

The Importance of Trust for Personalized Online Advertising

Alexander Bleier^{a,1}, Maik Eisenbeiss^{b,*}

^a Marketing Department, Carroll School of Management, Boston College, 140 Commonwealth Avenue, Fulton Hall 448, Chestnut Hill, MA 02467, United States

^b Department of Marketing, University of Bremen, Hochschulring 4, 28359 Bremen, Germany

Abstract

With the amount of online advertising on a steady rise, generic ads noticeably lose effectiveness. In order to break through the clutter, retailers employ a method called retargeting to tailor their advertisements to individual consumers based on inferred interests and preferences. However, while personalization should generally make ads more appealing, the authors use field data to show that the effectiveness of retargeting considerably hinges on consumers' trust in a respective retailer. To uncover the underlying mechanisms of this phenomenon, they investigate how trust moderates the impact of ad personalization on consumers' internal and external responses in the lab. They propose a two-dimensional conceptualization of ad personalization: First, a banner's personalization *depth* defines how *closely* the ad reflects a consumer's interests. Second, its personalization *breadth* determines how *completely* the banner reflects these interests. The lab results show that more trusted retailers can increase the perceived usefulness of their ads through a combination of high depth and narrow breadth of personalization without eliciting increased reactance or privacy concerns. On the other hand, for less trusted retailers, banners with higher depth are not perceived more useful, but instead trigger increased reactance and privacy concerns, regardless of their personalization breadth. These effects directly translate into consumers' click-through intentions so that retailers should adjust their personalization strategies accordingly in order to increase the effectiveness of their online advertising.

© 2015 New York University. Published by Elsevier Inc. All rights reserved.

Keywords: Retargeting; Online advertising; Personalization; Trust; Privacy; Reactance

Introduction

The retail industry constantly intensifies its online advertising efforts (eMarketer 2014b), so that standard display banners increasingly struggle to gain consumers' attention (Cho and Cheon 2004). As a result, their click-through rates have come down to as low as .08% (Sizmek 2014). In order to achieve greater effectiveness, many retailers now tailor their online ads to individual consumers with an instrument called retargeting (Helft and Vega 2010; Peterson 2013; Sengupta 2013).

Retargeting banners typically feature product images that match consumers' interests, inferred from their most recent shopping behavior in the advertising retailer's online store. Interestingly, while retargeting, as a special form of personalization, should make an ad more relevant and appealing (Ansari and Mela 2003; Chen, Pavlov, and Canny 2009), ad personalization

can also cause unfavorable responses. For example, consumers oftentimes feel manipulated, or deprived of their freedom of choice when perceiving a personalized advertisement inappropriately close to their preferences (King and Jessen 2010; Tucker 2012b; White et al. 2008). Personalized ads may also elicit privacy concerns (Goldfarb and Tucker 2011; Okazaki, Li, and Hirose 2009; Turow et al. 2009), especially since every personalized ad signals that the retailer has tracked and analyzed detailed information about the consumer's browsing activities and is willing to exploit this knowledge (Anand and Shachar 2009). It is therefore not surprising that a large majority of consumers still do not want retailers to adjust advertisements to their online behaviors across websites (Guild 2013; Turow et al. 2009).

Correspondingly, results concerning the ability of retargeting to increase ad effectiveness are mixed. While a number of industry studies proclaim its merits (Criteo 2015; eMarketer 2014a; Hunter et al. 2010), academic research finds retargeting to be only under certain circumstances superior to generic advertising (Lambrecht and Tucker 2013). While extant research has uncovered a number of favorable and unfavorable consumer responses to ad personalization (e.g., Van Doorn and Hoekstra 2013; White

* Corresponding author. Tel.: +49 421 218 66740; fax: +49 421 218 66741.

E-mail addresses: bleiera@bc.edu (A. Bleier), eisenbeiss@uni-bremen.de (M. Eisenbeiss).

¹ Tel.: +1 617 552 1870; fax: +1 617 552 6677.

et al. 2008), insight on moderating factors that strengthen or weaken these effects is still largely missing.

To fill this void, we examine trust in the retailer as a key moderator of the impact that retargeting exerts on favorable and unfavorable consumer responses. Trust is a major success factor in the online domain. In fact, on the Web trust often serves as the sole foundation on which consumers base their research and purchase decisions in lack of further information about firms (McStay 2011; Urban, Amyx, and Lorenzon 2009; Urban, Sultan, and Qualls 2000). At first glance, trust might seem to be a rather obvious means to enhance the performance of retargeting. For example, trust in a retailer may make consumers believe their personal data to be in safe hands and thus alleviate possible privacy concerns. Yet, its influence on other responses is less obvious. It is, for instance, not clear whether trust leads consumers to perceive a personalized ad as more useful or make them feel less manipulated by its appearance.

Moreover, retargeting is a complex and largely unexplored method of ad personalization. To derive meaningful insights about its impact on consumers we propose a new generalized perspective. Specifically, we suggest that ad personalization with retargeting can be described along the two dimensions of a banner's depth and breadth of personalization. Personalization depth, defined as the primary dimension of ad personalization, determines how *closely* a banner reflects a consumer's inferred interests. For example, an ad that features products a consumer previously placed in his or her virtual shopping cart reflects that person's interests closer than an ad that features products the consumer merely inspected without any further purchase-related actions. This dimension has been shown to significantly influence a banner's effectiveness (Bleier and Eisenbeiss 2015; Lambrecht and Tucker 2013). In addition, we introduce personalization breadth as a secondary dimension of ad personalization. Personalization breadth describes how *completely*, or exhaustively, a banner with a given personalization depth reflects a consumer's interests. Accordingly, a banner with wide breadth might show all products that the consumer placed in the shopping cart, while a banner with narrow breadth might only feature a subset of these items. So far this dimension has not been empirically studied. The moderating influence of trust in the retailer on consumers' responses to retargeting banners with different combinations of personalization depth and breadth is therefore even less obvious. With these aspects in mind, our study primarily contributes to the existent literature in three ways.

First, we empirically investigate whether and to which extent trust in the retailer moderates the effect of ad personalization through retargeting on consumers. To this end we use data from a quasi-experimental field study and a lab experiment. The field data reveal that a higher depth of personalization increases banners' click-through rates only for the more trusted retailer, but actually decreases click-through rates for the less trusted retailer. Second, we introduce personalization breadth as a second key dimension of ad personalization. We demonstrate its relevance in the lab and show how the effectiveness of banners' with specific depth/breadth combinations depends on consumers' trust in the advertising retailer. Third, and especially important from a theory perspective, we shed further

light on the internal mechanisms through which personalized online advertising influences consumers. In particular, we explicitly distinguish between perceived usefulness, reactance, and privacy concerns as internal responses that previous studies have not dissected and either examined in isolation or with certain ambiguity (e.g., Tucker 2014; Van Doorn and Hoekstra 2013; White et al. 2008). We then link these responses to consumers' click-through intentions. Altogether, our work helps explain the existent mixed findings about the impact of personalized online advertising on consumers.

The remainder of this article is organized as follows. We first report the results of a quasi-experimental field study. The purpose of this field study is to explore in a real-world retargeting context how trust in the retailer moderates the effectiveness of personalized banners with high and low depth as the primary dimension of ad personalization. Next, we introduce our conceptual framework that is grounded on the stimulus-organism-response paradigm. We introduce personalization breadth as the secondary dimension of ad personalization and link a retargeting banner stimulus with a given combination of personalization depth and breadth to internal and external consumer responses. We then propose how trust in the advertiser moderates these effects and derive corresponding hypotheses. We test the conceptual framework and our hypotheses based on the results from a lab experiment. Last, we interpret our findings, provide concrete managerial interpretations, and indicate possibilities for future research.

Pilot Study: The Moderating Role of Trust in the Retailer

Objective

We begin our research with a pilot study, using quasi-experimental field data on retargeting to explore whether, how, and to which extent consumers' trust in a retailer influences the effectiveness of its personalized online advertising in a real-world setting.

Data Description

We cooperated with an advertising agency that specializes in retargeting to obtain data on click-through responses to personalized display banner ads. As is common industry practice, the agency creates these banners and delivers them to consumers on behalf of its client retailers while the agency itself remains invisible to consumers. In retargeting, an ad agency tailors banner ads to individual consumers based on their most recent shopping behavior at a retailer's online store. At each visit, consumers' interactions with products are recorded in individual profiles through clickstream data (Bucklin and Sismeiro 2009). The ad agency then applies certain algorithms, so called assembly rules, to create banners based on these interactions. Assembly rules typically differ in their depth of personalization, i.e., how closely they reflect consumers' interests and preferences, inferred from their behavior in the retailer's online store.

In order to investigate whether trust in the retailer influences the effectiveness of retargeting, we selected two firms from our

Table 1
Results of the pilot study.

Retailer	Personalization depth	Ad impressions	Click-through rate	Std Dev	Min	Max
A (more trusted)	Low	1,358,002	0.0040	0.0629	0	1
	High	162,142	0.0051	0.0712	0	1
B (less trusted)	Low	1,438,960	0.0037	0.0607	0	1
	High	76,772	0.0020	0.0446	0	1

partnering ad agency's client base that were very similar in several regards, but differed substantially with respect to the average trust that consumers placed in them.² We then compared how well retargeting banners with different depths of personalization worked for them. We identified two firms that were comparable to the extent that both were major German retailers (over \$500m revenues p.a.; more than 1,500 employees) with a concentration on selling fashion merchandise for men, women, and children online. Both retailers engaged in substantial retargeting efforts of similar magnitude, i.e., the ad agency delivered about the same number of ad impressions for them every week.

Yet, the two retailers differed significantly in the average level of trust that consumers placed in them, according to a pretest we conducted with 72 online shoppers who were familiar with the respective retailers ($M_{\text{age}} = 27$, 51% female). We measured trust with the items "this retailer can be relied on to keep its promises," "this retailer puts the customer's interest first," "this retailer usually keeps the promises that it makes to me," and "I can count on this retailer to provide a good service" (Verhoef, Franses, and Hoekstra 2002) on a seven-point rating scale from 1 ("Strongly disagree") to 7 ("Strongly agree"). The results show that trust in retailer A was significantly above the scale average ($M = 4.578$, $t = 3.82$, $p < .001$) while trust in retailer B was significantly below the scale average ($M = 3.504$, $t = -3.54$, $p < .001$). Correspondingly, the difference between their levels of trust was also significant ($p < .001$).

At the time of data collection, our partnering ad agency employed the same assembly rules to create banners for both firms. For comparable time intervals, we analyzed two of these rules that differed with respect to their depth of personalization. The two rules were:

- (1) *Low personalization depth*: A banner features products sampled from those items that a consumer viewed during his or her most recent shopping session.
- (2) *High personalization depth*: A banner features products sampled from those items that a consumer placed in the virtual shopping cart during his or her most recent shopping session.

Drawing on choice theory, products that a consumer merely views during a shopping session represent his or her real interests and preferences with considerably lower accuracy, i.e., less closely, than products that he or she places in the virtual shopping cart (Hoeffler and Ariely 1999). Banners featuring

products according to rule (1) have therefore a lower depth of personalization than banners assembled according to rule (2). For both retailers and assembly rules, the ad agency provided us with consumer-level records of all ad impressions that they received during a time interval of one week. The data also included consumers' responses to each ad impression in terms of click-through (click/no click).

Results and Discussion

With regard to these two assembly rules, the agency delivered a total of 1,520,144 ad impressions on behalf of retailer A and 1,515,732 ad impressions on behalf of retailer B during the observation period, as shown in Table 1. Note that these almost identical ad impression levels underscore that both firms are comparable in terms of their general retargeting efforts.

To compare the effectiveness of banners with low and high personalization depth between the two retailers, we calculated the respective average click-through rates (including their 95% confidence levels). The results appear in Fig. 1.

For banners with low personalization depth, the average click-through rates are similar for both retailers ($M_{\text{Retailer A, low}} = .40\%$, $M_{\text{Retailer B, low}} = .37\%$). By contrast, for banners with high personalization depth the average click-through rates differ significantly between them ($M_{\text{Retailer A, high}} = .51\%$, $M_{\text{Retailer B, high}} = .20\%$). Most interestingly, a higher personalization depth increases banners' click-through rates for the more trusted retailer A significantly (relative change of +27%; $M_{\text{Retailer A, low}} = 0.0040$, $M_{\text{Retailer A, high}} = 0.0051$; $t = 6.750$, $p < .001$), while it significantly decreases click-through rates for banners of




Fig. 1. Click-through rates of the pilot study. Note: Error bars denote 95% confidence intervals.

