined whether the level of trust in the ad platform (i.e., trust in Facebook) interacted with ad transparency (i.e., the degree to which an ad practice adhered to norms of

information flows). In addition, we again measured our mediator in study 4, thus providing an additional test of hypothesis 2.

Procedure

Participants ($N = 462$, 52.5% male; $M_{\text{age}} = 33.5$, $SD = 10.2$) who indicated that they had a Facebook account were recruited from MTurk. We first assessed how much participants trusted Facebook by asking them to rate the extent to which they agreed or disagreed with the statements (Chaudhuri and Holbrook 2001): “I trust Facebook,” “I rely on Facebook,” “Facebook is an honest brand,” and “Facebook is safe,” each on a seven-point scale (1 = “Strongly disagree” to 7 = “Strongly agree”). We averaged these items into a composite trust measure ($\alpha = .88$).

Next, all participants were asked to log into their Facebook accounts. To ensure that they were logged in, we asked a question that required them to be logged in to know its answer (appendix I), and participants could not advance in the survey until they correctly answered this question. Next, we guided participants to locate the first ad in their newsfeed and to click on its “Why am I seeing this ad” message (appendix B). Participants were asked to copy and paste the message and the company name. Forty-eight participants failed to accurately complete this step; we excluded them from analysis (however, the results hold when they are included).

Next, participants coded the content of their transparency message with respect to our independent variable of interest: the degree to which the ad practice adhered to norms of information flows. Specifically, participants indicated “Yes,” “No,” or “I’m not sure” to each of the following four statements: “This ad was presented to me based on: 1) my activities within Facebook, 2) my activities on third-party websites (outside of Facebook), 3) the information that I stated about myself on Facebook, and 4) the information that Facebook inferred (i.e., guessed) about me.” This coding generated a transparency score for each ad, with higher transparency scores reflecting more acceptable practices. Specifically, we assigned a value of +1 to “Yes” responses to the first and third statements, which depict acceptable information flows (i.e., within-website tracking and stated attributes, respectively), and a value of -1 to “Yes” responses to the second and fourth statements, which depict unacceptable information flows (i.e., cross-website tracking and inferred attributes, respectively). We assigned a neutral value of 0 to all “No” and “I’m not sure” responses (see appendix J for descriptive statistics). Following an approach in psychology literature, we then summed these values to create a transparency score for each observation (Costa and McCrae 1980).

Finally, participants answered a series of questions based on the ad they had received. First, they rated how much they were interested in: 1) visiting the company website, 2)

spending their money at the company, 3) clicking the like button for the company’s Facebook page, and 4) spending their money on the advertised product on a seven-point scale (1 = “Strongly disagree” to 7 = “Strongly agree”). We averaged these items into a composite ad effectiveness measure ($\alpha = .93$). As in studies 2 and 3, participants indicated how much they were concerned about privacy over personalization. We also administered a trait measure of privacy concern ($\alpha = .85$; Malhotra, Kim, and Agarwal 2004), which did not interact with our independent variables.

Results

Ad Transparency and Ad Effectiveness. A linear regression revealed that transparency scores significantly predicted ad effectiveness ($b = .31$, $SE = .09$, $p < .001$). In other words, participants who received an ad that revealed acceptable practices were more likely to engage with the ad relative to those receiving an ad that revealed unacceptable practices.

Privacy over Personalization. A linear regression also revealed that transparency scores significantly predicted privacy concerns ($b = .25$, $SE = .11$, $p = .03$). That is, participants who received an ad that revealed acceptable practices were less likely to be concerned about privacy over personalization relative to those receiving an ad that revealed unacceptable practices.

Mediation. The impact of ad transparency on ad effectiveness was mediated by participants’ relative concern for their privacy over their interest in personalization. Consistent with previous studies, a 5,000 sample bootstrap analysis using PROCESS model 4 (Hayes 2013) suggested a significant indirect effect ($b = -.17$, $SE = .08$; 95% confidence interval: $[-.23, -.02]$).


Moderation. Next, we examined whether the relationship between transparency scores and ad effectiveness was moderated by platform trust. A linear regression revealed a significant interaction between trust and transparency score ($b = .15$, $SE = .06$, $p = .02$).

To decompose this interaction, we followed procedures recommended by Spiller et al. (2013) and used the Johnson-Neyman technique to identify the range(s) of trust for which the simple effect of revealing acceptable information flows was significant. This analysis revealed a significant effect of ads with acceptable transparency on ad effectiveness for any model with trust greater than 3.56 ($BJN = .18$, $SE = .09$, $p = .05$), but not for any model with trust less than 3.56. In other words, users who trust Facebook were more likely to engage with an ad that revealed acceptable information flows than those who distrust Facebook. See figure 2.

We also conducted a moderated mediation analysis to simultaneously test moderation by trust and mediation by privacy concern. A 5,000 sample bootstrap analysis showed that the index of moderated mediation excluded

FIGURE 2

AD TRANSPARENCY AND AD PLATFORM TRUST



zero (95% CI = [.002, .04]), suggesting a significant indirect effect (Hayes 2013).

STUDIES 5A AND 5B: THE BENEFIT OF AD TRANSPARENCY IN THE FIELD

Study 4 suggested that the benefit of revealing (vs. not revealing) acceptable information flows emerges in contexts in which trust is present. Our final studies tested the effect of ad transparency on ad effectiveness under favorable circumstances: when consumers both trust the platform and are shown norm-adherent transparency messages.

To do this, we conducted two different field experiments with a trusted platform. Previous research has shown that loyalty programs are a context in which consumer trust is high; members typically have higher trust than nonmembers (Ashley et al. 2011; Chaudhuri and Holbrook 2001; García Gómez, Gutiérrez Arranz, and Gutiérrez Cillán 2006). Thus, studies 5a and 5b were conducted within the context of two different loyalty program point-redemption sites, and sought to replicate the beneficial effect of ad transparency in the field. Specifically, these studies sought evidence of a boost for ad transparency when it was disclosed that personal information was obtained within the point-redemption site (i.e., within-website, study 5a) or explicitly by the user (i.e., stated attributes, study 5b).

Field Setting Description

We collaborated with a provider of a white-label online platform to companies that operate rewards programs. For example, if a consumer is enrolled in a hotel membership program and accumulates points, he/she can use those points to purchase various items on the hotel's rewards

redemption site, serviced by the provider with whom we partnered.

The provider therefore customizes the platform to the given loyalty program and features that program's own branding. Studies 5a and 5b involved two companies that use this reward site, one per experiment.

Study 5A Procedure

Users of a given company's rewards redemption site were randomly assigned to one of two conditions: baseline (no transparency) versus within-website. The experiment was live for two weeks in the spring of 2017; the sample consisted of 9,079 unique users who logged into the rewards site during the two-week period.

As a user browses products on this site, a sidebar is displayed, suggesting items for the user to buy (appendix K). These items are a form of behaviorally targeted advertising; they are tailored based on the given user's behavior. Our experiment leveraged this sidebar feature by manipulating its title: in the within-website condition, the sidebar was titled "Recommended based on your clicks on our site"; in the control condition, it was simply titled "Recommended."

In both studies 5a and 5b, the primary outcome was the behavioral measure most proximal to the manipulation: the propensity to click on the recommended items. As exploratory measures, the rewards website also provided data on additional outcomes: the number of seconds spent on the pages of recommended items, and revenue generated from the recommended items.

