

The Bulletproof Glass Effect: Unintended Consequences of Privacy Notices

Aaron R. Brough , David A. Norton, Shannon L. Sciarappa, and Leslie K. John

Journal of Marketing Research
 2022, Vol. 59(4) 739-754
 © American Marketing Association 2022
 Article reuse guidelines:
sagepub.com/journals-permissions
 DOI: 10.1177/00222437211069093
journals.sagepub.com/home/mrj



Abstract

Drawing from a content analysis of publicly traded companies' privacy notices, a survey of managers, a field study, and five online experiments, this research investigates how consumers respond to privacy notices. A privacy notice, by placing legally enforceable limits on a firm's data practices, communicating safeguards, and signaling transparency, might be expected to promote confidence that personal data will not be misused. Indeed, most managers expected a privacy notice to make customers feel more secure (Study 1). Yet, consistent with the analogy that bulletproof glass can increase feelings of vulnerability despite the protection offered, formal privacy notices undermined consumer trust and decreased purchase interest even when they emphasized objective protection (Studies 2, 3, and 5) or omitted any mention of potentially concerning data practices (Study 6). These unintended consequences did not occur, however, when consumers had an a priori reason to be distrustful (Study 4) or when benevolence cues were added to privacy notices (Studies 5 and 6). Finally, Study 7 showed that both the presence and conspicuous absence of privacy information are sufficient to trigger decreased purchase intent. Together, these results provide actionable guidance to managers on how to effectively convey privacy information (without hurting purchase interest).

Keywords

privacy, trust, information disclosure

Online supplement: <https://doi.org/10.1177/00222437211069093>

Consumers regularly encounter privacy notices explaining if and how their personal information will be collected, stored, used, and shared. Although privacy notices are mandated in many industries and locations by law, such as the European Union's General Data Protection Regulation, wide variation exists in the manner and extent to which details about a firm's privacy practices and handling of data are communicated to consumers. For example, some notices include a lengthy description of the company's privacy practices, while others consist of only a brief and often vague statement. Privacy-related information may even be absent (Culnan 2000) or unavailable, such as when a "privacy nutrition label" on Apple's App Store indicates that the developer has not provided details about its data-handling practices (Miller 2021). In this research, we address the question of how consumers respond to such differences in the availability and presentation of privacy-related information.

Privacy notices might be expected to help consumers feel more secure for several reasons. First, privacy notices place legally enforceable limits on how organizations can collect, store, use, and share consumers' personal data. To illustrate, the California Consumer Privacy Act allows consumers to sue companies that fail to fulfill promised privacy protections.

Second, privacy notices often communicate protective measures (e.g., encryption, firewalls) that guard against unauthorized use of consumer information. Third, prior research suggests that transparency in how a firm manages and protects customer data can reduce perceived vulnerability (Martin, Borah, and Palmatier 2017). Thus, by revealing exactly what personal data companies have access to and how it will be processed, managers may expect consumers to be more comfortable with a firm's handling of their data.

In contrast, we propose that privacy notices can, ironically, lead consumers to feel *more* rather than less vulnerable despite the protections they offer. In this sense, a privacy notice may be likened to bulletproof glass, which may increase

Aaron R. Brough is Associate Professor of Marketing and Harry M. Reid Endowed Professor of Research, Jon M. Huntsman School of Business, Utah State University, USA (email: aaron.brough@usu.edu). David A. Norton is Visiting Assistant Professor, The Ohio State University, USA (email: norton.253@osu.edu). Shannon Sciarappa is Research Associate, Harvard Business School, Harvard University, USA. Leslie K. John is Professor of Business Administration, Harvard Business School, Harvard University, USA (email: ljohn@hbs.edu).

feelings of vulnerability despite the protection it provides (particularly when encountered in a context of expected safety, such as an elementary school). If a privacy notice decreases consumers' willingness to trust a company with personal information, purchase interest is likely to decline. Accordingly, we refer to the "bulletproof glass effect" as the decreased purchase interest resulting from exposure to a privacy notice. In the following sections, we review relevant literature and delineate the theoretical basis for our contention that formal privacy notices can reduce trust, and, in turn, purchase interest. We then provide an overview of the studies that test our predictions.

Conceptual Development

Consumer Responses to Privacy-Related Information

Faced with common news reports of identity theft, leaked personal data, and corporate security breaches, it is not surprising that consumers, businesses, and policy makers are concerned with protecting personal information from unauthorized access, collection, storage, use, and sharing (Hazel and Slobogin 2018; Kamleitner et al. 2018; Phelps, Nowak, and Ferrell 2000; White 2004). When consumers realize that personal information has been collected without consent, click-through rates drop (Aguirre et al. 2015; Kim, Barasz, and John 2019), and when provided with privacy ratings for multiple websites, participants avoid purchasing from sites that offer lower levels of privacy protection (Tsai et al. 2011). In short, insufficient control over personal information can decrease consumers' willingness to make a purchase (Phelps, D'Souza, and Nowak 2001).

Given its consequential business implications, privacy has been identified as an area ripe for behavioral research (Brough and Martin 2020; Kim, Barasz, and John 2020; Krishna 2020; Lambertson and Stephen 2016), in part because of the disconnect between what consumers say and do with respect to privacy-related information. Surveys of consumers' attitudes toward privacy protections often produce sensible and predictable results. Such surveys typically ask consumers to indicate, in the abstract, whether they would like firms to present them with privacy policies, to encrypt their data, to offer control over the deletion of personal information, etc. As might be expected, when directly asked, consumers favor restrictions on the gathering and use of personal information (Turow et al. 2012; Westin 1991)—particularly information that is highly sensitive (Milne et al. 2017; Nowak and Phelps 1992). Similarly, consumers say they would be more comfortable with a firm's collection and use of their personal data when fair information practices are promised (Culnan and Armstrong 1999).

In light of these polls, in which consumers generally express preferences for privacy protections, it is reasonable to expect that a privacy notice might mitigate concerns about the potential misuse of personal information. Specifically, by transparently explaining how information will be collected, stored, used, and protected, a privacy notice could build trust and increase

willingness to purchase. In line with this logic, scholars have proposed that instead of treating privacy policies as a compliance cost, managers should approach privacy as an opportunity to give consumers a positive experience with a brand (Goldfarb and Tucker 2013). Of course, privacy notices differ in the level of privacy expectations they create and in the degree of objective protections they afford; some are consumer-protective, describing security measures and highlighting how the collected data will benefit consumers (e.g., through personalization), while others border on the exploitative (essentially giving firms "carte blanche" to do with consumers' data as they will) (Martin 2015; Reidenberg et al. 2016; Zeng et al. 2020).

Taken together, the research discussed above suggests that it would be sensible to expect consumers to be assured by protective privacy notices and alarmed by exploitative ones. By contrast, we posit that even objectively protective privacy notices can undermine, rather than enhance, consumers' trust in a firm. Whereas the results of consumer surveys generally portray a rational response to privacy-related information, consumers' responses to actual exposure to privacy-related information are malleable and less intuitive (Acquisti, Brandimarte, and Loewenstein 2015; Acquisti, John, and Loewenstein 2013; Nissenbaum 2004; Smith, Dinev, and Xu 2011). For example, consumers are quick to abandon privacy-protecting behaviors in response to choice architecture and framing (Adjerid, Acquisti, and Loewenstein 2019; Brandimarte, Acquisti, and Loewenstein 2013), small inconveniences or small incentives (Athey, Catalini, and Tucker 2017), or greater perceived control over personal information (Mourey and Waldman 2020; Tucker 2014).

Given that privacy-related information can have surprising effects on consumer behavior, it would be instructive to know whether privacy notices—either their specific content or their mere presence—affect consumers' purchase interest. Yet scant marketing research exists on this topic. For example, we know of no field study that has manipulated the presence or content of a privacy notice and measured resulting consumer behavior. Therefore, using data from the field and online experiments, we contribute to the privacy literature by examining the impact of exposure to privacy notices on consumer attitudes and behavior. We predict that in some contexts, privacy notices can reduce consumers' trust in a firm, resulting in decreased purchase interest. Next, we delineate the conceptual underpinnings of this prediction.

Privacy Notices, Trust and Purchase Interest

Privacy notices are formal legal contracts—binding agreements that dictate how a firm can collect, use, and store consumers' personal data (Martin 2012). Formal contracts are explicit, rigid, and literal; violations are resolved in the courts and penalized with economic sanctions (Martin 2016). So, it would seem sensible to posit, as privacy scholars have, that formal contract-based approaches to respecting consumer data, such as privacy notices and privacy seals, ought to enhance consumers' comfort in purchasing from a given firm (Martin 2018;

Martin and Murphy 2016; Pan and Zinkhan 2006; Rifon, LaRose, and Choi 2005; Wang, Beatty, and Foxx 2004). Accordingly, it would also seem sensible for managers to expect privacy notices—at least those that offer objective privacy protections—to enhance consumers' feelings of security. Thus, we predict,

H₁: Managers expect privacy notices to make consumers feel more secure.

However, the empirical evidence as to whether privacy seals and other formal contract-based approaches to privacy protection actually foster feelings of security has been mixed (Lauer and Deng 2007; Tang, Hu, and Smith 2008; Xu et al. 2011). Why? A growing body of work characterizes privacy as a social contract (Kim, Barasz, and John 2019; Martin 2012, 2016; Nissenbaum 2004). This perspective asserts that consumers' sense of whether their privacy is being respected or invaded is dictated by norms—consumers' expectations about how their information ought to be handled. These expectations are typically unspoken and implicit and vary across contexts. Firms that honor privacy expectations earn consumers' trust (McCole, Ramsey, and Williams 2010) and enhance purchase interest (Cases et al. 2010; Eastlick, Lotz, and Warrington 2006), whereas those that violate privacy norms suffer consumer backlash, such as reproach and negative word of mouth (Miyazaki 2009). Even when consumers benefit, such as by seeing customized ads for products that they want and need, they tend to react negatively if they perceive that the ads were generated using unsavory methods (Kim, Barasz, and John 2019).