² The retailers' names cannot be revealed for confidentiality reasons.




Fig. 2. Conceptual framework.

the less trusted retailer B (relative change of -46% ; $M_{\text{Retailer B, low}} = 0.0037$, $M_{\text{Retailer B, high}} = 0.0020$; $t = -7.690$, $p < .001$).

Since trust is one major differentiating factor between the two firms while they are similar with respect to many other characteristics (e.g., industry, firm size, retargeting efforts), the results provide a first empirical indication that trust in a retailer moderates the effect of ad personalization on consumers. Specifically, the results suggest that more trusted retailers benefit from tailoring their banners closely to consumers' interests with high-depth banners while less trusted retailers actually endanger the effectiveness of their banners when employing high personalization depth.

This pilot study is of course subject to certain limitations. First, the data represent only a quasi-experimental comparison, since the respective assembly rules were applied to consumers based on their prior shopping behavior. This, however, occurred equally for prospects of retailers A and B because the ad agency applied the same matching algorithm for both of them. Our relative comparisons between rules and retailers are therefore valid, while the absolute click-through rates might be inflated to some degree. Next, possible unobservable factors besides trust that also discriminate between both retailers might have confounded the observed effects. While we sought to rule out such influences as much as possible by selecting two retailers that are comparable along several dimensions, the difference in trust provides only one explanation for the observed results. To further explore the moderating role of trust in a controlled setting, we conducted a lab experiment that we describe after introducing our corresponding conceptual framework. In this course, we also expand the conceptualization of ad personalization through personalization depth with the secondary dimension of personalization breadth to obtain a more accurate representation of the myriad personalization approaches employed in today's retargeting practice.

Conceptual Framework

The empirical results from the pilot field study suggest that trust in the retailer can explain why the same depth of personalization leads to an increase in click-through effectiveness for one retailer, while it causes a decrease in effectiveness for another

one. However, to profoundly understand how trust moderates the impact of ad personalization on consumers, it is important to uncover the internal processes that govern these behavioral responses. The stimulus–organism–response (S-O-R) paradigm (e.g., Jai, Burns, and King 2013; Li, Sarathy, and Xu 2011) provides an appropriate general framework to this end. It proposes that environmental cues (stimuli) influence individuals' affective and cognitive responses (internal responses of an organism), which then translate into specific behavior (external responses) (Donovan and Rossiter 1982; Mehrabian and Russell 1974). Adapted to our research context, a retargeting banner represents a stimulus that affects consumers' internal responses which then trigger their external, behavioral responses. Fig. 2 depicts the corresponding conceptual framework that we explain next. We begin by introducing the stimulus, i.e., a retargeting banner's personalization approach. Besides using personalization depth as the primary dimension of ad personalization, we introduce personalization breadth as a secondary personalization dimension. It determines how completely banners with a given personalization depth reflect consumers' preferences. Next, we delineate how this stimulus can affect consumers' internal responses in terms of perceived usefulness, reactance, and privacy concerns and how these internal responses culminate in click-through, as the ultimate external response. Last, we discuss how trust in the retailer possibly moderates the effects of the stimulus on consumers' internal responses and derive corresponding hypotheses.

Stimulus: Depth and Breadth of Personalization

Retargeting has become a popular instrument to personalize display banners to individual consumers. Firms currently use different personalization algorithms which we propose can in essence be summarized along two key dimensions: a banner's depth and breadth of personalization. We define personalization depth as the *primary dimension* of personalization which constitutes how *closely* an ad reflects a consumer's inferred interests and preferences. It describes a personalization continuum as similarly discussed in Kreuter and Wray (2003) that features at the lowest level only the name of the retailer, instead of actual products, catering to the idea that a consumer who visits a retailer's online store is generally interested in the firm and

its offerings. A banner with a somewhat higher personalization depth might include products from the category that a consumer primarily viewed in the retailer's online store or contain even exactly these viewed items. At an even higher personalization depth, a banner might feature products that a consumer previously placed in his or her virtual wish list or shopping cart, thereby coming closest to that individual's revealed interests and preferences (see pilot study for an application of this example).

Next to personalization depth, we define personalization breadth as the secondary dimension of ad personalization. This dimension describes how *completely*, or exhaustively, an ad with a given personalization depth reflects a consumer's inferred interests and preferences. During the purchase decision process, consumers may browse through and interact with several products of interest. This group of products is commonly referred to as the consideration set. Typically, display banners do not cover the entirety of these products and only feature a subset of items. For instance, ads featuring products that a consumer previously placed in his or her virtual wish list (this represents the banner's personalization depth) may either include relatively many (wide personalization breadth) or relatively few (narrow personalization breadth) of these products.

Technically, a banner's breadth of personalization defines the featured subset of all eligible products determined by a given personalization depth. In this sense, depth forms the primary dimension and breadth the secondary dimension of ad personalization. Personalization breadth obviously only applies to banners exceeding the very low end of the personalization depth continuum, because these banners are, as explained, generic and do not feature any specific products (Lambrecht and Tucker 2013).

Previous research provides empirical indications that both the depth and breadth of personalization can influence banner effectiveness. Bleier and Eisenbeiss (2015) show how banners of differing depth trigger click-through to varying extents, depending on consumers' current positions in the purchase decision process and the time that has passed since their last visit of the advertising retailer's online store. Moreover, Lambrecht and Tucker (2013) find high-depth banners to only generate more sales than low-depth banners when consumers have narrowly construed preferences and are actively involved in the advertised category.

No study so far has investigated how personalization breadth influences consumers' responses in the context of online banner advertising. Yet, for recommendation agents, Murthi and Sarkar (2003) highlight that the number and combination of products featured in offerings pose a knapsack problem. In these settings, firms have to carefully decide which specific items from consumers' initial consideration sets to display in recommendations whose capacities of includable items are constrained. Relatedly, Tam and Ho (2005) find that the recommendation set size in advertising messages positively influences consumers' attention.

From a managerial perspective these two dimensions are highly relevant, as their varying configurations used in practice underscore. Yet, to the best of our knowledge, no research has

studied personalization depth and breadth in combination so far. We next discuss how these dimensions may affect consumers' responses to personalized online advertising.

Responses to Personalized Online Advertising

Previous research highlights that ad personalization may cause favorable as well as unfavorable consumer responses (Tucker 2012b). However, often especially the unfavorable evaluations are conceptualized ambiguously under identical constructs or as amalgams of qualitatively different aspects. For example, White et al. (2008) study under which circumstances tailored advertising emails trigger reactance to their personalization approach. The authors argue that ad messages containing personal information threaten consumers' perceived ability to avoid being closely observed by firms. Tucker (2014) investigates personalized display banners adjusted to consumers' behavior at a social network website. This study's interpretation of reactance differs to some extent from the understanding of White et al. (2008). While White et al. (2008) explicitly design their lab experiment as to not capture privacy concerns, Tucker (2014) argues consumers to experience reactance in response to perceived privacy intrusiveness of personalized ads. Moreover, Van Doorn and Hoekstra (2013) use perceived intrusiveness to measure consumers' negative responses to display banners which constitutes a combination of conceptually distinct consumer responses such as reactance and privacy concerns. In this research, we clearly differentiate between reactance and privacy concerns as two distinct negative internal responses to personalized advertising which are also generally recognized by Aguirre et al. (2015).

Moreover, we capture positive internal responses through perceived usefulness. Usefulness has long since been a crucial success measure for advertising (Shavitt, Lowrey, and Haefner 1998). It is also frequently employed in online personalization research (e.g., Nysveen and Pedersen 2004) and the closely related research stream on recommendation agents (e.g., Tam and Ho 2006; Xiao and Benbasat 2007) which underscores its importance for the examination of personalized online ads.

Conceptually, we link these internal responses to click-through as consumers' ultimate external, behavioral response and principal measure of online display advertising effectiveness (PricewaterhouseCoopers 2010). In the following, we describe how a banner's personalization depth and breadth affect these favorable and unfavorable internal responses.

Usefulness

Consumers perceive advertising communications oftentimes as useful for their purchase decisions (Shavitt, Lowrey, and Haefner 1998). In particular, adjusting ad messages to consumers' physical locations or personal information is regarded useful (Conti et al. 2012; Zhang et al. 2010). Adapting ads to a consumer's processing goals and self-schemas also enhances their relevance and perceived usefulness which makes recipients more inclined to further explore the advertised content (Tam and Ho 2006). A banner's depth of personalization describes how closely the ad reflects an individual's interests and preferences.

A higher depth corresponds to an increased relevance of its content to that individual and should thus be seen as more useful. We therefore expect a positive effect of personalization depth on usefulness.

Besides catering to individual preferences, personalization can enhance the value of ad communications by reducing information overload (Ansari and Mela 2003). In the related context of recommendation agents, previous research shows that presenting limited selections of relevant options can effectively simplify consumers' choice decisions (Häubl and Trifts 2000; Murthi and Sarkar 2003) and thereby increase their choice confidence and subsequent satisfaction (Chernev 2003; Iyengar and Lepper 2000). Since a banner's breadth of personalization defines the subset of all relevant products determined by a given personalization depth, a narrower breadth featuring a smaller selection of these products should make a banner more useful for consumers. We thus expect the positive effect of personalization depth on usefulness to be even stronger for banners with narrow breadth.

Reactance

By contrast, advertisements can also have negative effects on consumers. In particular personalized ads can lead consumers to feel manipulated, or threatened in their freedom of choice when perceiving a personalized ad inappropriately close to their preferences (King and Jessen 2010; Tucker 2012b; White et al. 2008). The objective of personalized advertising is to present consumers with individualized offerings that meet their personal likings. However, similar to recommendation agents, confronting consumers with ads featuring select products that a retailer deems relevant to them can cause consumers to perceive their free choice processes being intruded upon (Fitzsimons and Lehmann 2004). The result is a psychological state called reactance in which consumers are increasingly motivated to restore their freedom of choice by behaving counter to the intention of the threatening stimulus (Brehm 1966; Brehm 1989; Brehm and Brehm 1981). For example, when consumers' esthetic choices are predicted on the basis of a personality test, they tend to choose against those predictions the more they believe that the test would indeed be able to foretell their preferences (Hannah, Hannah, and Wattie 1975). Similar effects are well documented in social psychology (Goodstadt 1971; Kahn and Tice 1973; Worchel and Brehm 1971) and health communications (Grandpre et al. 2003; Miller et al. 2007) and have also been shown in marketing research on recommendation agents (Fitzsimons and Lehmann 2004; Lee and Lee 2009), promotions (Clee and Wicklund 1980; Kivetz 2005), and general marketing communications (Godfrey, Seiders, and Voss 2011; White et al. 2008). Since banners with high personalization depth closely reflect consumers' inferred interests and preferences, they should cause more reactance compared to banners with low personalization depth. We therefore expect a positive effect of personalization depth on reactance.

Moreover, a further antecedent of reactance is the unavailability of specific choice options which influences consumers' evaluations of remaining alternatives (Clee and Wicklund 1980; Kramer and Carroll 2009; Mazis, Settle, and Leslie 1973).

For example, research finds consumers to ignore unsolicited advice from recommendation agents that advertise specific options while eliminating others (Fitzsimons and Lehmann 2004). Accordingly, for a banner with a given level of personalization depth a more narrow personalization breadth should more likely trigger reactance, because it features only a subset of relevant products. We therefore expect the positive effect of personalization depth on reactance to be even stronger for banners with narrow breadth.

