MIS Quarterly

RESEARCH ARTICLE

BEYOND THE PRIVACY PARADOX: OBJECTIVE VERSUS RELATIVE RISK IN PRIVACY DECISION MAKING[1]

Idris Adjerid
Mendoza College of Business, University of Notre Dame,
Notre Dame, IN 46556 U.S.A. {iadjerid@nd.edu}

Eyal Peer
The Graduate School of Business Administration, Bar-Ilan University,
Ramat-Gan 5290002 ISRAEL {eyal.peer@biu.ac.il}

Alessandro Acquisti
Heinz College, Carnegie Mellon University, 5000 Forbes Avenue,
Pittsburgh, PA 15213-3890 U.S.A. {acquisti@andrew.cmu.edu}

Privacy decision making has been examined in the literature from alternative perspectives. A dominant "normative" perspective has focused on rational processes by which consumers with stable preferences for privacy weigh the expected benefits of privacy choices against their potential costs. More recently, a behavioral perspective has leveraged theories from decision research to construe privacy decision making as a process in which cognitive heuristics and biases predictably occur. In a series of experiments, we compare the predictive power of these two perspectives by evaluating the impact of changes in the ***objective*** *risk of disclosure and the impact of changes in the* ***relative*** *perceptions of risk of disclosure on both hypothetical and actual consumer privacy choices. We find that both relative and objective risks can, in fact, influence consumer privacy decisions. However, and surprisingly, the impact of objective changes in risk diminishes between hypothetical and actual choice settings. Vice versa, the impact of relative risk becomes more pronounced going from hypothetical to actual choice settings. Our results suggest a way to integrate diverse streams of the information systems literature on privacy decision making: in hypothetical choice contexts, relative to actual choice contexts, consumers may both overestimate their response to normative factors and underestimate their response to behavioral factors.*

Keywords: Privacy, privacy paradox, behavioral economics, prospect theory, reference dependence

Introduction

Sophisticated and increasingly ubiquitous technologies centered around the collection of consumer data are keeping the public debate over privacy at center stage. Within this debate, understanding consumer privacy preferences and behaviors has been a primary consideration for both firms and policy makers. Firms whose products and services rely on individuals' personal information face complex consumer concerns. For instance, data-driven marketing approaches such as targeted ads may increase the effectiveness of advertising (Farahat and Bailey 2012), but may also scare off some consumers (Tucker 2012). Policy makers, despite a focus on aims that are distinct from those motivating firms, similarly seek an understanding of the drivers of consumers' privacy decisions, in order to design effective mechanisms for their

[1]H. Raghav Rao was the accepting senior editor for this paper. Anthony Vance served as the associate editor.

The appendices for this paper are located in the "Online Supplements" section of the *MIS Quarterly*'s website (http://www.misq.org).

protection (Adjerid et al. 2015; FTC 2012). For example, the ability of self-regulatory *transparency and control* mechanisms to adequately protect consumer privacy critically depends on which factors and processes ultimately drive consumer privacy concerns and subsequent behavior.

Unsurprisingly, a significant and active stream of research across multiple disciplines (including management and information systems, economics, marketing, psychology, and human computer interaction) has attempted to disentangle the complex dynamics of privacy decision making (Acquisti et al. 2015; Bélanger and Crossler 2011; Smith et al. 2011). Historically, much of the information systems (IS) research in this area has focused on a *normative* perspective of consumer choice. Normative theories of consumer choice are those theories consistent with the classical economic view of consumers as deliberative, utility-maximizing, rational agents who possess reasonably stable, and therefore predictable, preferences for goods (Mullainathan and Thaler 2000; Simon 1959). Under this perspective, privacy decisions can be construed as the result of a mental calculus that weighs the expected benefits of privacy allowances against their resulting costs (Dinev and Hart 2006; Klopfer and Rubenstein 1977; Milne and Gordon 1993). This perspective is supported by prior work showing that consumers react to objective differences in privacy settings, including the use of fair information practices or firms' adoption of privacy seals (Culnan and Armstrong 1999; Miyazaki and Krishnamurthy 2002).