Pretests

We conducted two pretests: one to confirm that trust in a given brand is generally higher among members of the

brand's loyalty program relative to nonmembers (consistent with Ashley et al. 2011; Chaudhuri and Holbrook 2001; García Gómez et al. 2006), and the other to assess whether the sidebar was in fact perceived to be a targeted ad.

Pretest 1. To assess whether the field setting was indeed a context in which trust is present, participants on MTurk ($N = 164$, 40.2% male; $M_{\text{age}} = 37.5$, $SD = 12.08$) were randomly assigned to one of two conditions: loyalty versus non-loyalty. Those in the loyalty [nonloyalty] condition were told: "Please name one company that meets the following requirements: 1) you have made a purchase from this company at least once over the past six months and 2) you are [are not] enrolled in this company's loyalty rewards program." Participants then rated the extent to which they trusted the company they had listed on a seven-point scale (1 = "Not at all" to 7 = "Very much"). Those in the loyalty condition ($M = 5.99$, $SD = 1.01$) trusted the company they listed significantly more than those in the nonloyalty condition ($M = 5.34$, $SD = 1.38$, $t(162) = 3.41$, $p = .001$).

Pretest 2. To assess whether the sidebar was in fact perceived to be a targeted ad, participants on MTurk ($N = 153$, 50.6% male; $M_{\text{age}} = 34.6$, $SD = 11.1$) were shown a screenshot of the point-redemption site, with the product sidebar highlighted. Participants were told: "Targeted advertising is a type of advertising that allows companies to reach consumers based on various traits." They were asked to rate the extent to which they agreed with the statement: "The highlighted portion comes across as a targeted advertisement" (1 = "Strongly disagree" to 7 = "Strongly agree"). Results indicated that the sidebar was indeed perceived as a targeted ad; the mean rating was statistically significantly higher than the scale midpoint ($M = 5.10$, $SD = 1.40$, $t(152) = 9.79$, $p < .001$).

Study 5A Results

We conducted a mixed-effects regression analysis to evaluate the effect of revealing acceptable information flows (in this case, that the consumer was targeted based on information obtained within-website) on the propensity to click on recommended items. To account for individual-level differences, we included participant ID as a subject variable. As predicted, participants in the within-website condition were more likely to click on recommended items than those in the baseline condition ($M_{\text{within}} = .10$, $SD = .001$ vs. $M_{\text{baseline}} = .09$, $SD = .001$; Wald chi-square = 100.56, $p < .001$). Participants in the within-website condition also spent more time (seconds) on the recommended products' pages ($M_{\text{within}} = 8.52$, $SD = .14$ vs. $M_{\text{baseline}} = 6.35$, $SD = .14$; $F(1, 279, 006) = 117.73$, $p < .001$) and in turn, more money on the recommended products ($M_{\text{within}} = .22$, $SD = .02$ vs. $M_{\text{baseline}} = .16$, $SD = .02$; $F(1, 279, 006) = 4.24$, $p = .04$).

Study 5B Procedure

Study 5b was run concurrently with study 5a with another company that uses the rewards site. Users were randomly assigned to one of two conditions: baseline (no transparency) and stated attribute. The sample consisted of 1,862 participants who logged into the rewards site during the two-week period.

As in study 5a, we leveraged the sidebar feature by manipulating the title of the sidebar. In the stated attribute condition, the sidebar was titled, "Recommended based on what you've shared with us"; in the control condition, it was simply titled "Recommended."

Study 5B Results

We conducted a mixed-effects regression analysis to evaluate the effect of revealing acceptable information flows (in this case, that the consumer is being targeted based on information they explicitly stated about themselves) on the propensity to click on recommended items. To account for individual-level differences, we included participant ID as a subject variable. As predicted, participants in the stated attribute condition were more likely to click on recommended items than those in the baseline condition ($M_{\text{stated}} = .12$, $SD = .003$ vs. $M_{\text{baseline}} = .08$, $SD = .002$; Wald chi-square = 76.06, $p < .001$). Participants who were exposed to the stated attribute condition also spent more time (seconds) on the recommended products' pages ($M_{\text{stated}} = 8.31$, $SD = .43$ vs. $M_{\text{baseline}} = 5.98$, $SD = .42$; $F(1, 28, 136) = 14.98$, $p < .001$). There was no impact of condition on revenue ($M_{\text{stated}} = .11$, $SD = .04$ vs. $M_{\text{baseline}} = .17$, $SD = .04$; $F(1, 28, 136) = .86$, $p = .35$).

Discussion

Studies 5a and 5b, conducted within a loyalty program point-redemption website—a field setting in which ad platform trust was presumed to be, and which our pilot data suggested was, present—suggested that revealing acceptable information flows can increase ad effectiveness. However, we note that because we did not manipulate trust, alternative explanations are possible.

GENERAL DISCUSSION

In response to growing pressure from consumers and regulators, the number of firms adopting ad transparency is on the rise. Our research explores how ad transparency affects ad effectiveness, and is based on the premise that it depends on the perceived acceptability of the revealed practice. Because this territory is largely uncharted, we began with an inductive approach to capture how acceptable (or not) consumers perceive various data collection and use practices. Study 1 revealed that whether consumers deem information flows acceptable is driven by the extent to which the ad is

based on 1) consumers' activity tracked within versus outside of the website on which the ad appears, and 2) attributes explicitly stated by the consumer versus inferred by the firm (the latter of each pair is deemed less acceptable). Next, we conducted confirmatory experiments. Studies 2 and 3 confirmed that the two dimensions identified in study 1 indeed predict consumers' response to ad transparency. These studies also documented that consumers' relative concern for privacy over their interest in personalization mediates the relationship between the acceptability of the revealed practice and the effectiveness of the ad. Finally, we demonstrated the moderating role of platform trust: companies can reap the benefits of ad transparency that reveals acceptable information flows when consumers trust the ad platform.

We contribute to the literature on digital advertising. Existing research has focused on the effectiveness of online ads when consumers are unaware of underlying ad practices; however, as recent articles suggest, consumers are becoming increasingly conscious of and concerned about how marketers may be using their information (Lambrecht and Tucker 2013; Morey et al. 2015). Recent data breaches involving major companies—from Walmart to Ashley Madison to Equifax—have exacerbated this concern. Thus, it is critical to understand how consumers' willingness to engage with online ads is affected by their awareness of the data practices used to deliver such ads. In addition, much of the literature on behaviorally targeted ads has been correlational; our work adds to the small but growing body of experimental work on the topic (Bleier and Eisenbeiss 2015; Schumann, von Wangenheim, and Groene 2014; Summers, Smith, and Reczek 2016).

We also contribute to the literature on privacy. Privacy concerns have been shown to play a role in consumers' willingness to divulge information (Phelps et al. 2000; John et al., 2011; Hofstetter, Rueppel, and John 2017), but their role in targeted advertising is underexplored. Particularly as advertising becomes increasingly specific and well targeted, making salient to consumers that they are being tracked (e.g., seeing an ad featuring a product you just browsed; encountering an ad based on sensitive or private web searches you have done; Paul 2017), privacy concerns may likewise increase, with possible downstream consequences for disclosure, search behavior, and consumer-firm relationships. On the other hand, given research suggestive of consumers' capacity to adapt to privacy violations (Acquisti, John, and Loewenstein 2012), and the stickiness of those violations, the opposite may occur, with consumers becoming desensitized over time. As we describe below, future research could explore the complex and ever-evolving relationships between transparency, privacy concerns, and marketing outcomes.