Social contracts are held together by relational concerns; entities adhere to them not out of a desire to avoid legal and economic sanctions but out of a desire to promote harmonic interactions and to avoid social sanctions (Donaldson and Dunfee 1999; Martin 2012, 2016). Thus, social contracts enhance, and are enhanced by, trust (Kim, Barasz, and John 2019; Robinson 1996)—trust being defined as “a psychological state comprising the intention to accept vulnerability based upon positive expectations of the intentions or behavior of another” (Rousseau et al. 1998, p. 395). By contrast, formal contracts can actually *undermine* trust (Malhotra and Murnighan 2002; Martin 2016). Specifically, Malhotra and Murnighan (2002) found that participants who were induced to create formal, binding contracts at the outset of a multiround trust game demonstrated less rather than more trust in subsequent rounds, compared with participants who had not used contracts. Therefore, we surmise that privacy notices, as a kind of formal contract, can undermine trust. In turn, diminished trust has been one of the leading reasons for why some consumers are hesitant to shop online (Hoffman, Novak, and Peralta 1999a, b). Thus, despite managers' expectations to the contrary, we propose that privacy notices will decrease both trust and purchase interest. More formally:

H₂: In contrast to managers' expectations, we predict a bulletproof glass effect in which a salient (vs. absent or less salient) privacy notice decreases purchase interest, even when it emphasizes objective protection or omits any mention of potentially concerning data practices.

H₃: The bulletproof glass effect is mediated by decreased trust.

Prior research suggests that formal contracts may be especially likely to negatively impact trust when such formality is unexpected (Martin 2016; Puranam and Vanneste 2009). As illustrated by our analogy, observing bulletproof glass may have a greater negative impact on perceived security in an environment where safety is expected (e.g., an elementary school) than in an environment expected to be more dangerous (e.g., a prison). Consistent with this idea, potential survey respondents are less willing to complete a survey dealing with nonsensitive topics when provided with elaborate, and presumably unnecessary, assurances of confidentiality (Singer, Hippler, and Schwarz 1992; Singer, Von Thurn, and Miller 1995). Building on the logic that assurances can backfire when people do not already have the potential for harm in mind, we argue that privacy notices may decrease purchase interest when consumers expect safety, but not when consumers are already distrustful. Accordingly, we predict,

H₄: The bulletproof glass effect is likely to be observed when consumers expect safety, but not when consumers have an a priori reason to be distrustful.

The notion that privacy notices erode consumer trust and purchase interest raises an important practical question: how might firms present privacy notices in a way that does not produce these undesired effects? We argue that to avoid undermining purchase interest, privacy information must be communicated in a way that establishes trust. Thus, one potential solution may lie in modifying the written content of the privacy notice to build greater trust. Prior work has identified different components of trust; notably, these include a relational dimension as well as an ability-based dimension (Levin and Cross 2004; Mayer, Davis, and David Schoorman 1995). The former is typically referred to as benevolence-based trust and refers to the consumer's assessment of a firm's motivation to act in the consumer's best interest. The ability component refers to the consumer's assessment of the firm's capacity to execute its promises—for example, to competently encrypt consumer data.

Given our conceptualization of privacy as a social contract, we propose that privacy notices that include benevolence cues (e.g., statements such as “we care about you”) may be more effective at fostering consumer trust, or at least not undermining it, than those that rely only on ability cues (e.g., statements such as “we use 256-bit encryption”). Because benevolence cues appeal to the relational dimension of trust, they may encourage

consumers to view a privacy notice as more of a social than a formal contract. Building on this logic, we predict that incorporating benevolence cues into a privacy notice may mitigate the bulletproof glass effect. However, our intuition was that the legalese predominant in most privacy notices does not tend to foster the kind of relational, benevolence-based trust that underlies effective social contracts.

To assess the extent to which standard privacy notices include benevolence cues, we conducted a pilot study in which we analyzed the privacy notices of 50 publicly traded companies randomly selected from the NASDAQ stock exchange. This methodology ensured that our analysis covered a diverse set of companies, including those of different sizes and from a variety of industries. Relying on prior research showing that benevolence-based trust is affective in nature, whereas ability is a more cognitive dimension of trust (McAllister 1995; Schoorman, Mayer, and Davis 2007), we used the standard dictionaries included in Linguistic Inquiry and Word Count (LIWC) software to score each privacy notice for the proportion of words that reflect affect as well as for the proportion of words that reflect cognitive processes. A paired t-test showed that across all 50 privacy notices, the average cognitive processes score ($M = 16.80$, $SD = 1.84$) was significantly higher than the average affect score ($M = 4.16$, $SD = .98$; $t(49) = 45.21$, $p < .001$), suggesting a greater prevalence of ability (vs. benevolence) cues. In accordance with the findings of this pilot study that few companies currently seem to include benevolence cues in their privacy notices, we propose that adding benevolence cues to a standard privacy notice may attenuate, and possibly even reverse, its negative impact on trust and purchase interest. Thus,

H₅: The bulletproof glass effect is attenuated when a privacy notice incorporates (vs. omits) benevolence cues.

We further argue that the bulletproof glass effect is not limited to situations in which consumers read complete details about a firm's data management practices. Given their familiarity with the legalistic tone that is common among most privacy notices, consumers may respond to the mere concept of a formal contract—whether prompted by the presence or conspicuous absence of a privacy notice—with decreased trust and purchase interest. Thus, even opaque privacy notices that omit detailed descriptions, as well as standardized templates that draw attention to the absence of a privacy notice (e.g., Apple's privacy nutrition labels)—may be sufficient to produce the bulletproof glass effect. Our theoretical model is illustrated in Figure 1.

We tested our hypotheses in a field study as well as in multiple studies with externally valid stimuli and designs that included both attitudinal and behavioral measures. Study 1 tests H₁ by examining managers' intuitions of the effect of a privacy notice on consumer behavior. Study 2 demonstrates the bulletproof glass effect in a field experiment with a financial services firm and tests H₂ by showing that when a privacy notice was made more salient, enrollment rates declined. Study 3 replicates the bulletproof glass effect using both

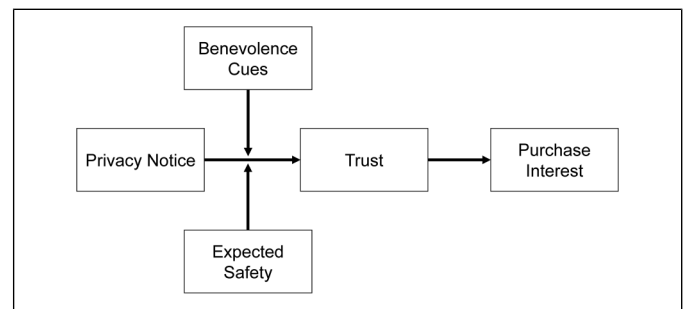


Figure 1. Theoretical model.

attitudinal and behavioral measures in a controlled online experiment and tests H₃, showing that trust mediates the decrease in purchase interest caused by exposure to a privacy notice. Study 4 tests H₄, showing that privacy notices negatively affect purchase interest when consumers expect to feel safe, but not when they are already distrustful, and that this effect is again mediated by trust. Studies 5 and 6 test H₅, showing that the negative effect of a privacy notice on purchase interest is attenuated, and can even be reversed, when it incorporates benevolence cues. Finally, consistent with the idea that the mere concept of a formal privacy notice can decrease trust, Study 7 uses Apple's privacy nutrition labels to show that both the presence and *conspicuous* absence of a privacy notice are sufficient to trigger decreased purchase interest. In all experiments, we preset our sample sizes¹ and/or the time period for data collection. We report all manipulations, measures, and data exclusions. Stimuli for all studies is available in Web Appendix A. Data for all studies are available on Open Science Framework (OSF).²

Study 1: Managers' Intuition

Study 1 tests H₁, that managers will expect privacy notices to make consumers feel more secure.

Method

We recruited 100 participants screened for management experience from Prolific, an online panel provider. Consistent with a preregistration plan (<https://aspredicted.org/blind.php?x=wy2ds9>), we excluded 30 participants who self-reported that they did not have experience working in a management position at their place of employment, leaving a final sample size of 70 participants (28.6% female; mean age = 31.96 years).³ Participants were told,

¹ Variation in sample size across studies is a function of experimental design as well as the time period in which the study was run (studies that were run more recently have larger sample sizes).

² See https://osf.io/7cz3s/?view_only=97d0e1e6f5704557a209fd4b5caa6e6. For the field experiment (Study 2), all available data are reported directly in the article.

³ The pattern and significance of the results do not change when all participants are included.

“Suppose you were working in a management position for an online retailer. Because consumers provide personal information (e.g., their credit card information, address) during the purchase process, your company has a privacy policy that tells consumers how their personal data will be used and protected.” The privacy notice specified practices used to safeguard personal information (e.g., “bank-level encryption”), promised never to share information without consent, and explained how personal information would be used to benefit customers.⁴

After reading the notice, participants were asked: “What, if any effect do you think displaying the privacy notice has on customers? Please select an option.” Participants chose between the following three options: “displaying the privacy notice will make customers feel **more secure**,” “displaying the privacy notice will make customers feel **less secure**,” and “displaying the privacy notice will have **no effect** on how secure customers feel.” Participants then completed demographic measures.