Privacy concerns

A related, yet conceptually different aspect refers to the fact that ad personalization involves collecting, analyzing, and leveraging consumers' private information beyond its original transaction purpose. Permission for this practice is seldom solicited, so that consumers are usually not aware of firms' undertakings until they receive an individualized communication (Aguirre et al. 2015; Rapp et al. 2009). The marketing literature regards these activities as the predominant triggers of consumers' concerns for privacy (Foxman and Kilcoyne 1993; Malhotra, Kim, and Agarwal 2004; Sheehan and Hoy 2000) where privacy is defined as "the ability of the individual to control the terms under which personal information is acquired and used" (Westin 1967, p. 7). A loss of control due to a privacy intrusion reflects a consumer's vulnerability to the data-collecting and advertising firm and prompts privacy concerns (Bandyopadhyay 2011; Dinev and Hart 2004; Raab and Bennett 1998). These effects are well-documented for general online personalization (Chellappa and Sin 2005; Kobsa 2007; Taylor, Davis, and Jillapalli 2009) and personalized online advertising. According to Turow et al. (2009), 84% of all consumers do not want advertisers to tailor ads to their behavior that was tracked on other websites than the one they currently visit. Correspondingly, Aguirre et al. (2015) show sharp declines in click-through rates when consumers realize their vulnerability to a retailer that has covertly collected their data to provide personalized banners. Moreover, even adjusting ads to a currently visited website elicits privacy concerns when banners are obtrusive (Goldfarb and Tucker 2011). The extent to which ads trigger privacy concerns depends on the sensitivity of consumer-specific information they are based on (Nowak and Phelps 1992). For any personalized ad, its depth of personalization likely serves as a proxy for the extent to which the advertiser has gained knowledge about a consumer (Anand and Shachar 2009; Winer 2001). It thereby reveals the level of vulnerability to the advertising firm and determines the amount of elicited privacy concerns. Suppose for instance a consumer browsed several cars in an automobile dealer's online store before receiving a personalized banner. With low personalization depth that ad might feature an image of the dealership's logo, signaling general knowledge about the consumer's visit to its online store. By contrast, a banner with high depth might specifically feature images of the inspected cars and thereby signal very detailed knowledge of the consumer's browsing activities in the online store. In the latter example with high depth the perceived vulnerability to the retailer should be higher and result in increased privacy concerns. Consequently, we expect a banner with high personalization depth to cause

more privacy concerns than a banner with low personalization depth, i.e., a positive effect of personalization depth on privacy concerns.

Moreover, for a given level of personalization depth, personalization breadth should be a further indicator for the amount of information that the advertising retailer acquired about the consumer. Continuing with the above example, if a high-depth banner had narrow breadth to feature only one of the inspected cars, the signaled knowledge about the consumer's browsing activities would be rather limited. A match between banner content and previous behavior could even be regarded coincidental. However, if the banner had wide breadth to exhaustively feature all inspected cars, a very profound knowledge of the consumer's activities would be revealed. We therefore expect the positive effect of personalization depth on privacy concerns to be even stronger for banners with wide breadth.

Summary of expected baseline effects of ad personalization

As outlined, we argue the impact of personalized banners on consumers' internal responses to be determined by the ads' underlying configurations of personalization depth and breadth. However, the previous discussion also reveals that these two dimensions affect consumers' internal responses differently. In particular, while personalization depth generally determines the main direction of the effect that ad personalization exerts on perceived usefulness, reactance, and privacy concerns, personalization breadth strengthens or weakens its effect at a given level of personalization depth. This reflects the fact that personalization depth, as the primary dimension of ad personalization, captures a banner's general closeness to a consumer's inferred interests and preferences, while personalization breadth, as the secondary dimension of ad personalization, depicts how completely a banner reflects preferences at a given personalization depth. Table 2 summarizes our expectations regarding the baseline effects of ad personalization with different depth/breadth configurations on consumers' internal responses.

Table 2
General expectations for baseline effects of personalized ads with different depth/breadth combinations compared to ads with low personalization depth.

Internal response	Depth	
	High	
	Breadth	
	Wide	Narrow
Usefulness	+	++
Reactance	+	++
Privacy concerns	++	+

Note: + and ++ represent respective medium or strong expected changes of the respective internal responses to ads with different depth/breadth combinations compared to ads with low personalization depth. Breadth only applies to banners with heightened personalization depth, since retargeting banners at the very low end of the depth continuum do not feature any products at all (Lambrecht and Tucker 2013).

We next derive specific hypotheses how trust in the retailer moderates these suggested baseline effects of ad personalization on the three internal responses, respectively. Note that for each response measure, we develop and test our moderation hypotheses for both configurations of personalization depth and breadth, i.e., high depth/narrow breadth and high depth/wide breadth. We do this because, as argued, the baseline effects can differ between these configurations, so that they are possibly also more or less sensitive to a moderation of trust.

Hypothesis Development: Moderating Effect of Trust

Every social or business interaction involves uncertainty about the behavior of the counterpart. The impossibility to fully predict or control in which actions the other will engage presents situations of considerable complexity that may inhibit certain behaviors in fear of potential ramifications (Gefen 2000). Trust is one of the most effective methods to reduce this complexity, especially in circumstances that are not completely governed by rules and regulations (Fukuyama 1995; Luhmann 1979; Olivero and Lunt 2004). While previous research has conceptualized trust in a variety of ways (see Gefen, Karahanna, and Straub 2003 for a comprehensive review), a well-accepted definition views it as a psychological state that comprises the intention to accept vulnerability based on positive expectations of the intentions or behaviors of another party (Rousseau et al. 1998). These positive expectations allow discarding possibly unfavorable behaviors of the other and thus reduce a situation's complexity (Gefen 2000). In particular, the marketing literature defines trust as the "willingness to rely on an exchange partner in whom one has confidence" (Moorman, Zaltman, and Deshpandé 1992, p. 315), where confidence primarily stems from the truster's belief in the trustee's reliability and integrity (Morgan and Hunt 1994). Trust is a crucial antecedent for participation in traditional commerce, but an even more important factor in e-commerce and the online environment where risks and uncertainties about retailers' true intentions are especially hard to assess (Gefen 2000; Hoffman, Novak, and Peralta 1999; Reichheld and Scheffer 2000). When further information about online retailers is lacking, trust serves as the key foundation on which consumers base their research and purchase decisions (McStay 2011; Urban, Amyx, and Lorenzon 2009; Urban, Sultan, and Qualls 2000).

Importantly, research finds trust to influence how useful consumers perceive retailers' marketing efforts. For instance, consumers regard retail websites as more useful when they have trust in the particular merchant, because they can rely on its expertise and be confident of its integrity and refraining from detrimental behavior (Gefen, Karahanna, and Straub 2003). Similarly, trust increases the perceived usefulness of recommendation agents, because consumers can rule out doubts about the agents' true abilities and intentions (Wang and Benbasat 2005). Trust thereby effectively reduces the situational complexity so that consumers can focus on evaluating a given recommendation without questioning the agent's expertise in its domain. In our similar advertising context, we therefore expect our reasoning for the baseline effect of ad personalization on a banner's perceived usefulness to hold when consumers trust

the banner-issuing retailer. For the respective configurations of personalization depth and breadth we thus postulate:

H1a. For a more trusted retailer, personalizing banners with high depth and narrow breadth increases their perceived usefulness compared to low-depth banners.

H1b. For a more trusted retailer, personalizing banners with high depth and wide breadth increases their perceived usefulness compared to low-depth banners.

In contrast, consumers find it harder to accomplish their product search and purchase tasks with a less trusted retailer (Gefen, Karahanna, and Straub 2003). They are also less inclined to accept advice from a recommendation agent that they do not trust well (Wang and Benbasat 2005). Without trust reducing situational complexity we do not expect consumers to focus on evaluating a banner's content so that our previous reasoning should not hold. We thus postulate:

H1c. For a less trusted retailer, personalizing banners with high depth and narrow breadth does not increase their perceived usefulness compared to low-depth banners.

H1d. For a less trusted retailer, personalizing banners with high depth and wide breadth does not increase their perceived usefulness compared to low-depth banners.

Trust might also moderate the suggested unfavorable internal responses to ad personalization. As argued, ad personalization likely triggers reactance primarily because confronting consumers with ads that feature select products a retailer deems relevant to them leads consumers to perceive their free choice processes intruded upon or restricted. Yet, by reducing situational complexity trust has generally been shown to propel consumers' likelihood to accept guidance (White 2005). Also, previous research shows that consumers value their freedom of choice not very highly when receiving advice from a trusted expert (Clee and Wicklund 1980) and are more likely to accept advice from a trusted online retailer (McKnight, Choudhury, and Kacmar 2002b). A trusted retailer's specific ad recommendation should therefore not lead consumers to experience increased reactance. We therefore postulate for the two configurations of personalization depth and breadth:

H2a. For a more trusted retailer, personalizing banners with high depth and narrow breadth does not increase reactance compared to low-depth banners.

H2b. For a more trusted retailer, personalizing banners with high depth and wide breadth does not increase reactance compared to low-depth banners.

Under low trust, however, consumers are more likely to suspect an advisor to predict their behavior in order to maximize its own profit and not to help them with their choice task. They also value their freedom to choose independently from that advisor's recommendations higher when trust in it is low (Clee and Wicklund 1980). We therefore expect our reasoning for the corresponding baseline effect to hold and consumers to experience increased reactance from banners with high personalization depth and narrow or wide breadth:

H2c. For a less trusted retailer, personalizing banners with high depth and narrow breadth increases reactance compared to low-depth banners.

H2d. For a less trusted retailer, personalizing banners with high depth and wide breadth increases reactance compared to low-depth banners.

Last, ad personalization causes privacy concerns primarily because it confronts consumers with their vulnerability to the data-collecting and advertising retailer. However, previous research argues that trust increases one's willingness to accept vulnerability based on positive expectations of the intentions or behaviors of another party (Rousseau et al. 1998). Similarly do trust building factors such as familiarity and positive past experiences with a firm assure consumers that the collection and usage of their personal information occurs under fair practices (Chellappa and Sin 2005). When consumers trust a retailer, they feel their collected personal data to be safe with the firm and believe it to be used in ethically correct ways (Brown and Muchira 2004; Taylor, Davis, and Jillapalli 2009). In high-trust situations consumers are therefore increasingly inclined to disclose personal information (Metzger 2004; Reichheld and Scheffer 2000). In sum, this implies that for a more trusted retailer consumers should not experience heightened privacy concerns from ad personalization as argued in our baseline effect argumentation. We therefore postulate:

H3a. For a more trusted retailer, personalizing banners with high depth and narrow breadth does not increase privacy concerns compared to low-depth banners.

H3b. For a more trusted retailer, personalizing banners with high depth and wide breadth does not increase privacy concerns compared to low-depth banners.

Under low trust, however, consumers are less inclined to accept vulnerability to a retailer (Mayer, Davis, and Schoorman 1995; Rousseau et al. 1998). In particular, they likely question whether the retailer will safeguard their personal information and use it under fair and ethically correct premises (Brown and Muchira 2004). For a less trusted retailer we therefore expect our reasoning for the ad personalization baseline effect to hold. We thus postulate:

H3c. For a less trusted retailer, personalizing banners with high depth and narrow breadth increases privacy concerns compared to low-depth banners.

H3d. For a less trusted retailer, personalizing banners with high depth and wide breadth increases privacy concerns compared to low-depth banners.

Experimental Lab Study

Experimental Design and Procedure

To test our hypotheses, we conducted a scenario-based online experiment which is a common approach in this domain (Fisher and Dubé 2005; Mitra and Lynch 1995; White et al. 2008). For

a relatively less trusted compared to a more trusted retailer, we investigated the proposed effects of ad personalization with different combinations of personalization depth and breadth. In line with our hypotheses, we first differentiated between low and high depth. In the case of high depth we further distinguished between narrow and wide breadth. As explained, breadth only applies to banners of heightened depth, because banners at the very low end of the personalization depth continuum are generic and contain no specific products. Thus, interests cannot be displayed more or less completely in low-depth ads. The resulting experimental setup was a 2 (more trusted vs. less trusted retailer) $\times$ 3 (high depth and narrow breadth vs. high depth and wide breadth vs. low depth) between-subjects design. After pre-testing the experiment we made slight adjustments to the scenario descriptions to ensure clarity.

The experiment consisted of 5 sections: (1) First, participants read a short introduction asking them to imagine they were looking to buy a compact camera at the online store of a retailer named Digicam.³ Depending on the intended trust manipulation, participants then received specific background information that described their prior experiences and interactions with Digicam. In particular, the wording of these descriptions reflected the three key trusting beliefs of competence, benevolence and integrity, as investigated by previous literature (McKnight, Choudhury, and Kacmar 2002a), in a positive or negative manner (for complete wording, see Appendix A). (2) Next, in a created Digicam online store participants had to select three digital cameras that they liked best from the retailer's assortment and place them in their virtual wish list. This interactive element also helped to increase the realism of the experiment (Eisenbeiss et al. 2014). (3) Subsequently, a cover story informed participants that they left the online store without purchasing any camera. (4) Participants were then shown a news website that they allegedly visited on the following day. Here, they encountered a banner ad from the retailer whose online store they had previously visited. Depending on the treatment group, this banner featured either one of the three cameras that were placed in the wish list (high depth and narrow breadth), all three cameras from the wish list (high depth and wide breadth), or just the retailer's logo (low depth). For screenshots of the banners, see Appendix B. (5) Last, participants filled out a questionnaire. Its measures included evaluations of the banner's usefulness as well as participants' elicited reactance, privacy concerns, and click-through intentions.