Although the literature has demonstrated the importance of privacy risks and benefits in influencing consumer behavior (Bélanger and Crossler 2011; Marthews and Tucker 2014; Smith et al. 2011), this account of decision making faces the challenge of explaining surprising, yet robust, empirical occurrences in privacy contexts. These include the dichotomy between stated privacy attitudes, or stated privacy intentions, and actual behaviors (Jensen et al. Jensen 2005; Spiekermann et al. 2001), as well as seemingly contradictory reactions to privacy trade-offs (Brandimarte et al. 2013; John et al. 2011). For example, prior work has found that default options that are trivial to unselect may have a major impact on consumers' privacy choices (Johnson et al. 2002), and that perceived control over information disclosures, even when it does not affect objective risks, can result in significantly more disclosure by consumers (Brandimarte et al. 2013). In the last few decades, a growing body of work in economics and decision research has proposed *behavioral* perspectives on decision making (Camerer et al. 2011; Mullainathan and Thaler 2000). Such perspectives are rooted in the notion that systematic (and therefore predictable and replicable) deviations from normative (e.g., rational-calculus based) accounts of consumer decision can arise due to limitations in consumers' cognitive ability, or their susceptibility to behavioral heuristics and decision-making biases. In the privacy context, this perspective suggests that factors independent of objective trade-offs associated with privacy choices, as well as variation in consumer preferences, can still significantly influence consumers' behavior.

Both accounts of privacy decisions stem from legitimate theoretical frameworks and have stimulated considerable bodies of empirical research. However, most IS privacy research has focused primarily on either the normative *or* behavioral perspectives of privacy decision making. As a result, comparisons between the results produced within the two literatures are *post hoc*, requiring meta-analysis across studies with diverse modeling assumptions and empirical methodologies. Given the significant interest by firms and policy makers in the nature of consumers' privacy decision making, as well as the significant and growing social and economic implications of consumer privacy choices, the absence of a bridge between the two streams of work represents a considerable gap in the literature and is thus the focus of this manuscript.

The research objective of this paper is to evaluate whether these alternative perspectives on privacy decision making can individually account for some of the variation in observed consumer privacy choices. Furthermore, we explore the conditions under which a normative or a behavioral account may differentially explain privacy choices. To meet this objective, we use a series of experiments to evaluate normative and behavioral perspectives within the same empirical settings.

We focus on informational privacy and the impact different degrees of data protection can have on consumers' willingness to reveal personal information—a construct common in both streams of literature. Across three experiments, we manipulate normative factors, behavioral factors, or both simultaneously. Many manipulations of either type are available; we lean on the extant privacy and behavioral literatures to identify established manipulations within each perspective to use in our experiments. Specifically, we manipulate a normative factor by varying the objective levels of protection afforded to disclosures of personal information (e.g., whether responses are identified or anonymous). We manipulate a behavioral factor by holding the objective levels of information protection constant, but varying the relative perception of changes in privacy protection afforded to disclosures of personal information, based on seminal research on prospect theory and reference dependence (Kahneman and Tversky 1979). We capture the impact of these factors on both hypothetical (self-reported) and actual disclosure behavior. This is relevant in light of existing literature suggesting con-

sumers may both overestimate their response to normative factors in hypothetical settings due to more favorable beliefs and attitudes about protecting one's privacy (Ajzen et al. 2004), and underestimate the effect of behavioral factors when considering a choice in hypothetical settings (Liberman et al. 2004; Loewenstein and Adler 1995). We explore whether these tendencies result in normative versus behavioral perspectives having differential effects on hypothetical versus actual behavior.