Results and Discussion

Our sample included managers in over 20 different industries with experience in upper, middle, and junior levels of management. In support of H_1 , approximately three in four managers (74.3%; $N = 52/70$) expected that displaying the privacy notice would make customers feel more secure, whereas only 11.4% ($N = 8/70$) expected that it would make customers feel less secure, and the remainder (14.3%; $N = 10/70$) expected that it would have no effect ($\chi^2(2) = 52.91, p < .001, \phi = .87$).

Study 2: Field Experiment

Study 2 was a field experiment designed to test whether (contrary to managers’ expectations in Study 1) a salient privacy notice can diminish consumers’ willingness to transact with a company despite the protections it offers.

Method

We partnered with Borrowell, a Canadian financial technology firm with over a million users. To sign up for Borrowell’s service, visitors must complete a nine-step enrollment process that involves providing sensitive personal information (e.g., name, address, birthdate, phone number, income, financial goals, access to credit report).

The experiment was conducted among 15,864 prospective customers during a seven-day period in May 2019. Each prospective customer who visited the site was randomly assigned to one of two conditions (privacy notice salience: high vs. low). In the low-salience condition, only a hyperlink to Borrowell’s privacy notice was provided on the first screen of

the sign-up process. In the high-salience condition, the link was preceded by an explanation of Borrowell’s commitment to the protection of customers’ personal information. This privacy notice was virtually identical to that used in Study 1 (with the addition of Borrowell’s name). To measure how the salience of the privacy notice impacted interest, we assessed the number of prospective customers who completed the enrollment process.

Results and Discussion

As we predicted, enrollment was significantly lower in the high-salience condition (39.66%; $N = 3,170/7,992$) than in the low-salience condition (41.48%; $N = 3,265/7,872$; $\chi^2(1) = 5.45, p = .020, \phi = .02$). Although this effect is rather small, it is meaningful—that such a subtle manipulation could change enrollment at all in the field is notable. Moreover, at scale, even a small change in enrollment rates can have substantial financial impact. Extrapolating from the seven-day period studied, one would expect roughly 825,000 prospective customers to visit Borrowell’s site annually. With that base, the observed decrease of 1.82% in the enrollment rate would translate to a difference of over 15,000 enrolled customers per year. If average annual revenue per customer were as low as \$15, these results suggest that a salient (vs. less salient) privacy notice could cost Borrowell nearly \$250,000 per year in lost revenue.

In summary, this field experiment provided evidence of the bulletproof glass effect (H_2), showing that prominently displaying detailed privacy protections can drive consumers away. The counterintuitive nature of this result is highlighted by Study 1’s finding that managers expected a nearly identical privacy notice to make customers feel more secure. In the next study, we explore the mechanism for the bulletproof glass effect.

Study 3: Mediation

In Study 3, we tested the hypothesis that the bulletproof glass effect is mediated by decreased trust. In addition to measuring overall trust, we also explored two subdimensions of trust: benevolence-based trust and ability-based trust. Doing so enabled us to explore whether the bulletproof glass effect is robust across different measures of trust.

Method

According to a preregistration plan (<https://aspredicted.org/z7s22.pdf>), we recruited 600 participants (56.3% female; mean age = 38.24 years) on Amazon Mechanical Turk (MTurk). All participants were first shown an identical image and product description. Participants were then randomly assigned to one of two conditions (privacy notice: absent vs. present). In the present condition, participants were asked to review the retailer’s privacy notice. The notice was crafted using language from retailers’ actual privacy notices and explicitly described protective measures such as storing information

⁴ As with most standard privacy notices, this protection was conveyed using language that was more cognitive than affective ($z = 2.1, p = .035$). LIWC scores for all privacy notices throughout the article are reported in Web Appendix B.

in securely encrypted log files, following established identity verification procedures, and adhering to guidelines deemed by the Privacy Shield Program to meet standards prescribed by the Data Privacy Commission.⁵ In the absent condition, participants did not view this notice.

Next, we captured both an attitudinal and a behavioral measure of participants' interest in the product. The attitudinal measure ("How interested would you be in learning more about these sunglasses?") used a sliding scale ranging from 0 = "not at all" to 100 = "extremely." For the behavioral measure, we measured respondents' willingness to spend extra time reading additional product information ("Would you like to see a little more information about these sunglasses?"), with the binary response options being "Yes, please show me a little more information" and "No, I'd like to finish the survey now." Those who selected "yes" were shown additional product information, and the amount of time they spent reading this information was surreptitiously recorded (as was the duration of the entire survey for all participants).

All participants then completed a single-item measure of trust: "For this purchase, how comfortable would you be with the way your data will be collected and stored?," measured using a sliding scale from 0 = "not at all" to 100 = "extremely." As a validity check of our single-item measure, each participant was also randomly assigned to complete one of three previously established measures of trust adapted from Mayer and Davis (1999): overall trust scale ($\alpha = .67$), benevolence-based trust subscale ($\alpha = .94$), or ability-based trust subscale ($\alpha = .95$). This measurement approach was designed to minimize respondent fatigue and to avoid the risk of cross-contamination between scales. Web Appendix C reports more details about these scales. A high correlation between our single-item measure of trust and overall trust ($r = .52, p < .001$), benevolence-based trust ($r = .64, p < .001$), and ability-based trust ($r = .68, p < .001$) suggests that the single-item measure successfully captures trust. Thus, for efficiency, we use only this item in subsequent studies.

Results and Discussion

Attitudinal measure. Consistent with H₂, product interest was significantly lower when the privacy notice was present ($M = 35.41, SD = 29.77, N = 298$) versus absent ($M = 52.48, SD = 28.31, N = 302; F(1, 594) = 50.25, p < .001, \eta_p^2 = .08$).

Behavioral measure. The behavioral measure of interest showed a similar pattern; a lower proportion of participants were willing to spend time reading additional product information when the privacy notice was present (38.3%; $N = 114/298$) versus absent (56.3%; $N = 170/302; \chi^2(1) = 19.57, p < .001, \phi = .18$). On

average, participants who opted to view additional product information spent 22.94 seconds doing so (approximately 15% of the median duration of the entire survey, suggesting that it was not a trivial cost to participants). Moreover, when the time for participants who opted not to view additional product information was recorded as zero (as preregistered), the number of seconds participants were willing to spend reading additional product information was significantly lower when the privacy notice was present ($M = 7.77, SD = 16.97, N = 298$) versus absent ($M = 13.90, SD = 23.08, N = 302; F(1, 598) = 13.72, p < .001, \eta_p^2 = .02$).

Trust. To determine whether trust mediated the effect of the privacy notice on purchase interest, we conducted eight separate mediation analyses using PROCESS Model 4 (Hayes 2012). As predicted, all four measures of trust mediated the effect for both attitudinal and behavioral measures of purchase interest, as illustrated in Table 1.

To summarize, in a controlled online experiment, Study 3 supported H₂ by replicating the bulletproof glass effect observed in the field experiment and also provided evidence consistent with H₃ that a reduction in trust is the underlying mechanism.

Study 4: Moderated Mediation

Study 4 further examines the role of trust in the bulletproof glass effect through moderated mediation. Specifically, it tested H₄, that the bulletproof glass effect is more likely to be observed when consumers expect safety than when they are already distrustful.

Method

We recruited 602 participants from MTurk who were randomly assigned to one of four conditions in a 2 (privacy notice: present vs. absent) \times 2 (expected safety: safe vs. unsafe) between-participants design. Consistent with our preregistration plan (<https://aspredicted.org/blind.php?x=cg8dr8>), 73 participants who failed the attention check were excluded, leaving a final sample size of 529 participants (48.4% female; mean age = 37.80 years).⁶ All participants evaluated a real product available from an online retailer, Ruggie (<https://ruggie.co/>). Product details were displayed in an image captioned "The Alarm Clock You Turn Off With Your Feet."

To manipulate expected safety, we then showed all participants a recent (fictitious) news headline from the *Wall Street Journal*. In the safe condition, the headline read, "Ruggie Praised by FTC for Zero Consumer Privacy Violations During 2020." In the unsafe condition, the headline read, "Ruggie Cited by FTC for Multiple Consumer Privacy Violations During 2020."

⁵ Consistent with the privacy notice used in the previous studies, the language was more cognitive than affective ($z = 2.0, p < .05$), as reported in Web Appendix B.

⁶ The pattern and significance of the results do not change when all participants are included.

Table 1. Mediation Analyses.

Mediator	N	Attitudinal Measure			Behavioral Measure		
		Indirect Effect	LLCI	ULCI	Indirect Effect	LLCI	ULCI
Single-item trust measure	600	-11.16	-14.24	-8.43	-.44	-.62	-.28
Overall trust scale	195	-3.26	-7.06	-.14	-.17	-.39	-.003
Benevolence-based trust subscale	194	-7.74	-11.98	-3.82	-.32	-.62	-.11
Ability-based trust subscale	211	-3.86	-8.02	-.49	-.17	-.44	-.01

Notes: LLCI = lower-level confidence interval; ULCI = upper-level confidence interval.

The indirect effect of the privacy notice on attitudinal and behavioral measures of purchase interest was negative in each model.

Next, we showed participants in the present condition an excerpt from Ruggie's actual privacy notice, whereas this was omitted in the absent condition. All participants then indicated their purchase interest ("How interested would you be in purchasing this product?") and completed a single-item measure of trust ("How comfortable would you be with the way your data is collected and managed by this retailer?"), each measured using a sliding scale from 0 = "not at all" to 100 = "extremely." As an attention check, participants were asked to identify which of the two news headlines they had previously read, with an option to select "I don't remember."