Key Limitations

The following limitations qualify the contribution of our work. First, we focused on one form of ad transparency—the

type that conspicuously reveals information to consumers. Doing so enabled us to test the relationship between ad transparency and ad effectiveness in an internally valid way (i.e., independent of consumer demand for transparency) and, as a result, may help marketers predict the impact of answering consumers' and regulators' calls for greater transparency. However, these advantages come with at least one distinct disadvantage: reduced external validity, as today's standard practice requires consumers to proactively seek out such information (e.g., Facebook users have to click on the "Why am I seeing this ad?" button to view ad generation practices). As we discuss below, our hope is that our findings serve as a stepping-stone for researchers to further investigate this topic, including the type of ad transparency that more closely mirrors today's practices (i.e., the type that relies on consumers to seek out information) and drivers of consumer demand for it.

Second, because we did not directly manipulate trust in studies 5a and 5b, there may have been other factors that led us to observe a positive relationship between ad transparency and ad effectiveness. Because alternative explanations are possible, future research could establish this relationship more definitely by manipulating trust—a method complementary to the approach taken in study 4, in which we measured individual differences in platform trust to test its moderating role.

The goal of this article was intentionally broad and high-level: to examine consumer perceptions of the ways firms collect and use personal data, and to identify overarching, widely applicable dimensions that characterize these perceptions. Thus, many questions remain and the topic is ripe for further investigation. Next, we outline what we see as some of the highest-priority areas for future research on ad transparency.

Moderators of Effectiveness

Consumer response to ad transparency may depend on a variety of yet-untested variables. For example, it may vary as a function of the ad content (e.g., quality, accuracy, or specificity of the ad or product displayed), basis for targeting (e.g., the sensitivity or confidentiality of the information on which consumer was targeted), language of transparency disclosure (e.g., perceived as trying to inform or trying to persuade; low or high elaboration in describing ad practices), and impetus for the disclosure transparency disclosure (e.g., voluntary or legally mandated). With respect to the latter, when the disclosure is voluntary, does it matter whether the decision to be transparent is—or is perceived to be—made by the advertiser versus the ad platform? Relatedly, when we explored the role of trust (studies 4, 5a, and 5b), we looked at trust in the platform, as opposed to trust in the advertised brand (and used "neutral" brands that were neither particularly trustworthy nor untrustworthy). Future research could examine when and why the source of the transparency disclosure matters,

and how this might interact with trust. For instance, if the platform is trusted but the brand in the advertisement is not, how would consumers react to ad transparency?

Future research could also examine ways in which firms could counteract the negative impact of unacceptable transparency (without having to change the ways they generate targeted ads)—for example, by enhancing consumers' perceived control over their personal information. However, the notion that perceived, and not necessarily meaningful, control might affect ad transparency effectiveness points to the potential for consumers to be misled. More broadly and as we discuss below, more research is needed to understand the consumer welfare implications of firms' use of consumers' personal information.

Mediators of Effectiveness

Our research suggests that the effectiveness of targeted advertising was driven in part by consumers' relative concern for their privacy over their interest in personalization. However, the relationship between ad transparency and ad effectiveness is likely to be multiply determined and may differ by context. Furthermore, there may be instances in which one's desire for privacy and for personalization are not in conflict. Thus, we leave open the possibility that other factors may also mediate the path from transparency to effectiveness. For example, ad transparency may affect the perceived credibility of targeted ads: when told that an ad was based on certain types of information, consumers may be more convinced of the product-consumer fit relative to other types of ad transparency, which could in turn increase consumer interest. On the other hand, learning about certain types of ad practices may heighten the degree to which consumers experience identity threat (Branscombe, Schmitt, and Harvey 1999). For instance, consider a consumer who learns that she received an ad for Planned Parenthood because of her particular demographic profile. In this case, ad transparency could make her feel reduced to a single membership category, which could in turn decrease her interest in the ad. Relatedly, consumers may resist even well-targeted ads when they threaten the expression of valued identities (Bhattacharjee, Berger, and Menon 2014).

Additional Implications for the Firm

Ad transparency may affect a range of consumer perceptions and behaviors, beyond the propensity to click on a given ad. Transparency, writ large, has been shown to have positive effects for the firm (Buell, Kim, and Tsay 2017; Buell and Norton 2011); ad transparency could have positive downstream effects. For example, given that disclosure breeds trust (John, Barasz, and Norton 2016), and that trust plays a role in consumers' acceptance of targeted ads (Aguirre et al. 2015; Bleier and Eisenbeiss 2015), it is worth investigating how ad transparency might affect consumers'

holistic view of the firm. Relatedly, there may be a positive feedback loop between privacy concerns and trust perceptions, such that enhancing trust may reduce privacy concerns, in turn shifting consumers toward valuing the increased personalization afforded by behaviorally targeted advertising. Another open question is whether encountering ad transparency encourages consumers to seek out more information from companies. In addition, as consumers become more sophisticated in their knowledge of targeted advertising, future research could investigate other secondary consequences of ad transparency for each of the parties involved.

Implications for Consumer Welfare

Our results also contribute to the debate surrounding the regulation of online privacy (Goldfarb and Tucker 2011). Proponents of industry self-regulation have been supportive of voluntary firm implementation of ad transparency (Alliance 2014). Indeed, the YourAdChoices icon is a collaborative, self-regulatory initiative led by the Digital Advertising Alliance (2014) and backed by several leading marketing and advertising associations. These entities tout the benefits of ad transparency in improving consumer welfare (Dienel 2015), with the presumption that arming consumers with more information will yield better (or at least, better-informed) decisions. However, it is conceivable that transparency could also have unforeseen adverse effects for consumers. For instance, the information-based approach to consumer empowerment stands in contrast to emerging evidence that consumers' privacy concerns can be affected by non-normative factors, making them prone to disclosing information in precisely the contexts in which it may be relatively dangerous to do so (John et al. 2011). Furthermore, if transparency (rightly or wrongly) increases consumer trust and complacency—such that consumers become less diligent about protecting their privacy or inadvertently divulging more than they otherwise would—the net benefits of transparency policies could be mitigated. Future research should examine the effects of ad transparency on consumer welfare.

Ad Transparency That Is Conspicuous versus Merely Available

This article investigated conspicuous ad transparency: after advertisers or ad platforms directly revealed transparency information to consumers, how did consumers respond? Future research might also investigate the antecedents of consumer demand for ad transparency. If left to their own devices, how likely are consumers to seek out this information when it is made available to them? This propensity could depend on a number of factors, including whether consumers are expecting it to be pleasant versus unpleasant information (Dana, Weber, and Kuang 2007; Loewenstein, Golman, and Hagmann 2017).

Relatedly, are certain types of consumers more likely to seek such information than others? How might such individual differences interact with the impact of ad transparency on ad effectiveness? While some industry surveys have suggested that consumers support ad transparency initiatives and believe it to be important in the online sphere (Dienel 2015), it is less clear when and why consumers demand such information. Future research could elucidate these causal mechanisms. For instance, are consumers more likely to demand transparency when they see a generic or highly personalized ad? An ad with neutral or sensitive content? An ad that they have encountered once or multiple times? And how (if at all) does consumer behavior change after the consumer seeks out (vs. being conspicuously provided with) transparency information? Future research may better establish the drivers of consumer demand for more transparency.