Across the experiments, we find evidence that both objective and relative differences in privacy protection affect consumer privacy decisions, lending credence to the notion that both perspectives (normative and behavioral) capture aspects of privacy choice. However, we also find that *objective* privacy protections have a diminished effect on privacy decision making in contexts that involve *actual* privacy choices relative to *hypothetical* privacy choices, whereas *relative* changes in privacy protection have a pronounced effect in *actual* settings relative to *hypothetical* ones. Specifically, we find that in a hypothetical context (Experiment 1), differences in objective privacy protection result in significant differences in participants' reported privacy concerns and their willingness to disclose personal information. By contrast, changes in relative risk result in smaller (although significant) differences in reported privacy concerns, and no differences in hypothetical willingness to disclose personal information. Mirroring what participants were asked to imagine in Experiment 1, we find the opposite effect in a context with actual disclosures (Experiment 2): differences in objective privacy protection have a small effect on participants' self-disclosure, whereas relative changes in privacy protections strongly influence participants' propensity for self-disclosure. In a final experiment (Experiment 3), we consider both normative and behavioral perspectives by manipulating all dimensions simultaneously (objective protection versus relative protection, and hypothetical versus actual choice). We find results consistent with the two prior experiments: both objective and relative changes in protection can have an impact on privacy decision making; however, objective changes have pronounced effects in hypothetical settings, whereas relative changes have pronounced effects in actual choice settings.

These findings contribute to the IS literature on the drivers and predictors of privacy decision making. Although the normative perspective for privacy decision making is now well studied, the behavioral perspective is still developing. While our findings bolster the role of the nascent but growing behavioral perspective, they also provide evidence of the simultaneous (yet uneven) role of both normative and behavioral factors. Such findings have implications beyond the privacy literature. Although the study of behavioral factors (e.g., applications of reference dependence and prospect theory) continues to garner interest from the broader research community across varied contexts,[2] their application to the IS literature is growing (Herrmann et al. 2014; Keith et al. 2014; Keith et al. 2012) but still relatively sparse. This area is ripe for exploration, considering that many technology choice contexts, related and unrelated to privacy, are highly dynamic, giving prospect theory a potentially important role in understanding consumer decision making in these contexts.

Conceptual Background and Theory

A prominent focus of economic research in the last half century has been understanding the bounds of rational consumer choice and reconciling traditional neoclassical theory with an accumulating body of empirical and theoretical research supporting behavioral accounts of consumer decision making (e.g., Camerer et al. 2011; Ho et al. 2006). In more recent years, behavioral research has started informing a number of other domains, including information systems. Goes (2013), for instance, highlights the need to incorporate insights from behavioral economics into theoretical and empirical IS research.

The interplay of rational choice and behavioral accounts of decision making is particularly prominent in the context of consumer privacy choice. Most of the literature in this area has been predicated on the notion that privacy decision making is largely a rational process driven by what we may refer to as *normative factors*—that is, factors that are normal to consider if an agent is attempting to maximize her utility. Such factors may include the objective benefits and costs of information disclosure, and the agent's stable, coherent preferences (Mullainathan and Thaler 2000; Simon 1959). For instance, a *privacy calculus* view of consumer decision posits that privacy is subject to interpretation in "economic terms" (Klopfer and Rubenstein 1977) and that a systematic weighing of the benefits of information disclosures against the perceived privacy risks from such disclosures drives consumer privacy choices (Dinev and Hart 2006; Milne and Gordon 1993). Along these lines, Westin (2000) posited that most consumers are shrewd privacy balancers who weigh the value to them and society of various business and government programs calling for personal information. Relatedly, a considerable body of work has focused on identifying systematic differences in privacy concerns between consumers (e.g.,

[2]Bartling et al. (2015) recently showed that reference-dependent decision making seems to hold in the context of coaches and players of professional soccer.

Smith et al. 1996) and has suggested elevated privacy concerns correspond to privacy-seeking behavior, such as a diminished willingness to disclose personal information (Malhotra et al. 2004).