Results and Discussion

Purchase interest. An analysis of variance (ANOVA) revealed a significant main effect of expected safety ($F(1, 525) = 68.19, p < .001, \eta_p^2 = .12$) as well as a marginally significant main effect of privacy notice on purchase interest ($F(1, 525) = 3.10, p = .079, \eta_p^2 = .01$). Consistent with H_4 , these main effects were qualified by a significant interaction ($F(1, 525) = 4.18, p = .042, \eta_p^2 = .01$). Specifically, in the safe condition, the bulletproof glass effect was replicated in that purchase interest was lower when the privacy notice was present ($M = 38.17, SD = 32.38, N = 135$) than when it was absent ($M = 48.25, SD = 33.45, N = 134; p = .007$). However, in the unsafe condition, the bulletproof glass effect was attenuated such that purchase interest did not differ when the privacy notice was present ($M = 21.70, SD = 27.46, N = 132$) or absent ($M = 20.95, SD = 27.93, N = 128; p = .842$). These results are illustrated in Figure 2.

Trust. We observed a similar pattern for trust; an ANOVA revealed a significant main effect of privacy notice on trust ($F(1, 525) = 4.18, p = .041, \eta_p^2 = .01$), as well as a significant main effect of expected safety ($F(1, 525) = 284.02, p < .001, \eta_p^2 = .35$). These main effects were also qualified by a significant interaction ($F(1, 525) = 28.16, p > .001, \eta_p^2 = .05$). Specifically, in the safe condition, trust was lower when the privacy notice was present ($M = 51.76, SD = 31.29, N = 135$) than when it was absent ($M = 69.68, SD = 28.11, N = 134; p < .001$). However, in the unsafe condition, the effect was reversed such that trust was higher when the privacy notice was

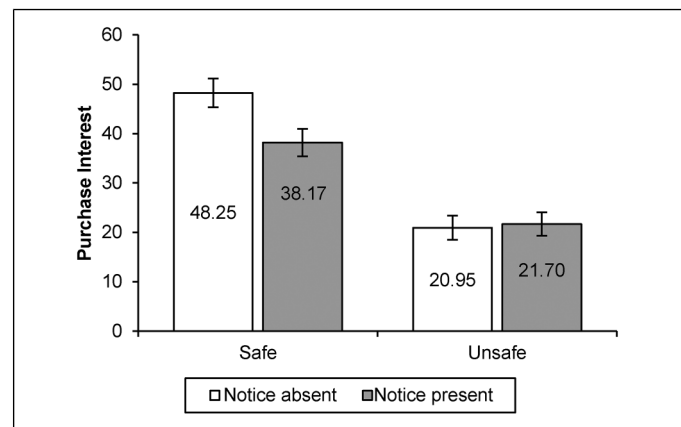


Figure 2. Expected safety moderates the bulletproof glass effect (Study 4).

present ($M = 23.61, SD = 26.91, N = 132$) than absent ($M = 15.66, SD = 25.29, N = 128; p = .023$).

Moderated mediation. To test whether the effect of privacy notice on purchase interest was mediated by trust and moderated by expected safety, we conducted a moderated mediation analysis using the PROCESS macro model 7 (Hayes 2018). Results indicated significant moderated mediation (95% confidence interval [CI]: [-22.09, -9.90]), with trust mediating the effect of a privacy notice on purchase interest in both the safe condition (95% CI: [-15.51, -6.60]) and, to a lesser degree, in the unsafe condition (95% CI: [.97, 8.85]). Note that the sign of the indirect effect reversed in the unsafe condition, suggesting that a privacy notice can help rather than hurt trust and purchase interest in a context where consumers have an a priori reason to be distrustful.

Not only does this study support H_3 by providing additional evidence of trust as the mechanism underlying the bulletproof glass effect, but it also supports H_4 by showing moderated mediation. Specifically, when the firm had a positive reputation for protecting customer data, exposure to a privacy notice reduced trust and purchase interest. However, this effect was attenuated when consumers had a reason to be distrustful before viewing the privacy notice.

Study 5: Benevolence Cues

Study 5 examined another potential moderator of the bulletproof glass effect. Specifically, we tested H_5 , that the negative effect of a privacy notice on purchase interest would be reduced by the addition of benevolence cues.

Method

We recruited 602 participants (59.6% female; mean age = 41.90 years) on MTurk. Participants were randomly assigned to one of three conditions (privacy notice: absent vs. standard vs. benevolent). All participants were shown an identical image and product description. In the absent condition, participants then proceeded directly to the dependent measure. In the standard condition, participants were asked to review the retailer's privacy notice and shown the same notice as was used in Study 3. In the benevolent condition, the notice was adapted slightly so as to subtly incorporate benevolence cues but add no objective information about data practices. These cues included the statements: "We care about your privacy," "We respect you and promise to treat you fairly," and "We are committed to the protection of your information."

The dependent measure, purchase interest ("How interested would you be in purchasing these sunglasses?"), was measured using a sliding scale ranging from 0 = "Not at all interested" to 100 = "Very interested." To control for any possible effect of time on the results, we also measured how long participants spent reading the notice as well as the overall duration of the survey.

Results and Discussion

An ANOVA revealed a significant main effect of condition on purchase interest; ($F(2, 599) = 42.67, p < .001, \eta_p^2 = .13$). Post hoc tests showed that, consistent with our previous studies, the bulletproof glass effect was replicated such that, compared with the absent condition ($M = 58.68, SD = 28.17, N = 200$), purchase interest was significantly lower after exposure to a privacy notice in both the standard ($M = 32.75, SD = 28.50, N = 202; p < .001$) and benevolent ($M = 41.40, SD = 29.25, N = 200; p < .001$) conditions. Moreover, consistent with our prediction in H_5 that incorporating benevolence cues would reduce the negative impact of a privacy notice, purchase interest was significantly higher in the benevolent versus standard condition ($p = .008$). These results are illustrated in Figure 3.

These effects remained significant when survey duration was included as a covariate, suggesting that the results cannot be explained by the additional time required by participants in the standard and benevolent (vs. absent) conditions to read the privacy notice. Furthermore, across the two conditions in which a privacy notice was shown, there was no significant difference in the number of seconds spent reading the notice ($M_{\text{standard}} = 35.13, SD = 30.67, N = 202$ vs. $M_{\text{benevolent}} = 36.23, SD = 32.85, N = 200; t(400) = .35, p = .73$). Together, the results of Study 5 provide support for H_5 and suggest another moderator




Figure 3. Benevolence cues can attenuate the bulletproof glass effect (Study 5).

of the bulletproof glass effect: namely, the addition of a benevolence cue to a privacy notice.

Study 6: Opaque Privacy Notices

This study provides another test of H_5 in a context where a privacy notice alludes to the existence of a full privacy policy but does not describe specific data management practices. As such, in addition to providing additional evidence of H_5 , Study 6 also addresses the possibility that the bulletproof glass effect is simply a product of consumers' distaste for the detailed description of specific data management practices in the privacy notices we have used thus far.

Method

We recruited 1,125 participants from MTurk (51.3% female; mean age = 40.17 years) who were randomly assigned to one of three conditions (privacy notice: absent vs. standard vs. benevolent) in a between-participants design that was preregistered (https://aspredicted.org/TVD_GTV). All participants were told, "Suppose you needed to buy some new clothes for an upcoming event and find some items you like on the website of an online retailer that you weren't previously familiar with. As you check out, you see the following screen." Participants were then shown a screenshot of a checkout page in which customer profile information was being collected. In the absent condition, there was no mention of a privacy notice. In the standard condition, the screenshot showed an arrow hovering over a question mark icon next to the cell phone number data field, with a pop-up window that stated, "Usage and sharing of this data is governed by the terms outlined in our Privacy Policy." The benevolent condition was identical, except that the message in the pop-up window included a benevolence cue that provided no objective information about data practices: "WE CARE about protecting your privacy!" (see Figure 4).

On the next page, all participants indicated their purchase interest ("How interested would you be in making a purchase from this retailer?"; 0 = "not at all interested," and 100 = "very interested").

Figure 4. Benevolent condition stimulus (Study 6).

Results and Discussion

An ANOVA revealed a significant effect of condition on purchase interest ($F(2, 1,122) = 23.46, p < .001, \eta_p^2 = .04$). A post hoc test showed that each contrast was significant. Specifically, the bulletproof glass effect was replicated in that purchase interest was lower when the standard privacy notice was present ($M = 45.54, SD = 25.66, N = 377$) than when it was absent ($M = 51.18, SD = 25.70, N = 377; p = .006$). However, purchase interest was higher in the benevolent condition ($M = 58.13, SD = 24.12, N = 371$) than in both the absent ($p < .001$) and standard ($p < .001$) conditions, demonstrating a reversal of the bulletproof glass effect when a benevolence cue was added. These results are illustrated in Figure 5.

Study 6 provided further support for H_5 , showing that the negative effect of a privacy notice on purchase interest was reversed when a benevolence cue was added. Importantly, this study also shows that the bulletproof glass effect may occur even when consumers do not read details about specific data management practices—a situation that prior research suggests is common even when such details are provided (Milne and Culnan 2004). Our conservative test suggests that unless tempered by benevolence cues, merely alerting consumers to the existence of formal privacy-related policies is sufficient to decrease purchase interest.

Study 7: Conspicuous Absence of Privacy Details

Study 7 offers further support for our contention that the bulletproof glass effect is not limited to situations in which consumers read complete details about a firm's data management practices. Whereas Studies 5 and 6 showed that the mere




Figure 5. A benevolence cue reversed the bulletproof glass effect (Study 6).