Indeed, this dynamic in supply versus demand for transparency between consumers and firms is a familiar one. For example, in food labeling, in response to the concern du jour of subsets of consumers (e.g., concerns about high fructose corn syrup), firms often respond with transparency (e.g., “no high fructose corn syrup” labels). In turn, consumers who did not initially seek such information may begin to pay attention to such labels and even update their preferences. Similarly, in the case of ad transparency, many consumers may not currently actively seek this information but start to demand it as transparency practices become more common. If consumer advocates succeed in requiring transparency messages from advertisers, we foresee this topic becoming even more relevant to the current discourse on consumer protection.

Tradeoff between Privacy and Personalization

Our studies demonstrated that consumer privacy concerns over interest in personalization mediate the relationship between ad transparency and ad effectiveness, highlighting the importance to marketers of understanding how to directly influence these attitudes, and thus influence ad effectiveness. However, the privacy versus personalization tradeoff has broader significance. Targeted advertising has the potential to enhance consumers’ online experiences: faced with the choice between viewing a website covered in entirely irrelevant ads or highly applicable and interesting ones, consumers would likely prefer the latter. But at what cost? How much personal information—be it demographic, stated preference, or behavioral—are consumers willing to divulge in exchange for better personalization? This question lies at the heart of the challenge facing modern advertisers and marketers. For firms, the benefits of personalization might ideally

render privacy concerns moot and encourage consumers to freely share their own data; however, for consumers, this reality still seems far off (Panel 2011; Rainie and Duggan 2016). Future research examining the relationship between privacy and personalization is therefore highly relevant, not only in the domain of ad transparency and effectiveness, but also in considerations of the more holistic relationship between consumer and firm.

Dynamic Attitudes and Practices

Consumer attitudes about the collection and use of personal information have changed over time (Affairs 2015; Greenblatt 2013), and will inevitably continue to change into the future. As devices and data rapidly enable firms to target consumers in increasingly specific ways and new targeting methods are developed, the invasive and unsavory practices of today may come to be seen as benign and acceptable. Moreover, if firms increasingly adopt ad transparency, thereby making targeting practices salient, consumers may adapt to the idea that they are being targeted, perhaps increasing the overall benefit of transparency (and possibly also consumer demand for it). While our overall framework offers flexibility necessary to mold into ever-morphing attitudes and practices, empirical research must keep pace.

Conclusion

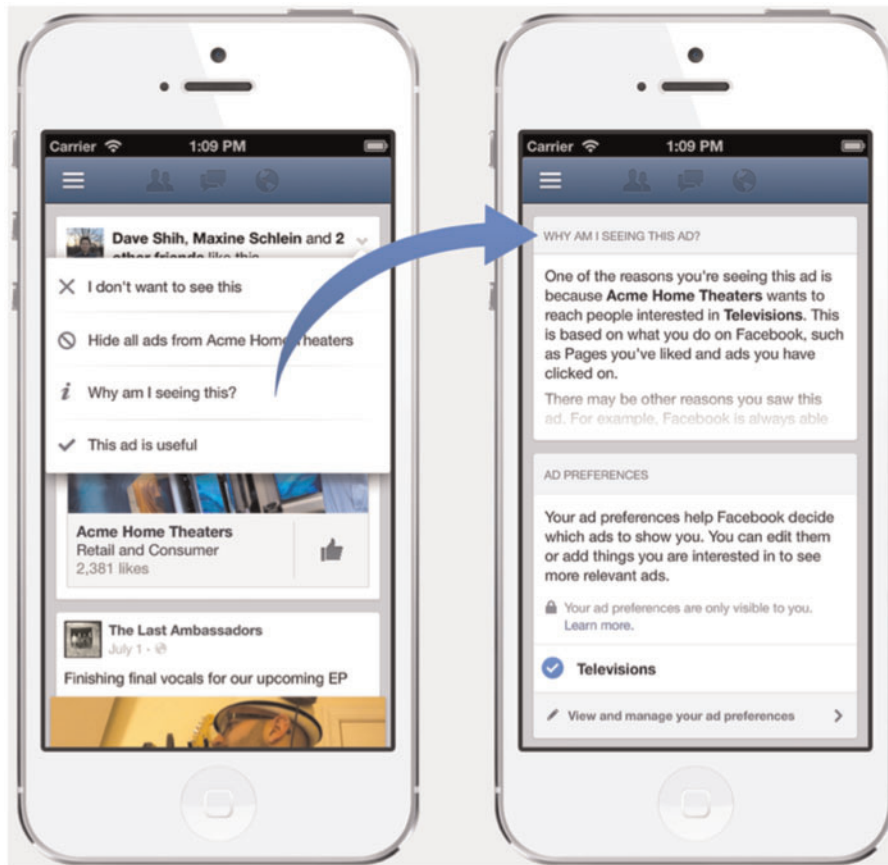
We opened by noting a growing trend of transparency in online advertising. While many advertisers may be slow to embrace ad transparency, our findings indicate that by considering norms of information flows, advertisers can mitigate the effects of exposing practices that consumers deem unsavory. For example, had Target understood and adhered to these norms, it could have avoided the now-infamous case of sending targeted, pregnancy-related coupons based on inferred information. We suggest that in addition to continually refining their targeting practices, firms might benefit by also being sensitive to consumers’ attitudes toward the process of generating ads. A perfectly targeted ad can be rendered ineffective if unsavory practices that underlie it are exposed.

DATA COLLECTION INFORMATION

The first author conducted the collection of data for studies 1–4 during 2016 and 2017 following the procedures described; studies 5a and 5b were conducted by our field site in 2017. All three authors oversaw analyses of these data.

Appendix A


AD TRANSPARENCY-FACEBOOK




APPENDIX B

STUDY 1 INSTRUCTIONS

1. Click on the link below. A new window will pop up. Please follow the rest of the instructions on the Facebook's website.
<http://www.facebook.com>
2. Log into your account.
3. The first page you see is your newsfeed. Please scroll down until you see an advertisement embedded in your newsfeed. It is likely to have a label, such as "Sponsored," "Suggested Post for you," and "Suggested Video for you" (example below).