Manipulation Checks

Post-experimental manipulation checks confirmed that trust in the retailer, averaged across the four items of Verhoef, Franses, and Hoekstra (2002) (see pilot study) on a seven-point scale from 1 ("Strongly disagree") to 7 ("Strongly agree"), was significantly higher in the more trusted retailer scenarios compared to the less trusted retailer scenarios ($\Delta M = 3.090$, $t = 20.680$, $p < .001$).

The manipulation of personalization depth was assessed with two items from Kalyanaram and Sundar (2006) on a seven-point scale from 1 ("Strongly disagree") to 7 ("Strongly agree"): "The content and information featured in the banner targeted me as a unique individual," and "this banner was 'personalized' according to my interests." The corresponding ratings were significantly higher in the high-depth treatment groups, compared to the low-depth group ($\Delta M = 2.354$, $t = 9.050$, $p < .001$). Last, the manipulation of personalization breadth within the high-depth conditions was assessed with two items from Ozcan and Sheinin (2012) on a seven-point scale from 1 ("Strongly disagree") to 7 ("Strongly agree"): "This banner contains all the products that I need," and "this banner is 'complete' in terms of products." Banners with high depth and wide breadth received significantly higher ratings than banners with high depth and narrow breadth ($\Delta M = 1.684$, $t = 8.270$, $p < .001$).

Sample

The final experiment was conducted with German university students who were randomly assigned to one of the experimental groups. 304 participants completed the survey. Before analyzing the data, we discarded the records from 24 participants who spent less than 6 min to complete the entire experiment (average completion time was 12.15 min). Moreover, we removed data from 28 participants who did not pass an attention check. In this check we asked participants to tick a specific answer box in order to determine whether they really answered questions with the necessary diligence. These two methods are frequently used to detect and remove participants that make insufficient effort to answer questions in survey research (Huang et al. 2012; Stenstrom and Curtis 2012). After this data cleansing, 252 valid questionnaires remained ($M_{age} = 21$, 56% female).

Measures

We measured usefulness with four items adapted from Tam and Ho (2006): "The banner enables me to accomplish shopping tasks more quickly," "the banner improves my shopping task performance," "the banner enhances my effectiveness in my shopping tasks," and "the banner makes it easier to do my shopping tasks." The corresponding seven-point rating scale ranged from 1 ("Strongly disagree") to 7 ("Strongly agree"). Alpha reliability was .91.

Reactance was measured with seven items adapted from Edwards, Li, and Lee (2002) and Hong and Faedda (1996). Participants indicated how strongly they agreed the banner to be "disturbing," "interfering," "intrusive," "forced upon them," and "unwelcomed". Moreover, they indicated to which extent they agreed wanting to "resist" and "dismiss" its content. The corresponding seven-point rating scale ranged from 1 ("Strongly disagree") to 7 ("Strongly agree"). Alpha reliability was .88.

We measured privacy concerns with four items according to Sheng, Nah, and Siau (2008), whose scale is based on Smith, Milberg, and Burke (1996) and Dinev and Hart (2004): "It bothers me that the firm is able to track information about me," "I am concerned that the firm has too much information about me," "it

³ We used a fictitious retailer name to increase the realism of the experiment and ensure that participants had no predetermined opinions about it.

Table 3
Measurement model fit indices.

Construct/item	More trusted retailer				Less trusted retailer			
	SL	<i>t</i> -value	CR	AVE	SL	<i>t</i> -value	CR	AVE
Usefulness			.748	.715			.772	.746
Use1	.827	22.407			.830	19.459		
Use2	.824	21.807			.823	16.532		
Use3	.813	17.897			.846	21.332		
Use4	.914	37.581			.949	39.324		
Reactance			.648	.590			.631	.569
R1	.911	38.930			.909	48.621		
R2	.803	20.978			.788	22.437		
R3	.855	26.775			.871	29.331		
R4	.773	16.980			.749	15.176		
R5	.859	30.315			.842	26.468		
R6	.535	8.556			.511	9.170		
R7	.553	8.814			.497	8.286		
Privacy concerns			.764	.735			.777	.751
PC1	.855	19.430			.870	26.649		
PC2	.863	27.940			.838	22.668		
PC3	.877	29.295			.918	41.214		
PC4	.834	26.581			.838	20.017		
CT intention			1.000	1.000			1.000	1.000
CT1	1.000	999.000			1.000	999.000		

Note: AVE, average variance extracted; CR, composite reliability; CT, click-through; SL, standardized factor loadings.

Table 4
AVE and squared correlations.

	More trusted retailer				Less trusted retailer			
	Usefulness	Reactance	Privacy concerns	CT intention	Usefulness	Reactance	Privacy concerns	CT intention
Usefulness	(.772)				(.715)			
Reactance	.116	(.570)			.005	(.590)		
Privacy concerns	.011	.154	(.751)		.000	.012	(.735)	
CT intention	.279	.231	.062	(1)	.144	.082	.066	(1)

Note: Average variance extracted (AVE) in brackets; CT, click-through.

bothers me that the firm is able to access information about me,” and “I am concerned that my information could be used in ways I could not foresee.” Items were rated on a seven-point scale ranging from 1 (“Strongly disagree”) to 7 (“Strongly agree”). Alpha reliability was .92.

We measured click-through intentions with a single-item according to Yoo (2007): “I would like to click-through the banner ad to acquire further information.” Participants responded on a seven-point scale ranging from 1 (“Strongly disagree”) to 7 (“Strongly agree”).⁴

All measurement constructs met convergent validity requirements, as Table 3 shows. For each construct, the composite-based reliabilities well exceeded the required threshold of 0.6 (Bagozzi and Yi 1988). The same applied to the average variance extracted (AVE) with the required threshold of 0.5 (Fornell and Larcker 1981). Finally, *t*-values greater than 1.645 for every item indicated that their factor loadings were significantly different from zero (Bagozzi, Yi, and Phillips 1991).

Discriminant validity was present, as Table 4 shows. Accordingly, the AVE for every latent construct was greater than the squared correlation of the construct with any other construct (Fornell and Larcker 1981).

To minimize the possibility of common method variance (CMV) and response biases we adopted several remedies as indicated by Podsakoff et al. (2003). In particular, we stated in the introduction to each scenario that participants’ responses would be treated anonymously, kept confidential, and be used only for academic purposes. In addition, we highlighted that there were no right or wrong answers and that we were merely interested in participants’ subjective impressions. Moreover, we ensured all items to be as unambiguous as possible, for instance by avoiding double-barreled questions, or complicated syntax. Since we manipulated the independent variables of the stimulus–organism part of our model, CMV poses no problem here. To ensure this empirically as well for the organism–response relationship, we compared the model fit of our measurement model with a single-factor confirmatory factor analysis (Podsakoff et al. 2003). A Satorra–Bentler scaled chi-square difference test indicated a significantly worse model fit of the single-factor model compared with the multifactor model in both groups (more trusted

⁴ The means and standard deviations of all measures appear in Appendix C.

Table 5
Estimated coefficients for the structural model.

Relationship	Path	(1) More trusted retailer				(2) Less trusted retailer			
		Coeff	Estimate	t-value	p	Coeff	t-value	p	
(A) Stimulus → Organism	NB → Usefulness	γ_{11}	.727	2.068	.039	γ_{21}	-.152	-.515	.606
	WB → Usefulness	γ_{12}	.259	.815	.415	γ_{22}	.314	.894	.371
	NB → Reactance	γ_{13}	.614	1.598	.110	γ_{23}	.909	2.612	.009
	WB → Reactance	γ_{14}	.666	1.861	.063	γ_{24}	.910	2.476	.013
	NB → Privacy concerns	γ_{15}	.226	.562	.574	γ_{25}	1.216	3.699	.000
	WB → Privacy concerns	γ_{16}	1.252	3.776	.000	γ_{26}	1.475	4.408	.000
(B) Organism → Response	Reactance → Usefulness	γ_{31}	-.316	-5.176	.000	γ_{31}	-.316	-5.176	.000
	Privacy concerns → Usefulness	γ_{32}	.033	.452	.651	γ_{32}	.033	.452	.651
	Usefulness → CT intention	γ_{33}	.428	5.579	.000	γ_{33}	.428	5.579	.000
	Reactance → CT intention	γ_{34}	-.215	-2.893	.004	γ_{34}	-.215	-2.893	.004
	Privacy concerns → CT intention	γ_{35}	-.137	-1.649	.099	γ_{35}	-.137	-1.649	.099
	Reactance ↔ Privacy concerns	γ_{36}	.800	4.186	.000	γ_{36}	.800	4.186	.000

Note: CT, click-through; NB, high depth and narrow breadth of personalization; WB, high depth and wide breadth of personalization.

retailer group: $TRd^5 = 275.545$, $\Delta df = 5$; less trusted retailer group: $TRd = 222.349$, $\Delta df = 5$) so that CMV was disproved (Korsgaard and Roberson 1995; Mossholder et al. 1998).

Last, as a precondition for comparing mean differences between groups, we tested whether full or partial metric and scalar invariance existed (Steenkamp and Baumgartner 1998). First, a confirmatory factor analysis with unconstrained loadings and intercepts for configural invariance had a good model fit ($\chi^2 = 285.479_{(200 \text{ df})}$, $p < .001$; SCF = 1.106; CFI = 0.964; RMSEA = 0.058; SRMR = 0.072). Here, within the *reactance* construct, we allowed the residuals of items R5 and R7 to covary. This adjustment was indicated by corresponding modification indices and reasonable due to the items' similar German meaning. Second, imposing constraints on the factor loadings across groups did not decrease model fit significantly which proved full metric invariance ($\chi^2 = 294.985_{(212 \text{ df})}$, $p < .001$, $TRd = 8.452$, $\Delta df = 12$; SCF = 1.098; CFI = 0.965; RMSEA = 0.056; SRMR = 0.073). Finally, full scalar invariance was present, as, compared to the second model, the fit of a model with all intercepts constrained across groups did not decrease significantly ($\chi^2 = 314.773_{(224 \text{ df})}$, $p < .001$, $TRd = 20.120$, $\Delta df = 12$; SCF = 1.092; CFI = 0.962; RMSEA = 0.057; SRMR = 0.074). Latent means were therefore comparable across groups.

Results

We applied a multigroup structural equation model (SEM) with a group code approach (Aiken, Stein, and Bentler 1994; Jöreskog 1971) to analyze the data. On the one hand, this model allows estimating group-specific effects within an SEM framework. In our setting, we are interested in comparing the effects of different configurations of personalization depth and breadth on consumer responses between the two trust conditions. Thus, we used the groups arising from the 2-way trust manipulation (more


trusted vs. less trusted retailer) to form two analysis groups. On the other hand, the modeling approach also allows estimating treatment effects within an SEM by including the treatment conditions as separate dummy variables into the model. Thus, we coded the treatments of the 3-way personalization manipulation (high depth and narrow breadth vs. high depth and wide breadth vs. low depth) into two separate dummy variables: The variable *NB* equals 1 in the high-depth/narrow-breadth condition and 0 otherwise. The variable *WB* equals 1 in the high-depth/wide-breadth condition. The low-depth condition therefore forms the reference condition.

In the first part of the model, representing the stimulus–organism relationship, we estimated group-specific effects (i.e., separate effects for the more trusted and less trusted retailer conditions, respectively) of the two binary coded personalization variables to determine the effects of high-depth/narrow-breadth banners (*NB*) and high-depth/wide-breadth banners (*WB*) compared to low-depth banners (reference condition). In particular, the estimated coefficient of the variable *NB* indicates the impact of a banner with high depth and narrow breadth of personalization on usefulness, reactance, and privacy concerns compared to the reference group with low depth of personalization. Correspondingly, the estimated coefficient of the variable *WB* indicates these impacts for banners with high depth and wide breadth of personalization compared to the reference group with low personalization depth. The results for this first part of the model appear in Panel (A) of Table 5. Column (1) contains the coefficients for the more trusted retailer and Column (2) shows the coefficients for the less trusted retailer.