Notes: Relative to the absence of a privacy notice, a standard privacy notice decreased purchase interest. However, the bulletproof glass effect was reversed when a benevolence cue was added.

presence of a privacy notice can decrease purchase interest; in Study 7, we test whether the conspicuous absence of a privacy notice can also decrease purchase interest. The logic behind this prediction is that trust can be undermined when consumers are made aware that information is conspicuously absent (John, Barasz, and Norton 2016). Thus, by formalizing the format of privacy information, such as adopting standardized templates for displaying privacy-related practices (e.g., Apple's privacy nutrition labels), consumers may become distrustful when privacy details are absent (in addition to when privacy details are presented as a formal contract, as shown in our previous studies).

Until recently, apps were not required to include a privacy notice; an analysis of over one million apps in the Google Play Store between August 2017 and May 2018 found that only about half (41.7%, 45.2%, and 51.8% on three separate crawls) included a privacy policy link (Story, Zimmeck, and Sadeh 2018). However, in December 2020, Apple made privacy nutrition labels mandatory in the App Store. Under these new regulations, when a developer has not provided privacy details to Apple, the absence of such information is obvious to consumers who view the privacy nutrition label and consider whether to download the app. Thus, in Study 7, we compare a control condition (in which no privacy-related information is provided) to two different treatment conditions that are both expected to reduce consumers' interest in downloading an app. One treatment condition examines how consumers respond when exposed to privacy details, and the other treatment condition examines how consumers respond when privacy information is conspicuously absent.

Method

We recruited 300 participants from MTurk (52.7% female; mean age = 39.26 years) who were randomly assigned to one of three conditions (privacy details: absent vs. present vs. conspicuously absent) in a between-participants design that was preregistered (<https://aspredicted.org/blind.php?x=ti9rf4>). All participants were told, "Imagine you planned to open a new retirement account and were evaluating different investment apps. One of the apps you are considering is Nest Egg, Inc., which uses a data-driven approach to help you meet your financial goals. Please take a moment to examine the screenshot of this app." They were then shown a mocked-up mobile phone screenshot of a fictitious investment app and told that it was one of the apps they were considering. This screenshot described the app's data-driven approach to investing and how a consumer's responses to a representative's questions about risk tolerance and investment objectives during a consultation would be combined with a large amount of personal data to provide personalized investment guidance (see Figure 6).

Participants in the present condition were then shown a screenshot of Apple's privacy nutrition label for the app, which described the types of data (e.g., contact info, location) that can be used to track the user across apps and websites owned by other companies. Participants in the conspicuously absent condition were shown a similar screenshot of Apple's privacy nutrition label but, consistent with what Apple actually displays on the App Store for developers that have not provided details about their privacy practices, the screenshot indicated that no details had been provided and that the developer will be required to provide privacy details when it submits its next app update. Participants in the absent condition proceeded directly from the app screenshot to the dependent measure.

As an attitudinal measure of purchase interest, all participants then responded to the question, "How interested would you be in downloading this app?" (0 = "not at all," and 100 = "extremely"). As a proxy for behavior, we also told participants

that the app normally costs \$1.99 and asked them, "At the conclusion of the study, would you like to receive a code to download the app for free?" (1 = "Yes, give me a free download code," 0 = "No thanks"). Finally, respondents completed demographic questions and were debriefed that the app was fictitious.

Results and Discussion

Attitudinal measure. We predicted that, compared with the mere absence of a privacy notice in the absent condition, the presence or conspicuous absence of a privacy notice would decrease interest in downloading the app. An ANOVA revealed a significant effect of condition on purchase interest ($F(2, 297) = 10.01$, $p < .001$, $\eta_p^2 = .06$). A post hoc test showed that the contrasts between both treatment conditions versus the control condition were significant. Specifically, the bulletproof glass effect was replicated in that, compared with the absent condition ($M = 50.84$, $SD = 28.10$, $N = 100$), purchase interest was lower when privacy details were present ($M = 34.98$, $SD = 32.12$, $N = 100$; $p < .001$) and when privacy details were conspicuously absent ($M = 32.96$, $SD = 32.43$, $N = 100$; $p < .001$). The present and conspicuously absent conditions did not differ significantly from each other ($p = 1.00$). These results are illustrated in Figure 7.

Behavioral proxy. As a further test of our hypothesis, we analyzed participants' desire to receive a code to download the app for free at the end of the study. The pattern of results matched that of the attitudinal measure of purchase interest; a chi-squared test revealed a significant effect of condition on the behavioral proxy ($\chi^2(2) = 8.53$, $p = .014$, $\phi = .169$). Specifically, the bulletproof glass effect was replicated in that, compared with the 46.0% ($N = 46/100$) of participants in the absent condition who chose to receive the free download code, significantly fewer participants opted to do so when privacy details were present (27.0%; $N = 27/100$; $\chi^2(1) = 7.79$, $p = .005$, $\phi = .197$) or conspicuously absent (32.0%; $N = 32/100$; $\chi^2(1) = 4.12$, $p = .042$, $\phi = .144$).

These results complement our previous findings by showing that, like the presence of a privacy notice, the conspicuous absence of privacy details is also sufficient to decrease purchase interest. This suggests that the higher purchase interest observed in previous studies when a privacy notice is missing is not because consumers prefer to avoid details about a firm's data management practices but, rather, because the concept of a formal privacy notice breeds distrust, and, in turn, reduces purchase interest. Our finding that the conspicuous absence of privacy information decreases purchase interest is consistent with Lwin, Wirtz, and Williams (2007), whereby participants who were explicitly told that "there was no mention of a privacy policy" exhibited greater privacy concerns than participants who were provided with a comprehensive privacy policy. Indeed, prior research suggests that dormant privacy concerns can be triggered by merely mentioning privacy-related topics (Marreiros et al. 2017). In one study, consumers who were explicitly primed to think of privacy were less willing to reveal their personal information on an unsafe

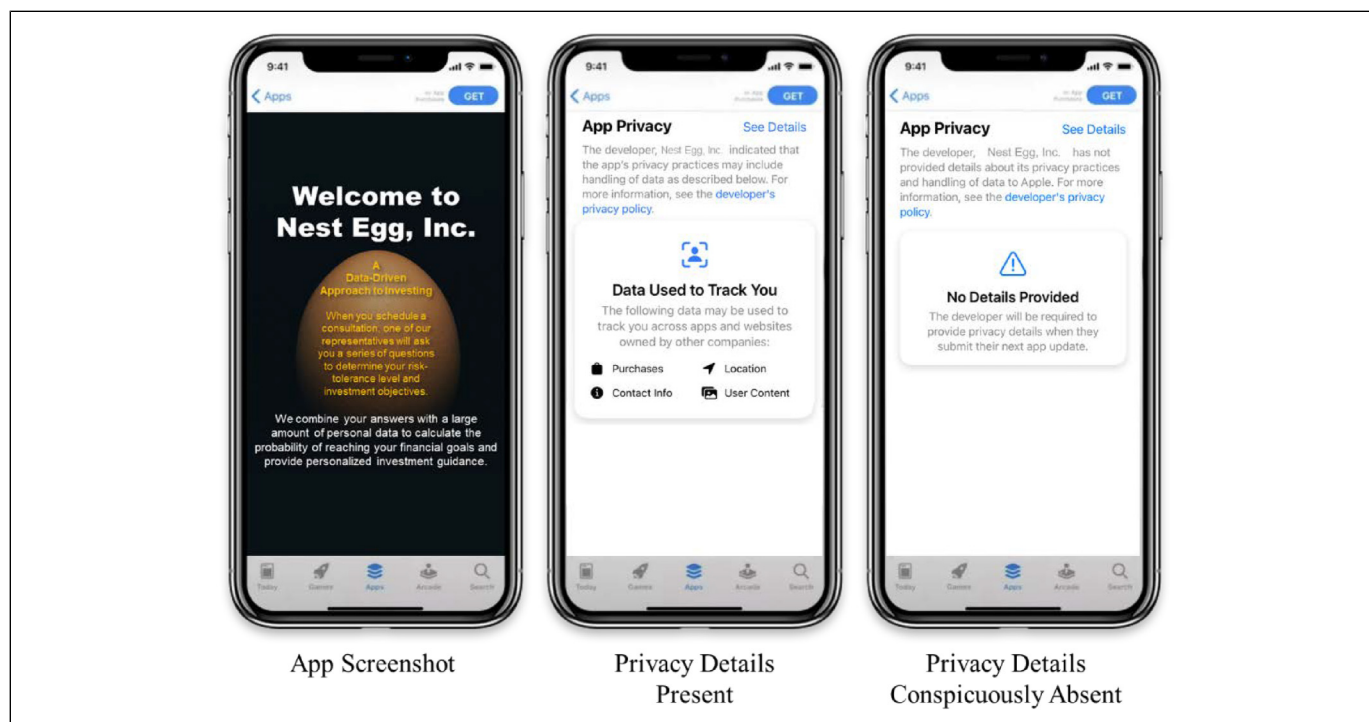


Figure 6. Privacy nutrition label stimuli (Study 7).




Figure 7. Purchase interest declines when privacy details are present or conspicuously absent (Study 7).

Notes: Relative to an absent control condition, interest in downloading an app was lower when privacy details were present or conspicuously absent.

website than consumers who had not been primed (John, Acquisti, and Loewenstein 2011). Our results build on these findings by suggesting that when the concept of a formal privacy contract is made salient—whether by the presence of a privacy notice or by its conspicuous absence—trust and purchase interest may decrease.

Study 7 also provides insight into how purchase interest may be affected differently by an absence of privacy information in regulated versus unregulated contexts. Specifically, although an absence (vs. presence) of privacy details can result in greater

purchase interest when attention is not drawn to the absence (e.g., in unregulated contexts, where the availability of privacy information may vary widely across firms, industries, and geographies), it is unlikely to do so in contexts such as the App Store, where current regulations standardize the presentation of privacy information and draw consumers’ attention to any unavailable information.