4. Once you locate the first advertisement, click on the downward arrow button.



5. Then, click on the "why am I seeing this ad?" button.



APPENDIX C

STUDY 1: LIST OF AD PRACTICES

Items

1. A company's request to target people like me using the information that I stated on my profile
2. A company's request to target people like me using what Facebook inferred based on my Facebook usage
3. Another company's website that I've logged into without using my Facebook ID
4. Another company's website that I've logged into using my Facebook ID
5. Facebook advertisements that I click on
6. Facebook groups that I am part of
7. Facebook networks that I belong to (e.g., school, workplace)
8. Facebook pages (of companies, celebrities, etc.) that I have liked
9. Facebook pages (of companies, celebrities, etc.) that I have visited
10. Facebook pages (of companies, celebrities, etc.) that my Facebook friends have liked
11. Facebook pages (of companies, celebrities, etc.) that my Facebook friends have visited
12. How often I log into Facebook
13. My age that Facebook inferred based on my Facebook usage
14. My age that I stated on my profile
15. My current location that Facebook inferred from my computer's unique IP address
16. My current location that I stated on my profile
17. My Facebook messaging activities
18. My family members that Facebook inferred from my Facebook usage
19. My family members that I listed on my profile
20. My gender that Facebook inferred based on my Facebook usage
21. My gender that I stated on my profile
22. My past browsing history on another company's website
23. My past purchase history on another company's website
24. My past search history on a search engine
25. My past visits to another company's website
26. My relationship status that Facebook inferred based on my Facebook usage
27. My relationship status that I stated on my profile
28. My sexual orientation that Facebook inferred based on my Facebook usage
29. My sexual orientation that I stated on my profile
30. Other people's Facebook profiles that I visit

*Items excluded from the final analysis: 1, 2, 10–12, 17, and 30

APPENDIX D

STUDY 1: FACTOR ANALYSIS RESULTS

<i>Dimension</i>	<i>Name</i>	<i>Variance explained</i>	<i>Items</i>	<i>Factor loading</i>
1	Within-website tracking	17.53%	Facebook pages (e.g., companies, celebrities) that I have liked	.78
			Facebook networks that I belong to (e.g., school, workplace)	.77
			Facebook groups that I am part of	.71
			Facebook pages (of companies, celebrities, etc.) that I have visited	.58
			Facebook advertisements that I click on	.78
			Another company's website that I've logged into using my Facebook ID	.56
2	Cross-website tracking	19.81%	My current location that I stated on my profile	.59
			My past purchase history on another company's website	.86
			My past browsing history on another company's website	.87
			My past search history on a search engine	.82
			My past visits to another company's website	.84
			Another company's website that I've logged in without using my Facebook ID	.78
3	Stated personal information	17.89%	My gender that I stated on my profile	.81
			My sexual orientation that I stated on my profile	.83
			My relationship status that I stated on my profile	.81
			My age that I stated on my profile	.76
			My family members that I listed on my profile	.71
4	Inferred personal information	19.94%	My sexual orientation that Facebook inferred based on my Facebook usage	.86
			My relationship status that Facebook inferred based on my Facebook usage	.73
			My age that Facebook inferred based on my Facebook usage	.86
			My gender that Facebook inferred based on my Facebook usage	.86
			My current location that Facebook inferred from my computer's unique IP address	.59
			My family members that Facebook inferred from my Facebook usage	.74

APPENDIX D

STUDY 1: CROSS-FACTOR LOADINGS

<i>ITEMS</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
My age that Facebook inferred based on my Facebook usage.	0.86	0.18	0.20	0.16
My sexual orientation that Facebook inferred based on my Facebook usage.	0.86	0.20	0.29	0.10
My gender that Facebook inferred based on my Facebook usage.	0.86	0.18	0.26	0.12
My family members that Facebook inferred from my Facebook usage.	0.74	0.36	0.27	0.07
My relationship status that Facebook inferred based on my Facebook usage.	0.73	0.35	0.25	0.15
My current location that Facebook inferred from my computer's unique IP address.	0.59	0.38	0.13	0.25
My past browsing history on another company's website.	0.25	0.87	0.03	0.16
My past purchase history on another company's website.	0.22	0.86	0.05	0.17
My past visits to another company's website.	0.22	0.84	0.07	0.20
My past search history on a search engine.	0.37	0.82	0.06	0.10
Another company's website that I've logged into without using my Facebook ID.	0.18	0.78	0.15	0.02
My sexual orientation that I stated on my profile.	0.29	0.01	0.83	0.30
My relationship status that I stated on my profile.	0.21	0.20	0.81	0.29
My gender that I stated on my profile.	0.26	0.03	0.81	0.37
My age that I stated on my profile.	0.23	0.05	0.76	0.40
My family members that I listed on my profile.	0.35	0.25	0.71	0.15
Facebook advertisements that I click on.	0.10	0.04	0.10	0.78
Facebook pages (of companies, celebrities, etc.) that I have "liked."	-0.04	0.22	0.25	0.78
Facebook networks that I belong to (e.g., school, workplace).	0.23	0.15	0.31	0.77
Facebook groups that I am part of.	0.27	0.08	0.31	0.71
My current location that I stated on my profile.	0.28	0.04	0.47	0.59
Facebook pages (of companies, celebrities, etc.) that I have visited.	0.15	0.39	0.34	0.58
Another company's website that I've logged into using my Facebook ID.	0.01	0.48	0.25	0.56

APPENDIX E

STIMULI FROM STUDY 2

For those tasked with browsing movies within-platform: a) interactive image, b) pop-up image for *Kong: Skull Island*. For those tasked with browsing movies cross-platform, the same images served as: a) the blog’s home page b) landing page for *Kong: Skull Island*.

(a)

MARCH 2017 BOX OFFICE

Please click on the movies below to see more details.

1. Kong: Skull Island

2. Logan

3. Get Out

4. The Shack

5. The LEGO Batman Movie

6. Before I Fall

7. Hidden Figures


8. John Wick: Chapter 2

9. La La Land

10. Fifty Shades Darker

(b)

KONG: SKULL ISLAND



A team of scientists explore an uncharted island in the Pacific, venturing into the domain of the mighty Kong, and must fight to escape a primal Eden.