In the second part of the structural model, representing the organism–response relationship, we held all parameters equal across groups, since trust was not expected to moderate effects here. Panel (B) of Table 5 shows the estimated results for the impact of usefulness, reactance, and privacy concerns on click-through intentions. We also control for possible interactions between the three internal responses. The structural model has a good fit ($\chi^2 = 349.032_{(266 \text{ df})}$, $p < .001$; CFI = 0.967; RMSEA = 0.050; SRMR = 0.071), allowing for the assessment of our conceptual model and the pertaining hypotheses.

⁵ TRd = Satorra–Bentler scaled chi-square difference test statistic. Throughout this study we used robust maximum likelihood methods which require model fit comparisons to be based on the TRd (Satorra and Bentler 2001).

(a) Structural model for a more trusted retailer.



(b) Structural model for a less trusted retailer.




Fig. 3. Structural models. *Note:* The respective upper path coefficients pertain to the effects of NB (high-depth/narrow-breadth banners compared to low-depth banners) on usefulness, reactance, and privacy concerns while the lower coefficients pertain to respective effects of WB (high-depth/wide-breadth banners compared to low-depth banners) on these measures. * $p < .1$; ** $p < .05$; *** $p < .001$.

Fig. 3 depicts the graphical illustrations for the path coefficients of the structural model. For presentation purposes we present results separately for the more trusted and the less trusted retailer in Fig. 3(a) and (b). For the stimulus–organism part of the model, the respective upper estimates pertain to the effects of NB (high-depth and narrow-breadth banners compared to low-depth banners) on usefulness, reactance, and privacy concerns while the lower estimates pertain to respective effects of WB (high-depth and wide-breadth banners compared to low-depth banners) on these measures.

In the following, we test our hypotheses to determine the extent to which trust in the retailer moderates the effects of ad personalization on perceived usefulness, reactance and privacy concerns. We also examine how these effects translate into click-through intentions. H1a suggests that for a more trusted retailer,

personalizing banners with high depth and narrow breadth should lead to an increase in perceived usefulness. The results support this hypothesis ($\gamma_{11} = .727$, $t = 2.068$, $p < .05$). Hence, featuring selected items from consumers' initial consideration sets effectively increases banner usefulness for more trusted retailers. By contrast, perceived usefulness does not increase when high depth is combined with wide breadth of personalization ($\gamma_{12} = .259$, $t = .815$, $p > .1$) so that we do not find support for H1b. This result reinforces our general reasoning that personalization breadth influences the impact of banners with a given level of personalization depth on consumers. It moreover underscores our argument that featuring only selected items provides additional value. By contrast, H1c and H1d propose that for a less trusted retailer, ad personalization with high depth and either narrow or wide breadth does not enhance a banner's perceived

usefulness. We find support for both hypotheses ($\gamma_{21} = -.152$, $t = -.515$, $p > .1$; $\gamma_{22} = .314$, $t = .894$, $p > .1$).

In terms of negative responses to ad personalization, **H2a** and **H2b** suggest that for a more trusted retailer, personalizing banners with high depth and either narrow or wide breadth should not increase reactance compared to banners with low depth. While the results confirm this for high-depth/narrow-breadth banners in support of **H2a** ($\gamma_{13} = .614$, $t = 1.598$, $p > .1$), we do find a significant effect for high-depth banners with wide breadth on reactance ($\gamma_{14} = .666$, $t = 1.861$, $p < .1$). We therefore cannot confirm **H2b**. Surprisingly, this finding is opposite to our baseline assumption that narrow breadth rather than wide breadth should add to the effects of high depth on reactance (see **Table 2**). Note, however, that the estimated coefficient is only marginally significant and close to the estimate of the insignificant coefficient for a combination of high depth and narrow breadth (.614 vs. .666). The differences should therefore not be over-interpreted.⁶ Instead, it seems that reactance is generally less likely to be triggered by ad personalization from a more trusted retailer.

For a less trusted retailer, **H2c** suggests that banners with high depth and narrow breadth should increase reactance compared to banners with low depth. The results support this hypothesis ($\gamma_{23} = .909$, $t = 2.612$, $p < .05$). Moreover, we also find support for **H2d** which proposes banners with high depth and wide breadth to cause increased reactance ($\gamma_{24} = .910$, $t = 2.476$, $p < .05$). Thus, for a less trusted retailer, attempting to very accurately predict consumers' inferred interests and preferences is indeed perceived as freedom threatening.

As a further negative internal response we examine privacy concerns. **Tables 3 and 4** clearly support our argument that privacy concerns are a conceptually distinct response, but nevertheless related to reactance ($\gamma_{36} = .800$, $t = 4.186$, $p < .001$). **H3a** proposes that, for a more trusted retailer, ad personalization with high depth and narrow breadth should not increase privacy concerns compared to low depth banners. We find empirical support for this argumentation ($\gamma_{15} = .226$, $t = .562$, $p > .1$). However, the data do not support **H3b** which suggests that ad personalization with high depth and wide breadth should not increase privacy concerns compared to low depth for a more trusted

retailer ($\gamma_{16} = 1.252$, $t = 3.776$, $p < .001$). Finding that ad personalization can even increase privacy concerns for well-trusted retailers is interesting and in line with our baseline assumption that wide breadth should add to the effects of high personalization depth compared to low-depth banners on privacy concerns (see **Table 2**). Hence, it appears that consumers are willing to accept vulnerability to a more trusted retailer without experiencing privacy concerns, but only up to a certain extent. When banners with high depth and wide breadth signal that the retailer has very profound knowledge of their previous shopping behavior, consumers experience privacy concerns even in a trusting relationship. Moreover, since this configuration is not associated with increased usefulness, consumers might also perceive the retailer to collect too much information for too little value in return which then does not offset privacy concerns, akin to the widely accepted privacy calculus (Dinev and Hart 2006).

By contrast, **H3c** and **H3d** suggest that, for a less trusted retailer, personalizing banners with high depth and either narrow or wide breadth should increase privacy concerns compared to low depth banners. The results confirm both hypotheses ($\gamma_{25} = 1.216$, $t = 3.699$, $p < .001$; $\gamma_{26} = 1.475$, $t = 4.408$, $p < .001$). Thus, irrespective of personalization breadth is a high-depth banner's signal that a retailer increasingly tracks and analyzes consumers' private information enough to elicit heavy privacy concerns when the retailer is less trusted.

According to the S-O-R paradigm, consumers' internal responses eventually translate into external responses (Donovan and Rossiter 1982; Mehrabian and Russell 1974). We captured these in terms of click-through intentions and find all three internal responses, i.e., usefulness, reactance and privacy concerns, to significantly affect this measure. In particular, and as to be expected, perceived usefulness exerts a positive impact on click-through intentions ($\gamma_{33} = .428$, $t = 5.579$, $p < .001$), while reactance ($\gamma_{34} = -.215$, $t = -2.893$, $p < .05$) and privacy concerns ($\gamma_{35} = -.137$, $t = -1.649$, $p < .1$) have a negative impact on them.

Overall, the results from the lab confirm our findings from the field that trust is a major determinant of the effectiveness of ad personalization through retargeting. Moreover, they show that the depth and breadth of personalization are two important dimensions that influence consumers' internal and external responses.

Discussion

In order to enhance the effectiveness of their online advertising retailers increasingly employ retargeting to tailor banners to individual consumers (Helft and Vega 2010; Peterson 2013; Sengupta 2013). Nevertheless, ad personalization has proven to be a delicate matter whose usage comes at the risk of triggering unfavorable responses (Guild 2013; Tucker 2012b; Turow et al. 2009). In this research we use data from a quasi-experimental field study and a lab experiment to investigate whether and to which extent trust in the retailer moderates the impact of ad personalization in retargeting on consumers.

Altogether, our results suggest that trust is an important moderator of substantial impact. In particular, for the more trusted

⁶ Yet, a possible explanation for this empirical finding might lie in the negative relationship between reactance and usefulness ($\gamma_{31} = -.316$, $t = -5.176$, $p < .001$). For more trusted retailers, consumers perceive high-depth banners with narrow breadth to be more useful and these ads also do not cause reactance. Potentially, consumers accept and even appreciate a banner's reduction of their initial consideration sets, weighing its usefulness in the current purchase decision heavier than a constraint of their freedom of choice. This finding aligns with previous research showing that consumers' compliance can outweigh opposition when ad messages are consistent with preferences and the sender is regarded competent (Clee and Wicklund 1980; Fitzsimons and Lehmann 2004). Moreover, it underscores findings that consumers are more likely to accept advice from benevolent sources and do not experience their freedom of choice threatened when they are accustomed to trust a sender's judgment (Clee and Wicklund 1980; Jodlbauer and Jonas, 2011). By contrast, consumers do not regard banners with wide personalization breadth as more useful. Here, it appears that they do not perceive them helpful enough to offset their freedom-threatening appearance.

retailer in our field study, we find banner click-through rates to increase by 27% (from .40% to .51%) as a result of employing a higher depth of personalization, i.e., tailoring banners more closely to consumers' interests and preferences inferred from their previous online shopping behaviors. By contrast, a higher depth of personalization actually decreases click-through rates for banners of the less trusted retailer by 46% (from .37% to .20%). To profoundly understand the mechanisms that drive these outcomes we conducted a follow-up lab experiment. In this study we uncover the processes through which personalization influences consumers' internal and external ad responses. We also introduce and investigate personalization breadth as a secondary dimension of ad personalization that determines how completely, or exhaustively, banners with a given personalization depth reflect consumers' preferences. The lab results show that, for banners of more trusted retailers, click-through intentions are particularly high when ads reflect a combination of high depth and narrow breadth of personalization. This occurs not only because consumers perceive these banners more useful than those with low depth, but also because these ads do not elicit increased reactance or privacy concerns compared to low-depth banners. By contrast, for banners of less trusted retailers, a higher depth of personalization decreases consumers' click-through intentions, irrespective of banners' personalization breadth. Consumers do not only fail to view these ads as more useful than low-depth banners, they also experience increased reactance and privacy concerns when encountering them. These findings provide clear implications for practice and important contributions for theory.

First and foremost, our results imply that retailers should carefully assess consumers' trust in them before deciding how to personalize banners to their inferred interests and preferences. To this end, retailers might evaluate their corporate standing in the marketplace relative to their competitors by estimating general consumer trust through market research. At best, however, they should implement ways to determine trust at the individual consumer level, since trust perceptions typically vary across consumers. For this, retailers may analyze whether specific consumers are willing to provide certain personal information during transactions or when browsing the retailer's online store (Phelps, Nowak, and Ferrell 2000; Schoenbachler and Gordon 2002). Based on these insights, more trusted retailers should then serve banners with a combination of high depth and narrow breadth of personalization to consumers that appear to have high trust in them. Also, retailers may incorporate these individual trust assessments into their real-time bidding strategies (Angwin 2012) to help determine bid values for ad impressions to be served to particular consumers.

Yet, while trust plays a crucial role in governing consumers' ad responses, being trusted is at best a necessary, but not a sufficient condition for retailers to ultimately benefit from ad personalization through retargeting. For example, our results reveal that even ads from more trusted retailers can elicit privacy concerns under a combination of high depth and wide breadth of personalization. Hence, besides determining optimal configurations of personalization depth and breadth for their banners,

even trusted retailers should take additional steps to prevent consumers from experiencing privacy concerns from specific ad encounters. Amazon is a prime example of a retailer that conducts such additional measures. Although the retailer is highly trusted among consumers (OC&C 2013), it provides recipients of retargeting ads with detailed information about which specific personal data it stores and uses to personalize its ads. Moreover, the firm fully discloses its privacy policy and offers clear instructions how consumers can either opt-out from being served personalized banner ads or improve the experience of personalized ads and recommendations (Amazon 2015). This way it maintains transparency and provides consumers with control about the collection and usage of their personal data which, in line with previous research (Dinev and Hart 2004; Sheehan and Hoy 2000; Tucker 2014), is an effective means to reduce consumers' privacy concerns and reactance in online advertising.

Moreover, our results further imply that less trusted retailers should refrain from closely tailoring their banners to consumers. These insights appear particularly relevant considering that retargeting is not only a popular tool among incumbent retailers such as Amazon, but also among small retailers and early-stage start-ups. Due to low marketing budgets, many of these firms focus on retargeting as a primary marketing instrument to bring visitors that did not yet convert back to their websites (500 Startups 2013). Since consumers are often not very familiar with these retailers, they likely have not yet formed trust in them. Importantly, our findings therefore not only apply to retailers that have a low standing in consumers' eyes, but especially to those about whom consumers have not yet built a clear opinion.