A more recent theme in the privacy literature has considered factors that ostensibly should have little (or even no) direct impact on objective risk and benefits of disclosure, but that nevertheless considerably affect people's privacy concerns and personal preferences for self-disclosure (e.g., Moon 2000). For example, people respond more honestly and with higher rates of disclosure to an online version, versus a paper-and-pencil version, of the same questionnaire (Tourangeau 2004), even though online responses are more likely to be tracked, duplicated, disseminated, shared with, or accessed by a larger number of parties than a questionnaire filled out on a single sheet of paper. Similar effects emerge when comparing online disclosures to those made during face-to-face communication (e.g., Harper and Harper 2006). Also, people seem to rely on contextual cues, such as a survey's look and feel or implicit social norms, when disclosing intimate details about themselves (John et al. 2011). Or, holding objective risk constant, the mere increase in perceived control over who can access and use online personal information can result in an increased likelihood of making sensitive, risky disclosures (Brandimarte et al. 2013). This behavioral perspective suggests consumer privacy preferences may be malleable rather than stable, and that privacy behavior is not just highly context-dependent (something that may be also predicted by rational calculus-grounded theories of privacy decision making), but can in fact be affected by factors with little relationship to changes in objective trade-offs from disclosure, such as order effects, framing, and other decision heuristics.

Much of the privacy literature has studied consumer privacy decision making by evaluating the impact of varying degrees of privacy protection and assurances on consumers' behavior, often with consumer disclosure or engagement with a commercial entity as outcomes of interest. However, most privacy IS research has focused on either the normative or the behavioral perspective of privacy decision making. In this article, we study the impact of changes in privacy protection and assurances on consumers' privacy behavior, but consider simultaneously normative and behavioral perspectives. We do so by manipulating either objective changes in privacy protection or relative changes in perceptions of protection, in two alternative contexts extensively examined in previous research: hypothetical self-disclosure choices (i.e., behavioral intentions) and actual disclosure decisions (see Figure 1). Later in this section, we highlight that support for normative and behavioral effects differs across hypothetical and actual choice.

Objective Changes in Privacy Protection

A substantial body of research suggests that changes in expected privacy benefits and risks can affect consumers' observed privacy choices. For example, disclosing personal information can lead to consumer benefits such as an improved experience in retail via customization of products, promotions, and even user interfaces (Ansari and Mela 2003), and can enable users to derive personal and economic value from social networks (Ellison et al. 2007). Similarly, the literature has noted a number of potential risks of loss due to these information disclosures, stemming from the misuse of disclosed data (Featherman and Pavlou 2003), sharing of personal information with third parties, or price discrimination (Viswanathan et al. 2007). Prior literature suggests the relationship between consumers' perceived risk and behavior is critically related to consumer trust, both in terms of its impact on consumers' perception of risk (Kim et al. 2008; Vance et al. 2008) and as a mediator explaining why shifts in privacy risk affect behavior. For example, Dinev and Hart (2006) show that variation in the perceived risk of a context affects behavioral intention both directly and through a strong effect on consumer trust and privacy concern for a particular context.

Within this general paradigm, it follows that privacy protections have the potential to influence consumer privacy decision making via their impact on the perceived risks of misuse of consumers' personal information. For example, Dinev and Hart (2006) suggest that assurances from salespeople can mitigate perceptions of risk—which, in their model, would diminish privacy concerns and increase trust, thus leading to changes in behavior. A number of studies support this conjecture by shifting the objective degree of privacy protection afforded to consumers' personal information, and find evidence for various facets of this model of consumer behavior. For instance, Culnan and Armstrong (1999) find that the use of fair information practices by firms can engender trust from consumers, reducing privacy concerns and perceived risks of disclosure. Xu et al. (2009) find that self-regulation and government regulation reduce perceived risk from participating in location-based services and increase consumers' intention to disclose personal information. Miyazaki and Krishnamurthy (2002) find a significant effect of privacy seals on consumer perception of firms' privacy practices and their stated willingness to disclose personal information. Finally, Xu et al. (2012) find that industry self-regulation and government regulation reduce consumer privacy concerns.