Director: Jordan Vogt-Roberts

Writers: Dan Gilroy, Max Borenstein, Derek Connolly, John Catin

Stars: Tom Hiddleston, Samuel L. Jackson, Brie Larson

Rating: PG-13

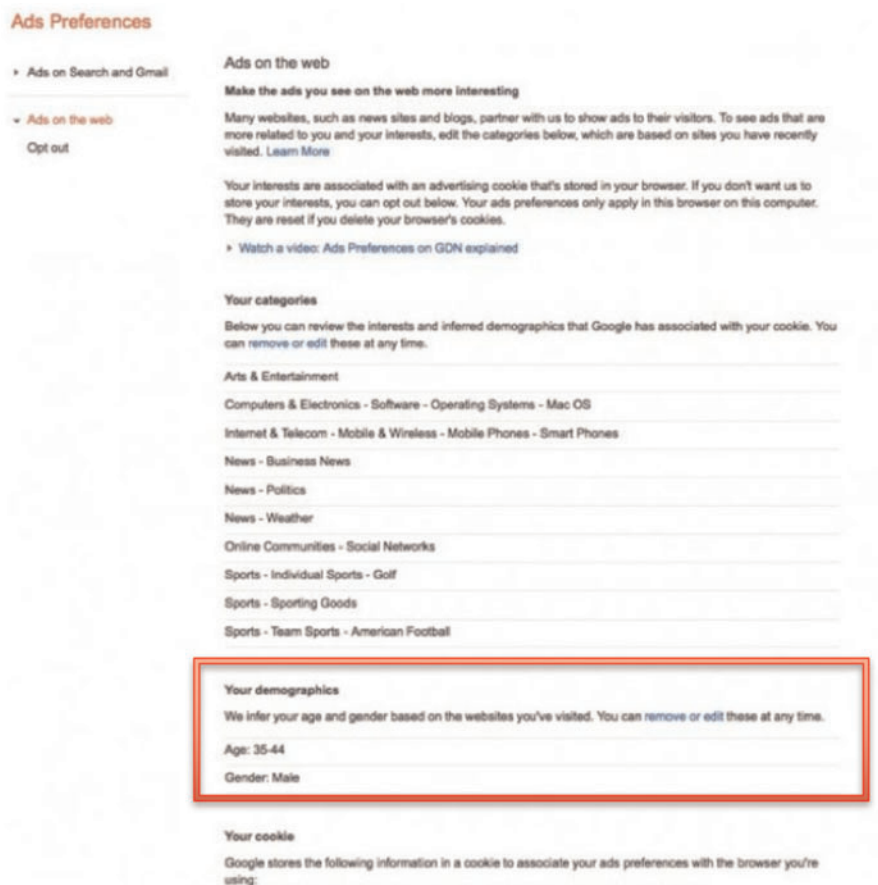
Runtime: 1h 58m

Genres: Action, Adventure, Fantasy

Please [click here](#) to return to the main page

APPENDIX G

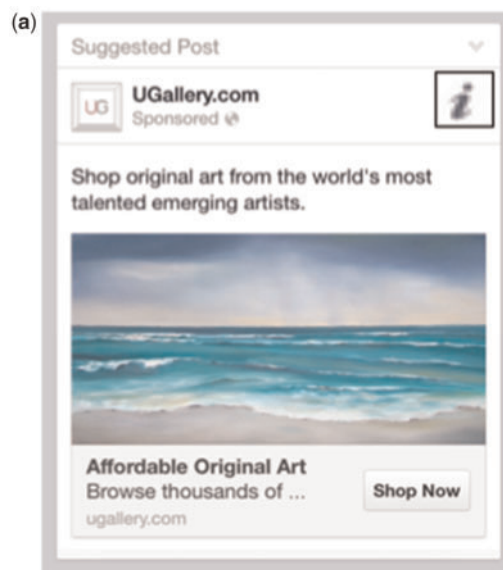
INFERRED INFORMATION BY GOOGLE

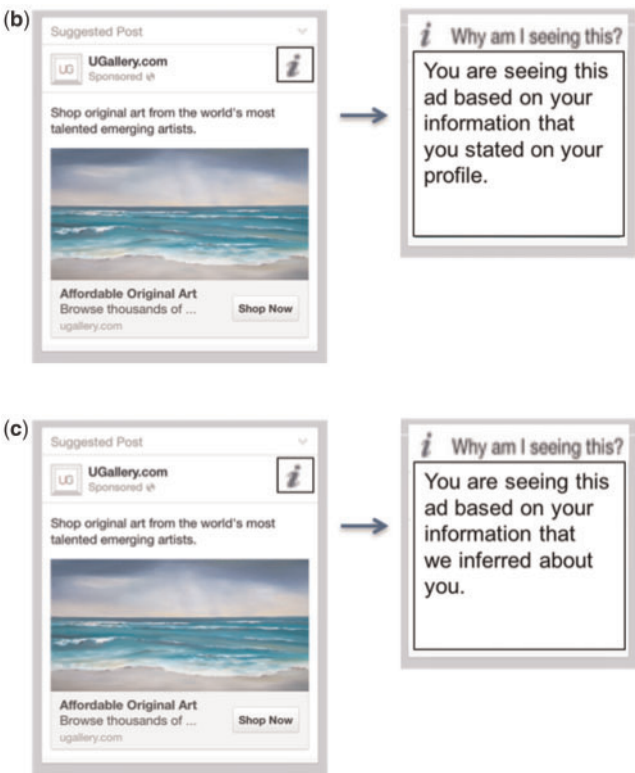


APPENDIX H

STIMULI FROM STUDY 3

Images used in study 3 for the (a) baseline, (b) stated, and (c) inferred conditions.





APPENDIX I

STUDY 4 DESIGN

Validation question:
Log into your account. Then, answer the following question. On the top-right corner of your Facebook account, there is a pull-down arrow. Which of the following does not appear on this pull down list? [Correct answer: “Help”]

- Activity log
- News Feed Preferences
- Settings
- Log Out
- Help

APPENDIX J

STUDY 4 DESCRIPTIVE STATISTICS

Percent of ad transparency messages ($N = 414$) coded as revealing the given information flow. Rows sum to 100%.


	Yes	No	I'm not sure
Within-website tracking	47.1%	37.7%	15.2%
Cross-website tracking	34.9%	46.2%	18.9%
Stated attribute	62.8%	25.6%	11.6%
Inferred attribute	25.1%	51.5%	23.4%

APPENDIX K
SCREENSHOTS FROM STUDIES 5A AND B

Study 5a

WHAT'S NEWELECTRONICSHOME & GARDENFASHIONSPORTINGKIDS & BABYGIFTS TREND MARKETPLACEHOT DEALSMORE

ELECTRONICS / CAMERAS / POINT & SHOOT / GOPRO HERO4 SILVER STANDARD



420 Points

GoPro Hero4 Silver Standard
14429

(Estimated S/H 42Points)