Besides these practical implications, from an academic perspective, our work also contributes to theory. Previous research has thus far regarded ad personalization as a one-dimensional construct where approaches only vary with respect to their personalization depth. In this research, we introduce the breadth of personalization as a second key dimension of personalized online advertising. We demonstrate its relevance to govern the effectiveness of banners with a given level of personalization depth and thereby contribute to a more holistic understanding of ad personalization.

Moreover, we provide a greater understanding of the internal psychological responses that mediate the impact of a personalized banner stimulus on consumers' behavioral responses. In particular, we identify perceived usefulness, reactance, and privacy concerns as distinct but interrelated constructs. Differentiating these internal responses is of high relevance for research on personalized online advertising. Specific ad personalization methods might create advertisements that elicit reactance because consumers perceive them as manipulative. Yet, these methods may at the same time not necessarily cause increased privacy concerns. In reverse, other ads may not cause reactance, but trigger heavy privacy concerns for instance by displaying very intimate consumer information. While our work advances current knowledge about the impact of ad personalization in retargeting on consumers, there are also limitations to be addressed that facilitate avenues for future research.

Limitations and Future Research

First, as previously indicated, our field data regarding the effects of ad personalization on click-through is quasi-experimental. While this should not influence the between-retailer comparison we conduct for banners with high and low personalization depth, the absolute click-through rates might be positively biased. To determine these rates more accurately, future research could replicate our study in a completely randomized field experiment. Such an experiment could also manipulate personalization breadth in addition to personalization depth. Moreover, while we control for a number of firm characteristics, the retailers in our field data can potentially differ on more dimensions than trust alone. A replication study could help rule out alternative explanations by incorporating trust assessments at the individual consumer level for instance through the measures we outlined above.

Next, the lab experiment provides only a static snapshot of reality. It confronts each participant with only a single banner while firms serve consumers usually multiple ad impressions over week-long or month-long time frames. Research shows that banner influences on consumers dynamically change with increasing impression amounts (Braun and Moe 2013). Future work might account for these effects and investigate to which extent our findings prove more differentiated in such settings. Also, we conducted our lab study with German participants. While the impact of cultural characteristics on the perception of personalized advertising has not yet been investigated, results might differ for participants from other countries or cultural groups, based on findings from traditional advertising research (e.g., Han and Shavitt 1994; Zhang and Gelb 1996). Studying how cultural factors might influence the effectiveness of ad personalization might also be worthwhile as the Internet

enables firms to easily reach consumers from various backgrounds worldwide.

Besides tackling these methodological limitations, future research might also expand our work conceptually. First, we focus on the personalization of banners based on individual consumers' shopping behaviors in retailers' online stores. Yet, this poses only one source of information firms can use to personalize their ads. In practice they also commonly incorporate knowledge about consumers' search queries (Rutz and Bucklin 2011), emails (Gibbs 2014), or social media activities (Tucker 2012a). As these data points represent information of specific sensitivities, the impact of trust on consumer responses might vary for ads that are personalized based on different information types (Nowak and Phelps 1992), another avenue for future research.

Moreover, while our finding that trust in the retailer strongly moderates the impact of personalized advertisements on consumers is an important contribution to theory and practice, further factors that we do not consider in this research such as situational characteristics (e.g., browsing modes, browsing motives) or consumer characteristics (e.g., age, personality traits, shopping habits) might be in play. Future research might investigate such moderators to shed additional light on the factors that drive the effectiveness of personalized online advertising.

Finally, given the importance of trust on the Web, building trust is crucial for retailers in e-commerce (Kim and Kim 2011). While previous studies have investigated different approaches to build and regain trust (e.g., Ba and Pavlou 2002; Matzat and Snijders 2012; Urban, Sultan, and Qualls 2000), choosing appropriate communication approaches, such as avoiding ads with high personalization depth in certain situations, is likely an important factor in building or regaining trust. Future research in this area can build on our results.

Appendix A. Introductory descriptions for a more/less trusted retailer in section 1 of the lab experiment

Trusting beliefs (McKnight, Choudhury, and Kacmar 2002a)	More trusted retailer	Less trusted retailer
<u>Competence</u>	Imagine you want to buy a digital camera. To this end you visit the online store of the retailer Digicam.de, known for its large assortment. During previous shopping sessions <u>the different search and comparison options that Digicam.de offers in its online store worked very well for you so that you always felt well-assisted and helped with your purchase decisions. Payment processes were always smooth and ordered goods arrived a</u>	Imagine you want to buy a digital camera. To this end you visit the online store of the retailer Digicam.de, known for its large assortment. During previous shopping sessions <u>the different search and comparison options that Digicam.de offers in its online store did not work well for you so that you oftentimes did not feel well-assisted or helped with your purchase decisions. Also, payment processes were seldom smooth and</u>
<u>Integrity</u>	<u>few days after you placed an order. You therefore regard Digicam.de as an honest retailer of integrity. Apparently, Digicam.de wants to</u>	<u>ordered goods often arrived with considerable delay. You therefore regard Digicam.de as a rather dishonest retailer of low integrity. Apparently, Digicam.de only wants to achieve the</u>
<u>Benevolence</u>	<u>achieve the best outcome for you as a customer. In addition to these positive aspects, you consider Digicam.de because of its large assortment.</u>	<u>best outcome for itself. Yet, despite these negative aspects, you still consider Digicam.de because of its large assortment.</u>

Appendix B. Examples of banners shown in section 4 of the lab experiment


High depth and narrow breadth



High depth and wide breadth



Low depth



Appendix C. Mean (Std Dev) results across experimental groups

Ad personalization	Item	More trusted retailer (N = 129)		Less trusted retailer (N = 123)	
		Mean	Std Dev	Mean	Std Dev
High depth/narrow breadth	Use1: The banner enables me to accomplish shopping tasks more quickly	3.2821	1.8202	2.1714	1.4849
	Use2: The banner improves my shopping task performance	2.6154	1.4621	1.6571	0.9983
	Use3: The banner enhances my effectiveness in my shopping tasks	2.7179	1.6214	1.9429	1.2589
	Use4: The banner makes it easier to do my shopping tasks	2.7692	1.4772	1.8000	1.1324
	R1: The banner is disturbing	4.5897	1.7729	4.7895	1.6134
	R2: The banner is interfering	5.4103	1.7580	5.6053	1.4054
	R3: The banner is intrusive	4.6923	1.8660	5.1579	1.5162
	R4: The banner is forced upon me	4.4615	1.9980	4.9474	1.8737
	R5: The banner is unwelcomed	5.0256	1.8135	5.7632	1.3443
	R6: I want to resist the banner	3.3590	1.7089	4.4737	1.8561
	R7: I want to dismiss the content of this banner	3.4615	1.6519	4.4474	1.8113
	PC1: It bothers me that the firm is able to track information about me	4.6923	2.0794	5.6053	1.2848
	PC2: I am concerned that the firm has too much information about me	4.5385	2.0755	5.4211	1.4262
	PC3: It bothers me that the firm is able to access information about me	4.2821	2.1020	5.5789	1.3682
	PC4: I am concerned that my information could be used in ways I could not foresee	4.5385	1.9980	5.5789	1.3076
High depth/wide breadth	CT: I would like to click-through the banner ad to acquire further information	2.7714	1.8956	1.8286	1.0977
	Use1: The banner enables me to accomplish shopping tasks more quickly	2.7647	1.8503	2.3947	1.6852
	Use2: The banner improves my shopping task performance	2.2353	1.5045	2.2368	1.6347
	Use3: The banner enhances my effectiveness in my shopping tasks	2.3922	1.4571	2.2368	1.6013
	Use4: The banner makes it easier to do my shopping tasks	2.3529	1.4398	2.2895	1.5923
	R1: The banner is disturbing	4.5686	1.9210	4.8158	1.8723
	R2: The banner is interfering	5.7059	1.6161	5.9737	1.3653
	R3: The banner is intrusive	4.4510	2.0718	4.9737	1.6520
	R4: The banner is forced upon me	5.2353	1.8065	5.4474	1.5717
	R5: The banner is unwelcomed	4.8235	1.7743	5.3947	1.7170
	R6: I want to resist the banner	3.7451	1.9884	4.5263	1.9135
	R7: I want to dismiss the content of this banner	3.6471	1.8091	4.4737	2.0761
	PC1: It bothers me that the firm is able to track information about me	5.6078	1.6623	6.0789	1.1942
	PC2: I am concerned that the firm has too much information about me	5.2745	1.7445	5.4474	1.7037
	PC3: It bothers me that the firm is able to access information about me	5.4902	1.5540	5.6842	1.3970
	PC4: I am concerned that my information could be used in ways I could not foresee	5.8039	1.4563	6.1053	1.0601
	CT: I would like to click-through the banner ad to acquire further information	2.3404	1.6717	2.1316	1.6793

Ad personalization	Item	More trusted retailer (N = 129)		Less trusted retailer (N = 123)	
		Mean	Std Dev	Mean	Std Dev
Low depth	Use1: The banner enables me to accomplish shopping tasks more quickly	2.5641	1.8609	2.5652	1.7720
	Use2: The banner improves my shopping task performance	2.4103	1.7125	2.0652	1.3889
	Use3: The banner enhances my effectiveness in my shopping tasks	2.1795	1.5021	2.2174	1.4593
	Use4: The banner makes it easier to do my shopping tasks	2.2821	1.6535	2.1957	1.4700
	R1: The banner is disturbing	3.7436	1.5848	4.0213	1.9052
	R2: The banner is interfering	5.1282	1.8235	4.7021	2.0314
	R3: The banner is intrusive	3.8974	1.7591	4.2128	1.9442
	R4: The banner is forced upon me	4.4359	1.7591	4.2979	2.0947
	R5: The banner is unwelcomed	4.4103	1.9428	4.7872	1.9553
	R6: I want to resist the banner	3.6410	1.6934	3.7872	1.6932
	R7: I want to dismiss the content of this banner	3.5897	1.4994	4.0638	1.7371
	PC1: It bothers me that the firm is able to track information about me	4.3077	1.9487	4.3617	1.9828
	PC2: I am concerned that the firm has too much information about me	4.3590	1.8423	4.1702	1.9596
	PC3: It bothers me that the firm is able to access information about me	3.8974	1.9028	4.2766	2.0184
	PC4: I am concerned that my information could be used in ways I could not foresee	4.6154	1.9278	4.5106	2.1353
	CT: I would like to click-through the banner ad to acquire further information	2.4722	1.7806	2.2609	1.7312

Note: Construct “usefulness” consists of the items Use1–Use4; construct “reactance” consists of the items R1–R7; construct “privacy concerns” consists of the items PC1–PC4; construct “click-through intentions” consists of the item CT.