General Discussion

Our results challenge a prevailing intuition among managers that privacy notices will cause consumers to feel more secure. Although privacy notices place legally enforceable limits on a firm’s data practices, communicate safeguards, and signal transparency, we find that instead of promoting a sense of confidence that personal data will not be misused, privacy notices often have the unintended consequence of causing consumers to become less trusting and less interested in making a purchase. We show that even explicitly protective privacy notices, as well as those that provide no objective information about data practices, can undermine consumer trust and potentially hurt sales. Notably, a field study shows that when privacy protections were made more salient, enrollment decreased. Despite the importance of the topic of privacy from both theoretical and managerial perspectives, we know of no other field study that has manipulated the salience of a privacy notice and measured resulting consumer behavior.

A decrease in purchase interest caused by exposure to a privacy notice was replicated in multiple studies, using both

attitudinal and behavioral measures, and multiple measures of trust were shown to mediate this effect. Although most of our studies focused on personal data collected at a single point in time during the process of conducting a transaction, Study 7 showed the bulletproof glass effect for an app that continues to collect, store, and transmit personal information on an ongoing basis.

Moreover, we identified several moderators, showing that the bulletproof glass effect is attenuated when consumers have a priori expectations that their personal data may not be safe and that the effect may even be reversed when benevolence cues are incorporated into a privacy notice. Given that our analysis of real privacy policies showed that most contain little affective language that can foster benevolence-based trust, this moderator is of great practical importance; indeed, as illustrated in Web Appendix B, LIWC analyses of our stimuli indicated that in all cases in which privacy notices decreased purchase interest, there was either a paucity (Studies 2–6) or complete absence (Study 7) of affective language. We also showed that purchase interest may decline in response to not only the presence of privacy details but also their conspicuous absence.

Contributions

Our findings offer several contributions to the marketing literature and have managerial and policy implications. First, we measured managers' expectations regarding how consumers will respond to privacy notices and documented a miscalibration between these expectations and consumer responses. Broadly, we contribute to the growing body of marketing literature by showing that consumers sometimes react to information about risks (e.g., privacy risks) in seemingly paradoxical ways. Although managers expected privacy notices to help consumers feel more secure, our studies suggest that consumers may view them more like warnings. In contrast to prior work suggesting that assurances increase compliance when survey respondents are asked to provide sensitive personal data (Singer, Hippler, and Schwarz 1992; Singer, Von Thurn, and Miller 1995), our findings illustrate conditions under which an opposite pattern may occur—privacy notices decreased interest in purchasing a product and providing the corresponding personal data.

While we recognize that not all privacy notices are necessarily intended to be assuring, documenting the unintended consequence of privacy notices on purchase interest adds to our understanding of the conditions in which backfire may occur. As documented by prior work, increasing the salience of risky behavior through measurement can be counterproductive (Fitzsimons and Moore 2008), and warning messages do not always achieve their intended effects, sometimes failing to increase consumer compliance (Argo and Main 2004; Menon, Block, and Ramanathan 2002; Stewart and Martin 1994) or even resulting in greater acceptance of the false claims that people were warned against (Skurnik et al. 2005). In addition, consumers seem to trust advisors who disclose conflicts of interest (Cain, Loewenstein, and Moore 2011; Sah, Malaviya,

and Thompson 2018) and tend to be more persuaded by messages that include negative information (Ein-Gar, Shiv, and Tormala 2012; Herr, Kardes, and Kim 1991; Ward and Brenner 2006). One mechanism that has been identified in the persuasion literature for these kinds of effects is peripheral or heuristic (vs. central or elaborative) processing (Herbst et al. 2012; Meyers-Levy and Malaviya 1999; Sah, Malaviya, and Thompson 2018). Our results suggest that using benevolence cues to foster trust may be a complementary mechanism.

Second, we provide converging evidence across multiple studies, including what we believe is the first manipulation of the salience of a privacy notice in the field, that a salient privacy notice may have unintended consequences by reducing consumers' trust and purchase interest. Indeed, our results across multiple studies showed that consumers were more likely to transact with an organization that lacked a privacy notice than with an organization that provided a transparent description of its data practices. Transparency in data practices, and the lack thereof, has been the source of much debate. Many of the transformational technologies that are influencing both marketers and consumers at an unprecedented rate, such as artificial intelligence and other forms of automation to collect and analyze consumer data, are deeply invasive of consumer privacy and obfuscate privacy risks (Leung, Paolacci, and Puntoni 2018; Mende et al. 2019; Puntoni et al. 2021; Wertenbroch 2019). Though regulators and consumer advocacy groups demand more transparency, we find that customers may react negatively to the transparency offered by formal privacy notices. These results are consistent with prior work in marketing communications that has demonstrated negative reactions to full transparency, finding that consumers may respond more favorably to imprecision than precision (Isaac, Brough, and Grayson 2016). They are also consistent with work in advertising, showing that ad performance declines when consumers are informed that an ad was generated using their personal information in privacy-invasive ways (Kim, Barasz, and John 2019).

By illustrating a situation in which consumers seem to respond more favorably to (quiet) omission than transparency, our findings are also conceptually related to the consumer research on information avoidance (Sweeny et al. 2010; Woolley and Risen 2021). This body of work shows that consumers often prefer ignorance to bad news. For example, "the ostrich effect" describes the tendency of investors who receive preliminary bad or ambiguous news to shield themselves from further news by monitoring their accounts less frequently (Galai and Sade 2006; Karlsson, Loewenstein, and Seppi 2009). Nonetheless, the negative reaction we observed when details about a firm's privacy practices are *conspicuously* absent suggests that consumers' hesitation to transact with organizations that have a privacy notice is not likely driven by an active aversion to privacy-related information. Instead, the effect seems to be due to the formality of privacy notices, and may be tempered when benevolence cues are incorporated into the notice.

Third, in contrast to the notion that consumers respond only to changes in the *content* of privacy notices, we show that consumers' purchase interest may also be affected by the mere

presence of a privacy notice, even if it provides no specific details about privacy practices. This finding may cause companies to hesitate to draw consumers' attention to privacy protections. However, our findings also offer an initial exploration of how policy makers and/or well-intentioned firms might mitigate the negative effects of a formal privacy notice on consumers' purchase interest. First, regulators could require the use of standardized templates that make an absence of privacy details conspicuous. While our findings suggest that such regulation could level the playing field by eliminating any advantage a company could gain by failing to disclose its privacy practices, mandating formal privacy notices could also have an unintended side effect of producing a climate of widespread distrust. Another potential solution suggested by our results is to add benevolence cues to (consumer-protective) privacy notices. Our content analysis of real companies' privacy notices found that the content of most notices tends to use more cognitive than affective language. Studies 5 and 6 indicated that merely prefacing mention of the privacy notice with benevolence cues such as "we care about protecting your privacy" was sufficient to attenuate or reverse the bulletproof glass effect. Together, these findings provide actionable guidance to managers on how to effectively convey privacy information (without hurting purchase interest).

Directions for Future Research

Finally, our work prompts many additional questions that could be explored in future research. The lack of privacy research in consumer behavior has been noted (Brough and Martin 2020; Kim, Barasz, and John 2020; Krishna 2020), and more work is needed to understand the multiplicity of factors that likely shape the effect of privacy notices on consumer behavior. For example, whereas regulation often influences the presence, content, and format of privacy notices, future research could explore how shifting requirements affect norms over time. As transparency becomes increasingly required, the absence of a privacy notice may become more conspicuous and, consistent with the results of Study 7, the negative impact of a standard privacy notice on consumers' purchase interest may decrease.

Another opportunity for future research lies in better understanding the relationship between trust and expected safety—although Study 4 focused on how expected safety impacts trust, it is possible that these constructs have a bidirectional influence on one another, and each may be affected by individual differences, prior experiences with a particular company, and/or prior experiences with privacy violations more generally. Another aspect that could be further explored is the relationship of these constructs with privacy concern. Although prior research has found privacy concern to be inversely correlated with trust and purchase intent (Eastlick, Lotz, and Warrington 2006), future research could directly measure how privacy notices impact privacy concern.

Other opportunities could lie in the exploration of individual differences; in particular, given the attenuation of the bulletproof glass effect by the addition of benevolence cues, it

seems plausible that the effect may be pronounced among consumers who chronically adopt an intuitive or experiential thinking style (Epstein et al. 1996). Further, whereas we focused exclusively on privacy notices, additional research might compare the relative impact (on purchase interest) of privacy notices versus other modes of communicating privacy-related information, such as privacy seals like an TRUSTe icon (Miyazaki and Krishnamurthy 2002; Rifon, LaRose, and Choi 2005; Wang, Beatty, and Foss 2004). Future research could also explore additional contexts in which measures designed to protect consumers, such as security screening at K–12 schools or armed guards in public settings, may undermine trust and evoke negative responses despite the protections they offer. Our results suggest that in such situations, benevolence cues may be a key to avoiding unintended consequences.

Acknowledgments

The authors thank the *JMR* review team, Mathew Isaac, Joseph Goodman, and Rebecca Reczek for their helpful comments on earlier versions of this manuscript and Anne Marie Green, Holly Howe, and Trevor Spelman for research assistance.

Associate Editor

Darren Dahl


Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship and/or publication of this article.