Using treatments of privacy assurances similar to those employed in this literature, we manipulate the objective degree

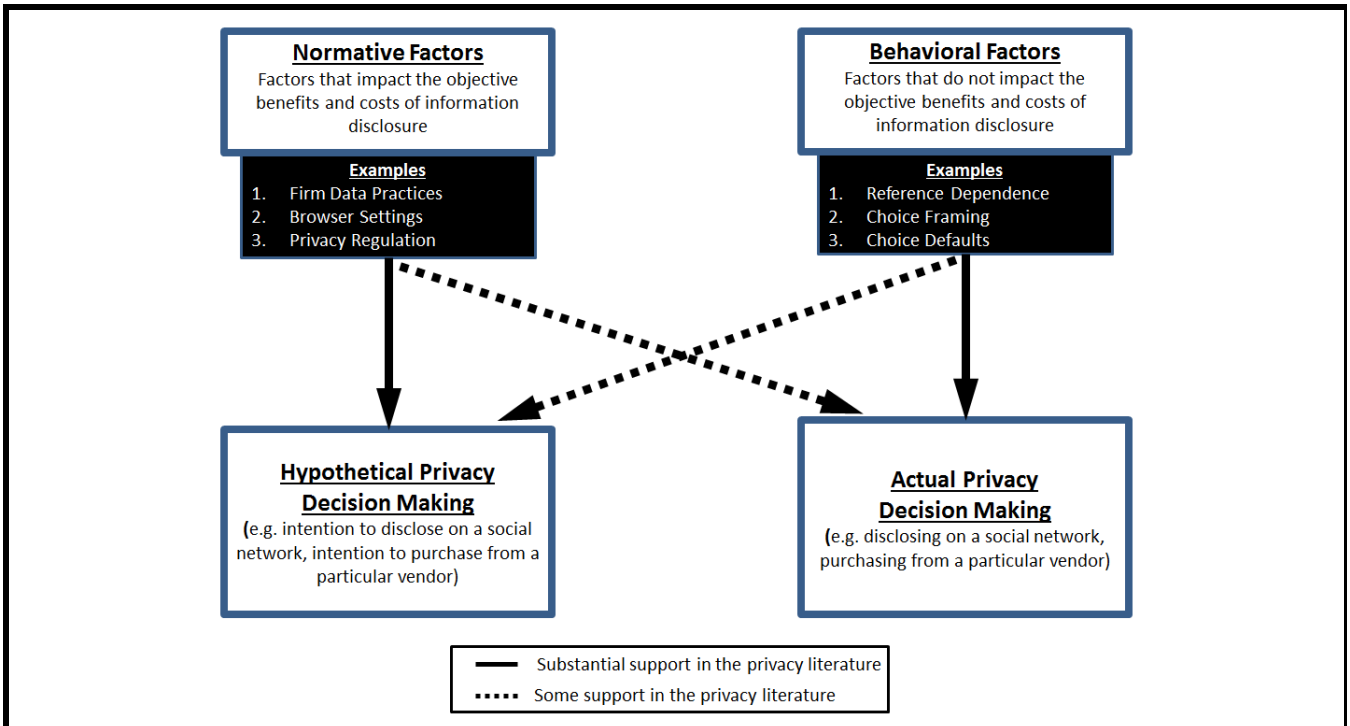

Figure 1. Research Overview

of protection afforded to consumers' disclosure (e.g., the breadth of access to personal information and the anonymity of responses) conveyed to participants via a privacy notice (similar to firms' privacy policies). We capture the role of normative factors (objective changes in risks and benefits) on privacy behavior through the following hypothesis:

> *H1: Changes in objective levels of privacy protection will affect disclosure: lower levels of privacy protection will lead to lower levels of disclosure of personal information.*