References

- 500 Startups (2013), “Distribution Tuesday: Two Quick Ways to Get More Conversion with Retargeting,” (accessed January 15, 2015), [available at <http://500.co/get-more-conversions-with-retargeting-two-quick-ways-to-start>].
- 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.
- Aiken, Leona S., Judith A. Stein and Peter M. Bentler (1994), “Structural Equation Analyses of Clinical Subpopulation Differences and Comparative Treatment Outcomes: Characterizing the Daily Lives of Drug Addicts,” *Journal of Consulting and Clinical Psychology*, 62 (3), 488–99.
- Amazon (2015), “Amazon Advertising Preferences,” (accessed February 24, 2015), [available at <http://www.amazon.com/gp/dra/info>].
- Anand, Bharat N. and Ron Shachar (2009), “Targeted Advertising as a Signal,” *Quantitative Marketing and Economics*, 7 (3), 237–66.
- Angwin, Julia (2012), “Online Tracking Ramps Up,” *Wall Street Journal*, June 17 (accessed February 25, 2014), [available at <http://www.wsj.com/articles/SB10001424052702303836404577472491637833420>].
- Ansari, Asim and Carl F. Mela (2003), “E-Customization,” *Journal of Marketing Research*, 60 (2), 131–45.
- Ba, Sulin and Paul A. Pavlou (2002), “Evidence of the Effect of Trust Building Technology in Electronic Markets: Price Premiums and Buyer Behavior,” *MIS Quarterly*, 26 (3), 243–68.
- Bagozzi, Richard P. and Youjae Yi (1988), “On the Evaluation of Structural Equation Models,” *Journal of the Academy of Marketing Science*, 16 (1), 74–94.
- Bagozzi, Richard P., Youjae Yi and Lynn W. Phillips (1991), “Assessing Construct Validity in Organizational Research,” *Administrative Science Quarterly*, 36 (3), 421–58.
- Bandyopadhyay, Soumava (2011), “Antecedents and Consequences of Consumers’ Online Privacy Concerns,” *Journal of Business & Economics Research*, 7 (3), 41–8.
- Bleier, Alexander and Maik Eisenbeiss (2015), “Personalized Online Advertising Effectiveness: The Interplay of What, When, and Where,” *Marketing Science*, (forthcoming).
- Braun, Michael and Wendy W. Moe (2013), “Online Display Advertising: Modeling the Effects of Multiple Creatives and Individual Impression Histories,” *Marketing Science*, 32 (5), 753–67.
- Brehm, Jack W. (1966), *A Theory of Psychological Reactance*, New York: Academic Press.
- (1989), “Psychological Reactance: Theory and Applications,” *Advances in Consumer Research*, 16 (1), 72–5.
- Brehm, Sharon S. and Jack W. Brehm (1981), *Psychological Reactance: A Theory of Freedom and Control*, New York: Academic Press.
- Brown, Mark and Rose Muchira (2004), “Investigating the Relationship between Internet Privacy Concerns and Online Purchase Behavior,” *Journal of Electronic Commerce Research*, 5 (1), 62–70.
- Bucklin, Randolph E. and Catarina Sismeiro (2009), “Click Here for Internet Insight: Advances in Clickstream Data Analysis in Marketing,” *Journal of Interactive Marketing*, 23 (1), 35–48.
- Chellappa, Ramnath K. and Raymond G. Sin (2005), “Personalization versus Privacy: An Empirical Examination of the Online Consumer’s Dilemma,” *Information Technology and Management*, 6 (2/3), 181–202.
- Chen, Ye, Dmitry Pavlov and John F. Canny (2009), “Large-Scale Behavioral Targeting,” in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 209–18.
- Chernev, Alexander (2003), “When More Is Less and Less Is More: The Role of Ideal Point Availability and Assortment in Consumer Choice,” *Journal of Consumer Research*, 30 (2), 170–83.
- Cho, Chang-Hoan and Hongsik J. Cheon (2004), “Why Do People Avoid Advertising on the Internet?,” *Journal of Advertising*, 33 (4), 89–97.
- Clee, Mona A. and Robert A. Wicklund (1980), “Consumer Behavior and Psychological Reactance,” *Journal of Consumer Research*, 6 (4), 389–405.
- Conti, Nicolette, Charlene Jennett, Jose Maestre and M. Sasse (2012), “When Did My Mobile Turn into A ‘Sellphone’?: A Study of Consumer Responses to Tailored Smartphone Ads,” in *Proceedings of the 26th Annual BCS Interaction Specialist Group Conference on People and Computers*, 215–20.
- Criteo (2015), “All Testimonials,” (accessed February 26, 2015), [available at <http://www.criteo.com/case-studies/>].
- Dinev, Tamara and Paul Hart (2004), “Internet Privacy Concerns and Their Antecedents—Measurement Validity and a Regression Model,” *Behaviour & Information Technology*, 23 (6), 413–22.
- and ——— (2006), “An Extended Privacy Calculus Model for E-Commerce Transactions,” *Information Systems Research*, 17 (1), 61–80.
- Donovan, Robert J. and John R. Rossiter (1982), “Store Atmosphere: An Environmental Psychology Approach,” *Journal of Retailing*, 58 (1), 34–57.

- Edwards, Steven, Hairong Li and Joo-Hyun Lee (2002), "Forced Exposure and Psychological Reactance: Antecedents and Consequences of the Perceived Intrusiveness of Pop-Up Ads," *Journal of Advertising*, 31 (3), 83–95.
- Eisenbeiss, Maik, Markus Cornelißen, Klaus Backhaus and Wayne D. Hoyer (2014), "Nonlinear and Asymmetric Returns on Customer Satisfaction: Do They Vary Across Situations and Consumers?," *Journal of the Academy of Marketing Science*, 42 (3), 242–63.
- eMarketer (2014a), "Marketers Respond to Consumer Demand for Personalization," (accessed February 2, 2015), [available at <http://www.emarketer.com/Article/Marketers-Respond-Consumer-Demand-Personalization/1010618>].
- eMarketer (2014b), "US Retail Industry Maintains Position as Digital Ad Spend Leader," (accessed January 27, 2015), [available at <http://www.emarketer.com/Article/US-Retail-Industry-Maintains-Position-Digital-Ad-Spend-Leader/1010871>].
- Fisher, Robert J. and Laurette Dubé (2005), "Gender Differences in Responses to Emotional Advertising: A Social Desirability Perspective," *Journal of Consumer Research*, 31 (4), 850–8.
- Fitzsimons, Gavan J. and Donald R. Lehmann (2004), "Reactance to Recommendations: When Unsolicited Advice Yields Contrary Responses," *Marketing Science*, 23 (1), 82–94.
- Fornell, Claes and David F. Larcker (1981), "Evaluating Structural Equation Models with Unobservable Variables and Measurement Error," *Journal of Marketing Research*, 18, 39–50.
- Foxman, Ellen R. and Paula Kilcoyne (1993), "Information Technology, Marketing Practice, and Consumer Privacy: Ethical Issues," *Journal of Public Policy & Marketing*, 12 (1), 106–19.
- Fukuyama, F. (1995), *Trust: The Social Virtues and the Creation of Prosperity*, New York: The Free Press.
- Gefen, David (2000), "E-commerce: The Role of Familiarity and Trust," *Omega*, 28 (6), 725–37.
- Gefen, David, Elena Karahanna and Detmar W. Straub (2003), "Trust and TAM in Online Shopping: An Integrated Model," *MIS Quarterly*, 27 (1), 51–90.
- Gibbs, Samuel (2014), "Gmail Does Scan All Emails, New Google Terms Clarify," *The Guardian*, April 15 (accessed February 18), [available at <http://www.theguardian.com/technology/2014/apr/15/gmail-scans-all-emails-new-google-terms-clarify>].
- Godfrey, Andrea, Kathleen Seiders and Glenn B. Voss (2011), "Enough Is Enough! The Fine Line in Executing Multichannel Relational Communication," *Journal of Marketing*, 75 (4), 94–109.
- Goldfarb, Avi and Catherine Tucker (2011), "Online Display Advertising: Targeting and Obtrusiveness," *Marketing Science*, 30 (3), 389–404.
- Goodstadt, Michael S. (1971), "Helping and Refusing to Help: A Test of Balance and Reactance Theories," *Journal of Experimental Social Psychology*, 7 (6), 610–22.
- Grandpre, Joseph, Eusebio M. Alvaro, Michael Burgoon, Claude H. Miller and John R. Hall (2003), "Adolescent Reactance and Anti-Smoking Campaigns: A Theoretical Approach," *Health Communication*, 15 (3), 349–66.
- Guild, Bill (2013), "2013 Choicestream Survey: Consumer Opinions on Online Advertising & Audience Targeting," (accessed February 18, 2015), [available at http://choicestream.com/2013_Staging/wp-content/uploads/2013/10/2013-Survey.pdf].
- Han, Sang-Pil and Sharon Shavitt (1994), "Persuasion and Culture: Advertising Appeals in Individualistic and Collectivistic Societies," *Journal of Experimental Social Psychology*, 30 (4), 326–50.
- Hannah, T. Edward, Elena R. Hannah and Barbara Wattie (1975), "Arousal of Reactance as a Consequence of Predicting an Individual's Behaviour," *Psychological Reports*, 37 (2), 411–20.
- Häubl, Gerald and Valerie Trifts (2000), "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids," *Marketing Science*, 19 (1), 4–21.
- Helft, Miguel and Tanzina Vega (2010), "Retargeting Ads Follow Surfers to Other Sites," *New York Times*, (August 30), A1.
- Hoeffler, Steve and Dan Ariely (1999), "Constructing Stable Preferences: A Look into Dimensions of Experience and their Impact on Preference Stability," *Journal of Consumer Psychology*, 8 (2), 113–39.
- Hoffman, Donna L., Thomas P. Novak and Marcos Peralta (1999), "Building Consumer Trust Online," *Communications of the ACM*, 42 (4), 80–5.
- Hong, Sung-Mook and Salvatora Faedda (1996), "Refinement of the Hong Psychological Reactance Scale," *Educational and Psychological Measurement*, 56 (1), 173–82.
- Huang, Jason L., Paul G. Curran, Jessica Keeney, Elizabeth M. Poposki and Richard P. DeShon (2012), "Detecting and Detering Insufficient Effort Responding to Surveys," *Journal of Business and Psychology*, 27 (1), 99–114.
- Hunter, Anne, Meredith Jacobsen, Richard Talens, and Tony Winders (2010), "When Money Moves to Digital, Where Should it Go? Identifying the Right Media-Placement Strategies for Digital Display," whitepaper, ComScore, September 21 (accessed March 3, 2014), [available at <http://www.iab.net/media/file/MediaPlacementWhitepaper.pdf>].
- Iyengar, Sheena S. and Mark R. Lepper (2000), "When Choice is Demotivating: Can One Desire Too Much of a Good Thing?," *Journal of Personality and Social Psychology*, 79 (6), 995–1006.
- Jai, Tun-Min, Leslie D. Burns and Nancy J. King (2013), "The Effect of Behavioral Tracking Practices on Consumers' Shopping Evaluations and Repurchase Intention Toward Trusted Online Retailers," *Computers in Human Behavior*, 29 (3), 901–9.
- Jodlbauer, Barbara and Eva Jonas (2011), "Forecasting Clients' Reactions: How Does the Perception of Strategic Behavior Influence the Acceptance of Advice?," *International Journal of Forecasting*, 27 (1), 121–33.
- Jöreskog, Karl G. (1971), "Simultaneous Factor Analysis in Several Populations," *Psychometrika*, 36 (4), 409–26.
- Kahn, Arnold and Thomas E. Tice (1973), "Returning a Favor and Retaliating Harm: The Effects of Stated Intentions and Actual Behavior," *Journal of Experimental Social Psychology*, 9 (1), 43–56.
- Kalyanaraman, Sriram and S. Shyam Sundar (2006), "The Psychological Appeal of Personalized Content in Web Portals: Does Customization Affect Attitudes and Behavior?," *Journal of Communication*, 56 (1), 110–32.
- Kim, Kyongseok and Jooyoung Kim (2011), "Third-Party Privacy Certification as an Online Advertising Strategy: An Investigation of the Factors Affecting the Relationship between Third-Party Certification and Initial Trust," *Journal of Interactive Marketing*, 25 (3), 145–58.
- King, Nancy J. and Pernille W. Jessen (2010), "Profiling the Mobile Customer – Is Industry Self-Regulation Adequate to Protect Consumer Privacy when Behavioural Advertisers Target Mobile Phones? – Part II," *Computer Law & Security Review*, 26 (6), 595–612.
- Kivetz, Ran (2005), "Promotion Reactance: The Role of Effort-Reward Congruity," *Journal of Consumer Research*, 31 (4), 725–36.
- Kobsa, Alfred (2007), "Privacy-Enhanced Personalization," *Communications of the ACM*, 50 (8), 24–33.
- Korsgaard, M. Audrey and Loriann Roberson (1995), "Procedural Justice in Performance Evaluation: The Role of Instrumental and Non-Instrumental Voice in Performance-Appraisal Discussions," *Journal of Management*, 21 (4), 657–69.
- Kramer, Thomas and Ryall Carroll (2009), "The Effect of Incidental Out-Of-Stock Options on Preferences," *Marketing Letters*, 20 (2), 197–208.
- Kreuter, Matthew W. and Ricardo J. Wray (2003), "Tailored and Targeted Health Communication: Strategies for Enhancing Information Relevance," *American Journal of Health Behavior*, 27 (3), 227–32.
- Lambrecht, Anja and Catherine E. Tucker (2013), "When Does Retargeting Work? Information Specificity in Online Advertising," *Journal of Marketing Research*, 50 (5), 561–76.
- Lee, Gyudong and Won Jun Lee (2009), "Psychological Reactance to Online Recommendation Services," *Information & Management*, 46 (8), 448–52.
- Li, Han, Rathindra Sarathy and Heng Xu (2011), "The Role of Affect and Cognition on Online Consumers' Decision to Disclose Personal Information to Unfamiliar Online Vendors," *Decision Support Systems*, 51 (3), 434–45.
- Luhmann, Niklas (1979), *Trust and Power: Two Works by Niklas Luhmann*, Chichester: Wiley.
- Malhotra, Naresh K., Sung S. Kim and James Agarwal (2004), "Internet Users' Information Privacy Concerns (IUPC): The Construct, the Scale, and a Causal Model," *Information Systems Research*, 15 (4), 336–55.