ORCID iD

Aaron R. Brough  <https://orcid.org/0000-0003-2457-3199>

References

- Acquisti, Alessandro, Laura Brandimarte, and George Loewenstein (2015), "Privacy and Human Behavior in the Age of Information," *Science*, 347 (6221), 509–14.
- Acquisti, Alessandro, Leslie K. John, and George Loewenstein (2013), "What Is Privacy Worth?" *Journal of Legal Studies*, 42 (2), 249–74.
- Adjerid, Idris, Alessandro Acquisti, and George Loewenstein (2019), "Choice Architecture, Framing, and Cascaded Privacy Choices," *Management Science*, 65 (5), 2267–90.
- 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.
- Argo, Jennifer J. and Kelley J. Main (2004), "Meta-Analyses of the Effectiveness of Warning Labels," *Journal of Public Policy and Marketing*, 23 (2), 193–208.

- Athey, Susan, Christian Catalini, and Catherine Tucker (2017), "The Digital Privacy Paradox: Small Money, Small Costs, Small Talk," working paper, National Bureau of Economic Research.
- Brandimarte, Laura, Alessandro Acquisti, and George Loewenstein (2013), "Misplaced Confidences: Privacy and the Control Paradox," *Social Psychology and Personality Science*, 4 (3), 340–47.
- Brough, Aaron R. and Kelly D. Martin (2020), "Critical Roles of Knowledge and Motivation in Privacy Research," *Current Opinion in Psychology*, 31, 11–15.
- Cain, Daylian M., George Loewenstein, and Don A. Moore (2011), "When Sunlight Fails to Disinfect: Understanding the Perverse Effects of Disclosing Conflicts of Interest," *Journal of Consumer Research*, 37 (5), 836–57.
- Cases, Anne-Sophie, Christophe Fournier, Pierre-Louis Dubois, and John F. Tanner (2010), "Web Site Spill Over to Email Campaigns: The Role of Privacy, Trust and Shoppers' Attitudes," *Journal of Business Research*, 63 (9/10), 993–99.
- Culnan, Mary J. (2000), "Protecting Privacy Online: Is Self-Regulation Working?" *Journal of Public Policy & Marketing*, 19 (1), 20–26.
- Culnan, Mary J. and Pamela K. Armstrong (1999), "Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation," *Organizational Science*, 10 (1), 104–15.
- Donaldson, Thomas and Thomas W. Dunfee (1999), *Ties That Bind: A Social Contract Approach to Business Ethics*. Cambridge, MA: Harvard Business School Press.
- Eastlick, Mary Ann, Sherry L. Lotz, and Patricia Warrington (2006), "Understanding Online B-to-C Relationships: An Integrated Model of Privacy Concerns, Trust, and Commitment," *Journal of Business Research*, 59 (8), 877–86.
- Ein-Gar, Danit, Baba Shiv, and Zakary L. Tomala (2012), "When Blemishing Leads to Blossoming: The Positive Effect of Negative Information," *Journal of Consumer Research*, 38 (5), 846–59.
- Epstein, Seymour, Rosemary Pacini, Veronika Denes-Raj, and Harriet Heier (1996), "Individual Differences in Intuitive-Experiential and Analytical-Rational Thinking Styles," *Journal of Personality and Social Psychology*, 71 (2), 390–405.
- Fitzsimons, Gavan J. and Sarah G. Moore (2008), "Should We Ask Our Children About Sex, Drugs, and Rock & Roll? Potentially Harmful Effects of Asking Questions About Risky Behaviors," *Journal of Consumer Psychology*, 18 (2), 82–95.
- Galai, Dan and Orly Sade (2006), "The 'Ostrich Effect' and the Relationship Between the Liquidity and the Yields of Financial Assets," *Journal of Business*, 79 (5), 2741–59.
- Goldfarb, Avi and Catherine Tucker (2013), "Why Managing Consumer Privacy Can Be an Opportunity," *MIT Sloan Management Review*, 54 (3), 10–12.
- Hayes, Andrew F. (2012), "PROCESS: A Versatile Computational Tool for Observed Variable Mediation, Moderation, and Conditional Process Modeling," <http://www.afhayes.com/public/process2012.pdf>.
- Hayes, Andrew F. (2018), "Partial, Conditional, and Moderated Moderated Mediation: Quantification, Inference, and Interpretation," *Communication Monographs*, 85 (1), 4–40.
- Hazel, James W. and Christopher Slobogin (2018), "Who Knows What, and When? A Survey of the Privacy Policies Proffered by U.S. Direct-to-Consumer Genetic Testing Companies," *Cornell Journal of Law and Public Policy*, 28 (5), 35–66.
- Herbst, Kenneth C., Eli J. Finkel, David Allan, and Gráinne M. Fitzsimons (2012), "On the Dangers of Pulling a Fast One: Advertisement Disclaimer Speed, Brand Trust, and Purchase Intention," *Journal of Consumer Research*, 38 (5), 909–19.
- Herr, Paul M., Frank R. Kardes, and John Kim (1991), "Effects of Word-of-Mouth and Product-Attribute Information on Persuasion: An Accessibility-Diagnosticity Perspective," *Journal of Consumer Research*, 17 (4), 454–62.
- Hoffman, Donna L., Thomas P. Novak, and Marcos A. Peralta (1999b), "Building Consumer Trust Online," *Communications of the ACM*, 42 (4), 80–85.
- Hoffman, Donna L., Thomas P. Novak, and Marcos A. Peralta (1999a), "Information Privacy in the MarketSpace: Implications for the Commercial Uses of Anonymity on the Web," *The Information Society*, 15 (2), 129–40.
- Isaac, Mathew S., Aaron R. Brough, and Kent Grayson (2016), "Is Top 10 Better Than Top 9? The Role of Expectations in Consumer Response to Imprecise Rank Claims," *Journal of Marketing Research*, 53 (3), 338–53.
- John, Leslie K., Alessandro Acquisti, and George Loewenstein (2011), "Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information," *Journal of Consumer Research*, 37 (5), 858–73.
- John, Leslie K., Kate Barasz, and Michael I. Norton (2016), "Hiding Personal Information Reveals the Worst," *Proceedings of the National Academy of Sciences*, 113 (4), 954–59.
- Kamleitner, Bernadette, Vincent W. Mitchell, Andrew T. Stephen, and Ardi Kolah (2018), "Your Customers May Be the Weakest Link in Your Data Privacy Defenses," *MIT Sloan Management Review* (May 22), <https://sloanreview.mit.edu/article/your-customers-may-be-the-weakest-link-in-your-data-privacy-defenses/>.
- Karlsson, Niklas, George Loewenstein, and Duane Seppi (2009), "The Ostrich Effect: Selective Attention to Information," *Journal of Risk and Uncertainty*, 38 (2), 95–115.
- Kim, Tami, Kate Barasz, and Leslie K. John (2019), "Why Am I Seeing This Ad? The Effect of Ad Transparency on Ad Effectiveness," *Journal of Consumer Research*, 45 (5), 906–32.
- Kim, Tami, Kate Barasz, and Leslie K. John (2020), "Consumer Disclosure," *Consumer Psychology Review*, 4 (1), 59–69.
- Krishna, Aradhna (2020), "Privacy Is a Concern: An Introduction to the Dialogue on Privacy," *Journal of Consumer Psychology*, 30 (4), 733–35.
- Lamberton, Cait and Andrew Stephen (2016), "A Thematic Exploration of Digital, Social Media, and Mobile Marketing: Research Evolution from 2000 to 2015 and an Agenda for Future Inquiry," *Journal of Marketing*, 80 (Special Issue), 146–72.
- Lauer, Thomas W. and Xiaodong Deng (2007), "Building Online Trust Through Privacy Practices," *International Journal of Information Security*, 6, 323–31.
- Leung, Eugina, Gabriele Paolacci, and Stefano Puntoni (2018), "Man Versus Machine: Resisting Automation in Identity-Based