★★★★★

5.0 | 1 Review

1 out of 1 (100%) reviewers recommend this product


Quantity

1


Add to Shopping Bag

Add to Wish List

Recommended based on your clicks on our site



Magellan eXplorist 310 Handheld GPS
230 Points




Titan Pioneer Men's Hybrid Bicycle
178 Points

Study 5b

WHAT'S NEWELECTRONICSHOME & GARDENFASHIONSPORTINGKIDS & BABYGIFTS TREND MARKETPLACEHOT DEALSMORE

GADGETS / ELECTRONICS / GOPRO HERO4 SILVER STANDARD



17,181 Credits

GoPro Hero4 Silver Standard
14429

★★★★★

5.0 | 1 Review

1 out of 1 (100%) reviewers recommend this product


Quantity

1


Add to Shopping Bag

Add to Wish List

Recommended based on what you've shared with us



Magellan eXplorist 310 Handheld GPS
7,068 Credits



Titan Pioneer Men's Hybrid Bicycle
6,937 Credits

REFERENCES

- Aaker, Jennifer L. (1997), "Dimensions of Brand Personality," *Journal of Marketing Research*, 34 (3), 347–56.
- Acquisti, Alessandro, Leslie K. John, and George Loewenstein (2012), "The Impact of Relative Standards on the Propensity to Disclose," *Journal of Marketing Research*, 49 (2), 160–74.
- Aguirre, Elizabeth, Dominik Mahr, Dhruv Grewal, Ko de Ruyter, and Martin Wetzels (2015), "Unraveling the Personalization Paradox: The Effect of Information Collection and Trust-Building Strategies on Online Advertisement Effectiveness," *Journal of Retailing*, 91 (1), 34–49.
- Altman, Irwin and Dalmas A. Taylor (1973), *Social Penetration: The Development of Interpersonal Relationships*, Oxford, UK: Holt, Rinehart & Winston.
- Archer, Richard L. and John H. Berg (1978), "Disclosure Reciprocity and Its Limits: A Reactance Analysis," *Journal of Experimental Social Psychology*, 14 (6), 527–40.
- Ashley, Christy, Stephanie M. Noble, Naveen Donthu, and Katherine N. Lemon (2011), "Why Customers Won't Relate: Obstacles to Relationship Marketing Engagement," *Journal of Business Research*, 64 (7), 749–56.
- Baumeister, Roy F., Ellen Bratslavsky, Catrin Finkenauer, and Kathleen D. Vohs (2001), "Bad Is Stronger than Good," *Review of General Psychology*, 5 (4), 323–70.
- Bhattacharjee, Amit, Jonah Berger, and Geeta Menon (2014), "When Identity Marketing Backfires: Consumer Agency in Identity Expression," *Journal of Consumer Research*, 41 (2), 294–309.
- Bleier, Alexander and Maik Eisenbeiss (2015), "Personalized Online Advertising Effectiveness: The Interplay of What, When, and Where," *Marketing Science*, 34 (5), 669–88.
- Branscombe, Nyla, Michael Schmitt, and Richard Harvey (1999), "Perceiving Pervasive Discrimination among African Americans: Implications for Group Identification and Well-Being," *Journal of Personality and Social Psychology*, 77 (1), 135–49.
- Buell, Ryan W., Tami Kim, and Chia-Jung Tsay (2017), "Creating Reciprocal Value through Operational Transparency," *Management Science*, 63 (6), 1673–95.
- Buell, Ryan W. and Michael I. Norton (2011), "The Labor Illusion: How Operational Transparency Increases Perceived Value," *Management Science*, 57 (9), 1564–79.
- Caltabiano, Marie Louise and Michael Smithson (1983), "Variables Affecting the Perception of Self-Disclosure Appropriateness," *Journal of Social Psychology*, 120 (1), 119–28.
- Chaudhuri, Arjun and Morris B. Holbrook (2001), "The Chain of Effects from Brand Trust and Brand Affect to Brand Performance: The Role of Brand Loyalty," *Journal of Marketing*, 65 (2), 81–93.
- Clifford, Stephanie and Quentin Hardy (2013), "Attention Shoppers: Store Is Tracking Your Cell," <https://www.nytimes.com/2013/07/15/business/attention-shopper-stores-are-tracking-your-cell.html>.
- Communications Consumer Panel (2011), "Online Personal Data: The Consumer Perspective," [http://www.communication-consumerpanel.org.uk/Online personal data final 240511.pdf](http://www.communication-consumerpanel.org.uk/Online%20personal%20data%20final%20240511.pdf).
- Costa, Paul T. and Robert R. McCrae (1980), "Influence of Extraversion and Neuroticism on Subjective Well-Being: Happy and Unhappy People," *Journal of Personality and Social Psychology*, 38 (4), 668–78.
- Culnan, Mary J. (2000), "Protecting Privacy Online: Is Self-Regulation Working?" *Journal of Public Policy & Marketing*, 19 (1), 20–6.
- Dana, Jason, Roberto A. Weber, and Jason Xi Kuang (2007), "Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness," *Economic Theory*, 33 (1), 67–80.
- Derlega, Valerian J. and Alan L. Chaikin (1977), "Privacy and Self-Disclosure in Social Relationships," *Journal of Social Issues*, 33 (3), 102–15.
- Desai, Preyas S., Woochoel Shin, and Richard Staelin (2014), "The Company That You Keep: When to Buy a Competitor's Keyword," *Marketing Science*, 33 (4), 485–508.
- Dienel, Eva (2015), "Transparency, Purpose, and the Empowered Consumer: A New Paradigm for Advertising," https://www.bsr.org/reports/BSR_Participant_Responsibility_Transparency_in_Advertising_March_2015.pdf.
- Digital Advertising Alliance (2014), "YourAdChoices," <http://youradchoices.com/>.
- Doney, Patricia M. and Joseph P. Cannon (1997), "An Examination of the Nature of Trust in Buyer-Seller Relationships," *Journal of Marketing*, 61 (2), 35–51.
- Edwards, Jim (2009), "Advertisers Lobby FTC for Right to Not Tell the Truth," <https://www.cbsnews.com/news/advertisers-lobby-ftc-for-right-to-not-tell-the-truth/>.
- Farmer, Gemma (2015), "What's the Latest on the Ico Privacy Seals?" <https://iconewsblog.wordpress.com/2015/08/28/whats-the-latest-on-the-ico-privacy-seals/>.
- Fehr, Ernst and Urs Fischbacher (2004), "Social Norms and Human Cooperation," *Trends in Cognitive Sciences*, 8 (4), 185–90.
- Fehr, Ernst and Simon Gächter (2002), "Altruistic Punishment in Humans," *Nature*, 415 (6868), 137–40.
- Fishbein, Martin (1979), "A Theory of Reasoned Action: Some Applications and Implications," *Nebraska Symposium on Motivation*, 27, 65–116.
- Foreh, Mark R. and Sonya Grier (2003), "When Is Honesty the Best Policy? The Effect of Stated Company Intent on Consumer Skepticism," *Journal of Consumer Psychology*, 13 (3), 349–56.
- García Gómez, Blanca, Ana Gutiérrez Arranz, and Jesús Gutiérrez Cillán (2006), "The Role of Loyalty Programs in Behavioral and Affective Loyalty," *Journal of Consumer Marketing*, 23 (7), 387–96.
- Goldfarb, Avi and Catherine E. Tucker (2011), "Privacy Regulation and Online Advertising," *Management Science*, 57 (1), 57–71.
- Golman, Russell, David Hagmann, and George Loewenstein (2017), "Information Avoidance," *Journal of Economic Literature*, 55 (1), 96–135.
- Greenblatt, Alan (2013), "When It Comes to Online Privacy, a Disconnect for the Young," <http://www.npr.org/sections/all-techconsidered/2013/06/10/190433719/when-it-comes-to-online-privacy-a-disconnect-for-the-young>.
- Greiff, Felicia (2016), "FTC to Ad Tech Industry: More Transparency, User Control over Data Necessary," <http://www.mediapost.com/publications/article/267211/ftc-to-ad-tech-industry-more-transparency-user-c.html>.
- Grice, H. Paul, Peter Cole, and Jerry Morgan (1975), "Logic and Conversation," in *Syntax and Semantics 3: Speech Acts*, ed.