- Matzat, Uwe and Chris Snijders (2012), "Rebuilding Trust in Online Shops on Consumer Review Sites: Sellers' Responses to User-Generated Complaints," *Journal of Computer-Mediated Communication*, 18 (1), 62–79.
- Mayer, Roger C., James H. Davis and F. David Schoorman (1995), "An Integrative Model of Organizational Trust," *Academy of Management Review*, 20 (3), 709–34.
- Mazis, Michael B., Robert B. Settle and Dennis C. Leslie (1973), "Elimination of Phosphate Detergents and Psychological Reactance," *Journal of Marketing Research*, 10 (4), 390–5.
- McKnight, D. Harrison, Vivek Choudhury and Charles Kacmar (2002a), "Developing and Validating Trust Measures for E-Commerce: An Integrative Typology," *Information Systems Research*, 13 (3), 334–59.
- , ——— and ——— (2002b), "The Impact of Initial Consumer Trust on Intentions to Transact with a Web Site: A Trust Building Model," *Journal of Strategic Information Systems*, 11 (3), 297–323.
- McStay, Andrew (2011), *The Mood of Information: A Critique of Online Behavioural Advertising*, New York: Continuum.
- Mehrabian, Albert and James A. Russell (1974), *An Approach to Environmental Psychology*, Cambridge, MA: MIT Press.
- Metzger, Miriam J. (2004), "Privacy, Trust, and Disclosure: Exploring Barriers to Electronic Commerce," *Journal of Computer-Mediated Communication*, 9 (4),
- Miller, Claude H., Lindsay T. Lane, Leslie M. Deatrick, Alice M. Young and Kimberly A. Potts (2007), "Psychological Reactance and Promotional Health Messages: The Effects of Controlling Language, Lexical Concrete-ness, and the Restoration of Freedom," *Human Communication Research*, 33 (2), 219–40.
- Mitra, Anusree and John G. Lynch Jr. (1995), "Toward a Reconciliation of Market Power and Information Theories of Advertising Effects on Price Elasticity," *Journal of Consumer Research*, 21 (4), 644–59.
- Moorman, Christine, Gerald Zaltman and Rohit Deshpandé (1992), "Relationships Between Providers and Users of Market Research: The Dynamics of Trust," *Journal of Marketing*, 29 (3), 314–28.
- Morgan, Robert M. and Shelby D. Hunt (1994), "The Commitment-Trust Theory of Relationship Marketing," *Journal of Marketing*, 58 (3), 20–38.
- Mossholder, Kevin W., Nathan Bennett, Edward R. Kemery and Mark A. Wesolowski (1998), "Relationships Between Bases of Power and Work Reactions: The Mediation Role of Procedural Justice," *Journal of Management*, 24 (4), 533–52.
- Murthi, B.P.S. and Sumit Sarkar (2003), "The Role of the Management Sciences in Research on Personalization," *Management Science*, 49 (10), 1344–62.
- Nowak, Glen J. and Joseph Phelps (1992), "Understanding Privacy Concerns: An Assessment of Consumers' Information-Related Knowledge and Beliefs," *Journal of Direct Marketing*, 6 (4), 28–39.
- Nysveen, Herbjørn and Per E. Pedersen (2004), "An Exploratory Study of Customers' Perception of Company Web Sites Offering Various Interactive Applications: Moderating Effects of Customers' Internet Experience," *Decision Support Systems*, 37 (1), 137–50.
- OC&C (2013), "Would You Trust a Machine? The OC&C Retail Proposition Index 2013," (accessed February 10, 2015), [available at <http://www.occstrategy.com/insights/vwxyz/would-you-trust-a-machine>].
- Okazaki, Shintaro, Hairong Li and Morikazu Hirose (2009), "Consumer Privacy Concerns and Preference for Degree of Regulatory Control," *Journal of Advertising*, 38 (4), 63–77.
- Olivero, Nadia and Peter Lunt (2004), "Privacy versus Willingness to Disclose in E-Commerce Exchanges: The Effect of Risk Awareness on the Relative Role of Trust and Control," *Journal of Economic Psychology*, 25 (2), 243–62.
- Ozcan, Timucin and Daniel A. Sheinin (2012), "Effects of Complete Products on Consumer Judgments," *Journal of Product & Brand Management*, 21 (4), 246–54.
- Peterson, Tim (2013), "EBay Opens Up its Data for Ad Targeting—Follows Lead of Amazon, Google and Facebook," *Adweek*, (April 8).
- 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.
- Podsakoff, Philip M., Scott B. MacKenzie, Jeong-Yeon Lee and Nathan P. Podsakoff (2003), "Common Method Bias in Behavioral Research: A Critical Review of the Literature and Recommended Remedies," *Journal of Applied Psychology*, 88 (5), 879–903.
- PricewaterhouseCoopers (2010), "Measuring the Effectiveness of Online Advertising," research report, PricewaterhouseCoopers LLP.
- Raab, Charles D. and Colin J. Bennett (1998), "The Distribution of Privacy Risks: Who Needs Protection?," *The Information Society*, 14 (4), 263–74.
- Rapp, Justine, Ronald P. Hill, Jeannie Gaines and R. Mark Wilson (2009), "Advertising and Consumer Privacy," *Journal of Advertising*, 38 (4), 51–61.
- Reichheld, Frederick F. and Phil Scheffer (2000), "E-Loyalty," *Harvard Business Review*, 78 (4), 105–13.
- Rousseau, Denise M., Sim B. Sitkin, Ronald S. Burt and Colin Camerer (1998), "Not So Different After All: A Cross-Discipline View of Trust," *Academy of Management Review*, 23 (3), 393–404.
- Rutz, Oliver J. and Randolph E. Bucklin (2011), "From Generic to Branded: A Model of Spillover in Paid Search Advertising," *Journal of Marketing Research*, 48 (1), 87–102.
- Satorra, Albert and Peter M. Bentler (2001), "A Scaled Difference Chi-Square Test Statistic for Moment Structure Analysis," *Psychometrika*, 66 (4), 507–14.
- Schoenbachler, Denise D. and Geoffrey L. Gordon (2002), "Trust and Customer Willingness to Provide Information in Database-Driven Relationship Marketing," *Journal of Interactive Marketing*, 16 (3), 2–16.
- Sengupta, Somini (2013), "What You Didn't Post, Facebook May Still Know," *New York Times*, March 25 (accessed November 17, 2014), [available at <http://www.nytimes.com/2013/03/26/technology/facebook-expands-targeted-advertising-through-outside-data-sources.html?pagewanted=all>].
- Shavitt, Sharon, Pamela Lowrey and James Haefner (1998), "Public Attitudes Toward Advertising: More Favorable Than You Might Think," *Journal of Advertising Research*, 38 (4), 7–22.
- Sheehan, Kim B. and Marica G. Hoy (2000), "Dimensions of Privacy Concern among Online Consumers," *Journal of Public Policy & Marketing*, 19 (1), 62–73.
- Sheng, Hong, Fiona F.-H. Nah and Keng Siau (2008), "An Experimental Study on Ubiquitous Commerce Adoption: Impact of Personalization and Privacy Concerns," *Journal of the Association for Information Systems*, 9 (6), 344–76.
- Sizmek (2014), "Sizmek Benchmarks Report 2014," (accessed February 10, 2015), [available at http://www.sizmek.com/assets/pdf/Sizmek_2014-03_Benchmarks_report.pdf].
- Smith, H. Jeff, Sandra J. Milberg and Sandra J. Burke (1996), "Information Privacy: Measuring Individuals' Concerns about Corporate Practices," *MIS Quarterly*, 20 (2), 167–96.
- Steenkamp, Jan-Benedict E.M. and Hans Baumgartner (1998), "Assessing Measurement Invariance in Cross-National Consumer Research," *Journal of Consumer Research*, 25 (1), 78–90.
- Stenstrom, Douglas M. and Mathew Curtis (2012), "Heroism and Risk of Harm," *Psychology*, 3 (12), 1085–90.
- Tam, Kar Y. and Shuk Y. Ho (2005), "Web Personalization as a Persuasion Strategy: An Elaboration Likelihood Model Perspective," *Information Systems Research*, 16 (3), 271–91.
- and ——— (2006), "Understanding the Impact of Web Personalization on User Information Processing and Decision Outcomes," *MIS Quarterly*, 30 (4), 865–90.
- Taylor, David G., Donna F. Davis and Ravi Jillapalli (2009), "Privacy Concern and Online Personalization: The Moderating Effect of Information Control and Compensation," *Business and Economics*, 9 (3), 203–23.
- Tucker, Catherine E. (2012a), "Social Advertising," working paper, MIT.
- (2012b), "The Economics of Advertising and Privacy," *International Journal of Industrial Organization*, 30 (3), 326–9.
- Tucker, Catherine E. (2014), "Social Networks, Personalized Advertising, and Privacy Controls," *Journal of Marketing Research*, 51 (5), 546–62.
- Turov, Joseph, Jennifer King, Chris J. Hoofnagle, Amy Bleakley, and Michael Hennessy (2009), "Americans Reject Tailored Advertising and Three Activities that Enable It," technical report, Annenberg School for Communication, University of Pennsylvania.

- Urban, Glenn L., Fareena Sultan and William J. Qualls (2000), "Placing Trust at the Center of Your Internet Strategy," *MIT Sloan Management Review*, 39–48.
- Urban, Glenn L., Cinda Amyx and Antonio Lorenzon (2009), "Online Trust: State of the Art, New Frontiers, and Research Potential," *Journal of Interactive Marketing*, 23 (2), 179–90.
- Van Doorn, Jenny and Janny C. Hoekstra (2013), "Customization of Online Advertising: The Role of Intrusiveness," *Marketing Letters*, 24 (4), 1–13.
- Verhoef, Peter C., Philip H. Franses and Janny Hoekstra (2002), "The Effect of Relational Constructs on Customer Referrals and Number of Services Purchased From a Multiservice Provider: Does Age of Relationship Matter?," *Journal of the Academy of Marketing Science*, 30 (3), 202–16.
- Wang, Weiquan and Izak Benbasat (2005), "Trust in and Adoption of Online Recommendation Agents," *Journal of the Association for Information Systems*, 6 (3), 72–100.
- Westin, Alan F. (1967), *Privacy and Freedom*, New York, NY: Atheneum.
- White, Tiffany B. (2005), "Consumer Trust and Advice Acceptance: The Moderating Roles of Benevolence, Expertise, and Negative Emotions," *Journal of Consumer Psychology*, 15 (2), 141–8.
- White, Tiffany B., Debra Zahay, Helge Thorbjørnsen and Sharon Shavitt (2008), "Getting too Personal: Reactance to Highly Personalized Email Solicitations," *Marketing Letters*, 19 (1), 39–50.
- Winer, Russell S. (2001), "A Framework for Customer Relationship Management," *California Management Review*, 43 (4), 89–105.
- Worchel, Stephen and Jack W. Brehm (1971), "Direct and Implied Social Restoration of Freedom," *Journal of Personality and Social Psychology*, 18 (3), 294–304.
- Xiao, Bo and Izak Benbasat (2007), "E-Commerce Product Recommendation Agents: Use, Characteristics, and Impact," *MIS Quarterly*, 31 (1), 137–209.
- Yoo, Chan Y. (2007), "Implicit Memory Measures for Web Advertising Effectiveness," *Journalism & Mass Communication Quarterly*, 84 (1), 7–23.
- Zhang, Yong and Betsy D. Gelb (1996), "Matching Advertising Appeals to Culture: The Influence of Products' Use Conditions," *Journal of Advertising*, 25 (3), 29–46.
- Zhang, Harry, Claudia Guerrero, David Wheatley and Young Seok Lee (2010), "Privacy Issues and User Attitudes towards Targeted Advertising: A Focus Group Study," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 54 (19), 1416–20.