- Consumer Behavior,” *Journal of Marketing Research*, 55 (6), 818–831.
- Levin, Daniel Z. and Rob Cross (2004), “The Strength of Weak Ties You Can Trust: The Mediating Role of Trust in Effective Knowledge Transfer,” *Management Science*, 50 (11), 1477–90.
- Lwin, May, Jochen Wirtz, and Jerome D. Williams (2007), “Consumer Online Privacy Concerns and Responses: A Power-Responsibility Equilibrium Perspective,” *Journal of the Academy of Marketing Science*, 35, 572–85.
- Malhotra, Deepak and J. Keith Murnighan (2002), “The Effects of Contracts on Interpersonal Trust,” *Administrative Science Quarterly*, 47 (3), 534–59.
- Marreiros, Helia, Mirco Tonin, Michael Vlassopoulos, and M.C. Schraefel (2017), “‘Now That You Mention It’: A Survey Experiment on Information, Inattention and Online Privacy,” *Journal of Economic Behavior & Organization*, 140, 1–17.
- Martin, Kirsten E. (2012), “Diminished or Just Different? A Factorial Vignette Study of Privacy as a Social Contract,” *Journal of Business Ethics*, 111, 519–39.
- Martin, Kirsten E. (2015), “Privacy Notices as Tabula Rasa: An Empirical Investigation into How Complying with a Privacy Notice Is Related to Meeting Privacy Expectations Online,” *Journal of Public Policy & Marketing*, 34 (2), 210–27.
- Martin, Kirsten E. (2016), “Do Privacy Notices Matter? Comparing the Impact of Violating Formal Privacy Notices and Informal Privacy Norms on Consumer Trust Online,” *Journal of Legal Studies*, 45 (S2), S191–215.
- Martin, Kirsten E. (2018), “The Penalty for Privacy Violations: How Privacy Violations Impact Trust Online,” *Journal of Business Research*, 82, 103–16.
- Martin, Kelly D., Abhishek Borah, and Robert W. Palmatier (2017), “Data Privacy: Effects on Customer and Firm Performance,” *Journal of Marketing*, 81 (1), 36–58.
- Martin, Kelly D. and Patrick E. Murphy (2016), “The Role of Data Privacy in Marketing,” *Journal of the Academy of Marketing Science*, 45 (2), 145–55.
- Mayer, Roger C. and James H. Davis (1999), “The Effect of the Performance Appraisal System on Trust for Management: A Field Quasi-Experiment,” *Journal of Applied Psychology*, 84 (1), 123–36.
- 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.
- McAllister, Daniel J. (1995), “Affect- and Cognition-Based Trust as Foundations for Interpersonal Cooperation in Organizations,” *Academy of Management Journal*, 38 (1), 24–59.
- McCole, Patrick, Elaine Ramsey, and John Williams (2010), “Trust Considerations on Attitudes Towards Online Purchasing: The Moderating Effect of Privacy and Security Concerns,” *Journal of Business Research*, 63 (9/10), 1018–24.
- Mende, Martin, Maura L. Scott, Jenny van Doorn, Dhruv Grewal, and Ilana Shanks (2019), “Service Robots Rising: How Humanoid Robots Influence Service Experiences and Elicit Compensatory Consumer Responses,” *Journal of Marketing Research*, 56 (4), 535–56.
- Menon, Geeta, Lauren G. Block, and Suresh Ramanathan (2002), “We’re at as Much Risk as We Are Led to Believe: Effects of Message Cues on Judgments of Health Risk,” *Journal of Consumer Research*, 28 (4), 533–49.
- Meyers-Levy, Joan and Prashant Malaviya (1999), “Consumers’ Processing of Persuasive Advertisements: An Integrative Framework of Persuasion Theories,” *Journal of Marketing*, 63 (4), 45–60.
- Miller, Chance (2021), “Report Speculates That Google Hasn’t Updated Its iOS Apps in Weeks to Avoid Providing Privacy Details,” 9to5Mac (January 5), <https://9to5mac.com/2021/01/05/google-privacy-details-app-store-apple/>.
- Milne, George R. and Mary J. Culnan (2004), “Strategies for Reducing Online Privacy Risks: Why Consumers Read (or Don’t Read) Online Privacy Notices,” *Journal of Interactive Marketing*, 18 (3), 15–29.
- Milne, George R., George Pettinico, Fatima M. Hajjat, and Ereni Markos (2017), “Information Sensitivity Typology: Mapping the Degree and Type of Risk Consumers Perceive in Personal Data Sharing,” *Journal of Consumer Affairs*, 51 (1), 133–61.
- Miyazaki, Anthony D. (2009), “Perceived Ethicality of Insurance Claim Fraud: Do Higher Deductibles Lead to Lower Ethical Standards?” *Journal of Business Ethics*, 87, 589–98.
- Miyazaki, Anthony D. and Sandeep Krishnamurthy (2002), “Internet Seals of Approval: Effects on Online Privacy Policies and Consumer Perceptions,” *Journal of Consumer Affairs*, 36 (1), 28–49.
- Mourey, James A. and Ari Ezra Waldman (2020), “Past the Privacy Paradox: The Importance of Privacy Changes as a Function of Control and Complexity,” *Journal of the Association for Consumer Research*, 5 (2), 162–80.
- Nissenbaum, Helen (2004), “Privacy as Contextual Integrity,” *Washington Law Review*, 79 (1), 119–57.
- Nowak, Glen J. and Joseph E. Phelps (1992), “Understanding Privacy Concerns: An Assessment of Consumers’ Information-Related Knowledge and Beliefs,” *Journal of Direct Marketing*, 6 (4), 28–39.
- Pan, Yue and George M. Zinkhan (2006), “Exploring the Impact of Online Privacy Disclosures on Consumer Trust,” *Journal of Retailing*, 82 (4), 331–38.
- Phelps, Joseph E., Giles D’Souza, and Glen J. Nowak (2001), “Antecedents and Consequences of Consumer Privacy Concerns: An Empirical Investigation,” *Journal of Interactive Marketing*, 15 (4), 2–17.
- 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.
- Puntoni, Stefano, Rebecca Walker Reczek, Markus Giesler, and Simona Botti (2021), “Consumers and Artificial Intelligence: An Experiential Perspective,” *Journal of Marketing*, 85 (1), 131–51.
- Puranam, Phanish and Bart S. Vanneste (2009), “Trust and Governance: Untangling a Tangled Web,” *Academy of Management Review*, 34 (1), 11–31.
- Reidenberg, Joel R., Jaspreet Bhatia, Travis Breaux, and Thomas B. Norton (2016), “Ambiguity in Privacy Policies and the Impact of Regulation,” *Journal of Legal Studies*, 45 (2), 163–90.

- Rifon, Nora J., Robert LaRose, and Sejung Marina Choi (2005), "Your Privacy Is Sealed: Effects of Web Privacy Seals on Trust and Personal Disclosures," *Journal of Consumer Affairs*, 39 (2), 339–62.
- Robinson, Sandra L. (1996), "Trust and Breach of the Psychological Contract," *Administrative Science Quarterly*, 41 (4), 574–99.
- 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.
- Sah, Sunita, Prashant Malaviya, and Debora Thompson (2018), "Conflict of Interest Disclosure as an Expertise Cue: Differential Effects Due to Automatic Versus Deliberative Processing," *Organizational Behavior and Human Decision Processes*, 147 (July), 127–46.
- Schoorman, F. David, Roger C. Mayer, and James H. Davis (2007), "An Integrative Model of Organizational Trust: Past, Present, and Future," *Academy of Management Review*, 32 (2), 344–54.
- Singer, Eleanor, Hans-Jürgen Hippler, and Norbert Schwarz (1992), "Confidentiality Assurances in Surveys: Reassurance or Threat?" *International Journal of Public Opinion Research*, 4 (3), 256–68.
- Singer, Eleanor, Dawn R. Von Thurn, and Esther R. Miller (1995), "Confidentiality Assurances and Response: A Quantitative Review of the Experimental Literature," *Public Opinion Quarterly*, 59 (1), 66–77.
- Skurnik, Ian, Carolyn Yoon, Denise C. Park, and Norbert Schwarz (2005), "How Warnings About False Claims Become Recommendations," *Journal of Consumer Research*, 31 (4), 713–24.
- Smith, H. Jeff, Tamara Dinev, and Heng Xu (2011), "Information Privacy Research: An Interdisciplinary Review," *MIS Quarterly*, 35 (4), 989–1015.
- Stewart, David W. and Ingrid M. Martin (1994), "Intended and Unintended Consequences of Warning Messages: A Review and Synthesis of Empirical Research," *Journal of Public Policy & Marketing*, 13 (1), 1–19.
- Story, Peter, Sebastian Zimmeck, and Norman Sadeh (2018), "Which Apps Have Privacy Policies?" in *Annual Privacy Forum*. Cham, Switzerland: Springer, 3–23.
- Sweeny, Kate, Darya Melnyk, Wendi Miller, and James A. Shepperd (2010), "Information Avoidance: Who, What, When, and Why," *Review of General Psychology*, 14 (4), 340–53.
- Tang, Zhulei, Yu (Jeffrey) Hu, and Michael D. Smith (2008), "Gaining Trust Through Online Privacy Protection: Self-Regulation, Mandatory Standards, or Caveat Emptor," *Journal of Management Information Systems*, 24 (4), 153–73.
- Tsai, Janice Y., Serge Egelman, Lorrie Cranor, and Alessandro Acquisti (2011), "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study," *Information Systems Research*, 22 (2), 254–68.
- Tucker, Catherine E. (2014), "Social Networks, Personalized Advertising, and Privacy Controls," *Journal of Marketing Research*, 51 (5), 546–62.
- Turow, Joseph, Michael X. Delli Carpini, Nora A. Draper, and Rowan Howard-Williams (2012), "Americans Roundly Reject Tailored Political Advertising," Annenberg School for Communication, University of Pennsylvania, https://repository.upenn.edu/asc_papers/398.
- Wang, Sijun, Sharon E. Beatty, and William Foxx (2004), "Signaling the Trustworthiness of Small Online Retailers," *Journal of Interactive Marketing*, 18 (1), 53–69.
- Ward, Andrew and Lyle Brenner (2006), "Accentuate the Negative: The Positive Effects of Negative Acknowledgement," *Psychological Science*, 17 (11), 959–62.
- Wertenbroch, Klaus (2019), "From the Editor: A Manifesto for Research on Automation in Marketing and Consumer Behavior," *Journal of Marketing Behavior*, 4 (1), 1–10.
- Westin, Alan F. (1991), *Harris-Equifax Consumer Privacy Survey*. Atlanta, GA: Equifax Inc.
- White, Tiffany Barnett (2004), "Consumer Disclosure and Disclosure Avoidance: A Motivational Framework," *Journal of Consumer Psychology*, 14 (1), 41–51.
- Woolley, Kaitlin and Jane L. Risen (2021), "Hiding from the Truth: When and How Cover Enables Information Avoidance," *Journal of Consumer Research*, 47 (5), 675–97.
- Xu, Heng, Tamara Dinev, Jeff Smith, and Paul Hart (2011), "Information Privacy Concerns: Linking Individual Perceptions with Institutional Privacy Assurances," *Journal of the Association for Information Systems*, 12 (12), <https://aisel.aisnet.org/jais/vol12/iss12/1/>.
- Zeng, Fue, Qing Ye, Zhilin Yang, Jing Li, and Yiping Amy Song (2020), "Which Privacy Policy Works, Privacy Assurance or Personalization Declaration? An Investigation of Privacy Policies and Privacy Concerns," *Journal of Business Ethics* (in press), <https://doi.org/10.1007/s10551-020-04626-x>.