- Peter Cole and J. L. Morgan, New York: Academic Press, 41–58.
- Higgins, E. Tory (1997), “Beyond Pleasure and Pain,” *American Psychologist*, 52 (12), 1280–300.
- Hill, Kashmir (2012), “How Target Figured Out a Teen Girl Was Pregnant before Her Father Did,” <https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/#24ce85596668>.
- Hofstetter, Reto, Roland Rüppell, and Leslie K. John. (2017), “Temporary sharing prompts unrestrained disclosures that leave lasting negative impressions,” *Proceedings of the National Academy of Sciences* 201706913.
- Huang, Karen, Michael Yeomans, Alison Wood Brooks, Julia Minson, and Francesca Gino (2017), “It Doesn’t Hurt to Ask: Question-Asking Increases Liking,” *Journal of Personality and Social Psychology* 113 (3), 430–52.
- John, Leslie K., Alessandro Acquisti, and George Loewenstein (2011), “Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information,” *Journal of Consumer Research*, 37 (5), 858–73.
- John, Leslie K., Kate Barasz, and Michael I. Norton (2016), “Hiding Personal Information Reveals the Worst,” *Proceedings of the National Academy of Sciences*, 113 (4), 954–9.
- Kroft, Steve (2013), “The Data Brokers: Selling Your Personal Information,” <https://www.cbsnews.com/news/data-brokers-selling-personal-information-60-minutes/>.
- Lambrecht, Anja and Catherine Tucker (2013), “When Does Retargeting Work? Information Specificity in Online Advertising,” *Journal of Marketing Research*, 50 (5), 561–76.
- Learmonth, Michael (2009), “Sensing Coming Regulation, Online Ad Groups Unite,” <http://adage.com/article/digital/online-ad-groups-fight-regulation-behavioral-targeting/133730/>.
- Lin-Healy, Fern and Deborah A. Small (2013), “Nice Guys Finish Last and Guys in Last Are Nice: The Clash between Doing Well and Doing Good,” *Social Psychological and Personality Science*, 4 (6), 692–8.
- 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.
- Marketing Charts (2014), “Consumers Say They Find Personalized Ads More Engaging and Memorable,” <http://www.marketingcharts.com/online/consumers-say-they-find-personalized-ads-more-engaging-and-memorable-43005/>.
- Milne, George R. and Mary Ellen Gordon (1993), “Direct Mail Privacy-Efficiency Trade-Offs within an Implied Social Contract Framework,” *Journal of Public Policy & Marketing*, 12 (2), 206–15.
- Montgomery, Alan L. and Michael D. Smith (2009), “Prospects for Personalization on the Internet,” *Journal of Interactive Marketing*, 23 (2), 130–7.
- Moon, Youngme (2000), “Intimate Exchanges: Using Computers to Elicit Self-Disclosure from Consumers,” *Journal of Consumer Research*, 26 (4), 323–39.
- Morewedge, Carey K. (2009), “Negativity Bias in Attribution of External Agency,” *Journal of Experimental Psychology: General*, 138 (4), 535–45.
- Morey, Timothy, Theodore Theo Forbath, and Allison Schoop (2015), “Customer Data: Designing for Transparency and Trust,” *Harvard Business Review*, 93 (5), 96–105.
- Newman, George E. and Daylian M. Cain (2014), “Tainted Altruism: When Doing Some Good Is Evaluated as Worse than Doing No Good at All,” *Psychological Science*, 25 (3), 648–55.
- Nissenbaum, Helen (2004), “Privacy as Contextual Integrity,” *Washington Law Review*, 79 (1), 119–58.
- (2011), “A Contextual Approach to Privacy Online,” *Daedalus*, 140 (4), 32–48.
- Paul, Kari (2017), “When Facebook and Instagram Think You’re Depressed,” https://www.vice.com/en_us/article/when-facebook-and-instagram-thinks-youre-depressed.
- Pavlou, Paul A., Huigang Liang, and Yajiong Xue (2006), “Understanding and Mitigating Uncertainty in Online Environments: A Principal-Agent Perspective,” *MIS Quarterly*, 31 (1), 105–36.
- Phelps, Joseph, Glen Nowak, and Elizabeth Ferrell (2000), “Privacy Concerns and Consumer Willingness to Provide Personal Information,” *Journal of Public Policy & Marketing*, 19 (1), 27–41.
- Public Affairs (2015), “Is Online Privacy Over? Findings from the USC Annenberg Center for the Digital Future Show Millennials Embrace a New Online Reality,” <http://annenberg.usc.edu/news/faculty/online-privacy-over-findings-usc-annenberg-center-digital-future-show-millennials>.
- Rainie, Lee and Maeve Duggan (2016), “Privacy and Information Sharing,” <http://www.pewinternet.org/2016/01/14/privacy-and-information-sharing/>.
- Ramirez, Edith, Julie Brill, Maureen K. Ohlhausen, Joshua D. Wright, and Terrell McSweeney (2014), *Data Brokers: A Call for Transparency and Accountability*, Washington, DC: Federal Trade Commission.
- Rozin, Paul and Edward B. Royzman (2001), “Negativity Bias, Negativity Dominance, and Contagion,” *Personality and Social Psychology Review*, 5 (4), 296–320.
- Rubin, Zick (1975), “Disclosing Oneself to a Stranger: Reciprocity and Its Limits,” *Journal of Experimental Social Psychology*, 11 (3), 233–60.
- Sayed, Amin, Kinshuk Jerath, and Kannan Srinivasan (2014), “Competitive Poaching in Sponsored Search Advertising and Its Strategic Impact on Traditional Advertising,” *Marketing Science*, 33 (4), 586–608.
- Schumann, Jan H., Florian von Wangenheim, and Nicole Groene (2014), “Targeted Online Advertising: Using Reciprocity Appeals to Increase Acceptance among Users of Free Web Services,” *Journal of Marketing*, 78 (1), 59–75.
- Singer, Eleanor, Hans-Juergen Hippler, and Norbert Schwarz (1992), “Confidentiality Assurances in Surveys: Reassurance or Threat?” *International Journal of Public Opinion Research*, 4 (3), 256–68.
- Smith, H. Jeff, Sandra J. Milberg, and Sandra J. Burke (1996), “Information Privacy: Measuring Individuals’ Concerns about Organizational Practices,” *MIS Quarterly*, 20 (2), 167–96.
- Stewart, Katherine J. (2003), “Trust Transfer on the World Wide Web,” *Organization Science*, 14 (1), 5–17.
- Summers, Christopher A., Robert W. Smith, and Rebecca Walker Reczek (2016), “An Audience of One: Behaviorally Targeted Ads as Implied Social Labels,” *Journal of Consumer Research*, 43 (1), 156–78.
- Tam, Kar Yan and Shuk Ying Ho (2006), “Understanding the Impact of Web Personalization on User Information Processing and Decision Outcomes,” *MIS Quarterly*, 30 (4), 865–90.

- Tsai, Janice Y., Serge Egelman, Lorrie Cranor, and Alessandro Acquisti (2011), "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study," *Information Systems Research*, 22 (2), 254–68.
- Tucker, Catherine (2014), "Social Networks, Personalized Advertising, and Privacy Controls," *Journal of Marketing Research*, 51 (5), 546–62.
- Turow, Joseph, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy (2009), "Americans Reject Tailored Advertising and Three Activities That Enable It," Departmental Paper, Annenberg School for Communication, University of Pennsylvania, Philadelphia, PA 19104.
- Waldman, Ari Ezra (2016), "Privacy, Sharing, and Trust: The Facebook Study," *Case Western Reserve Law Review*, 67 (1), 193–233.
- Wang, Wei-quan and Izak Benbasat (2005), "Trust in and Adoption of Online Recommendation Agents," *Journal of the Association for Information Systems*, 6 (3), 72–101.
- Wasserman, Todd (2012), "Facebook Now Has 901 Million Users," https://mashable.com/2012/04/23/facebook-now-has-901-million-users/#ZM.bMfVz_8qL.
- White, Tiffany Barnett (2004), "Consumer Disclosure and Disclosure Avoidance: A Motivational Framework," *Journal of Consumer Psychology*, 14 (1–2), 41–51.
- White, Tiffany Barnett, Debra L. Zahay, Helge Thorbjørnsen, and Sharon Shavitt (2008), "Getting Too Personal: Reactance to Highly Personalized Email Solicitations," *Marketing Letters*, 19 (1), 39–50.
- Yang, Sha, Shijie Lu, and Xianghua Lu (2014), "Modeling Competition and Its Impact on Paid-Search Advertising," *Marketing Science*, 33 (1), 134–53.
- Yoon, Yeosun, Zeynep Gürhan-Canli, and Norbert Schwarz (2006), "The Effect of Corporate Social Responsibility (CSR) Activities on Companies with Bad Reputations," *Journal of Consumer Psychology*, 16 (4), 377–90.
- Zhang, Kaifu and Zsolt Katona (2012), "Contextual Advertising," *Marketing Science*, 31 (6), 980–94.