Relative Shifts in Privacy Protection

As noted in the "Introduction," a growing body of empirical evidence has suggested that behavioral factors with little or no direct impact on objective risk and benefits of disclosure can, in fact, considerably affect people's privacy concerns and personal preferences for self-disclosure. For example, decision biases and heuristics, affect, and emotions can significantly affect privacy decision making (Acquisti et al. 2015; Li et al. 2011; Li et al. 2008). The behaviorally grounded body of research on privacy has paid increasing attention to the fact that privacy judgments can be relative in nature (Acquisti et al. 2012; Egelman et al. 2013): consumers may compare their (current) situation to that of other people, or to their situation in the past, and phenomena such as habituation or coherent arbitrariness may affect their decisions. This may be particularly salient in privacy contexts where heterogeneity in data practices abounds (across firms and over time). For instance, firms that aggregate consumer privacy information often highlight improvements (i.e., relative changes) to consumer privacy over time[3] and sometimes notify consumers of their privacy protections in a manner that highlights the relative privacy gains from their services compared to those of their competitors.[4]

Prospect theory (PT), introduced by Kahneman and Tversky (1979), is a useful framework for studying the potentially relative nature of consumer privacy choices, and has been increasingly applied to privacy contexts. Using PT as a framework, Keith et al. (2012) find an inverse relationship between initial perceptions of risk and subsequent intent to

[3]Facebook, for example, has been known to advertise updates to privacy policies to highlight gains to consumer privacy (https://www.facebook.com/notes/10150251867797131).

[4]Microsoft's "Scroogled" ad campaign sought to highlight the privacy protectiveness of their services (e.g., search, email, etc.) relative to those of Google (https://en.wikipedia.org/wiki/Scroogled).

Table 1. Relative Judgments and Prospect Theory in Privacy Research

Study	Method	Participant Pool	Key Findings
Acquisti, John, and Loewenstein (2012)	Experiment	*New York Times* Readers	Individuals are more likely to disclose sensitive information if they believe others have disclosed similar information and if questions are presented in decreasing order of intrusiveness.
Keith, Thompson, Hale, and Greer (2012)	Experiment	University Students	Prospect theory explains various dimensions of privacy decision making through effects of reference dependence and hyperbolic time discounting. Keith et al. find that relative increases/decreases in perception of risk influence subsequent intentions to disclose.
Egelman, Felt, and Wagner (2013)	Experiment	Amazon Mechanical Turk	Individuals were willing to pay a higher premium for applications that requested fewer permissions when these applications were compared to other applications side-by-side.
Keith, Babb, and Lowry (2014)	Experiment	University Students	Relative increases in benefits of disclosure result in higher levels of perceived privacy risk and lower levels of disclosure. This occurs because participants become risk averse when they perceive a "gain" in benefit but risk seeking when they observe a "loss" in benefit.
Dinev, McConnell, and Smith (2015)	Research Framework	N/A	Revised the APCO model of privacy decision making to integrate insights from behavioral economics and psychology. Many of these insights relate to prospect theory and reference-dependent judgments, including loss aversion, message framing, and endowment effects.

disclose: those who perceived their risks to be already high were more likely to share personal information in a subsequent disclosure setting. Their reasoning is that those who perceive their privacy to already be lost may perceive less new risk from an additional data request. In addition, Keith et al. (2014) find that relative increases in the benefit of disclosure surprisingly diminish, rather than increase, the likelihood of sensitive disclosure. Finally, the revised APCO (Antecedents→Privacy Concerns→Outcomes) model proposed by Dinev et al. (2015) leans on insights from PT to introduce behavioral dimensions of privacy decision making. Table 1 provides examples of privacy research that focuses on relative privacy judgments and PT.

Of particular relevance to our work is PT's proposition that consumers evaluate outcomes both with respect to objective levels of consumption and with respect to a reference point, treating outcomes above or below the reference point as gains or losses, respectively. In other words, PT allows for the impact of both the objective features of a particular choice context that should influence choice (e.g., price of a product) and the features of a particular context that, according to classic accounts of economically rational decision making (e.g., von Neumann and Morgenstern 1944), should not have an impact on behavior. Therefore, PT offers a framework for analyzing and comparing the impact of normative versus behavioral factors on privacy choice, that is, how objective changes in privacy risk can affect privacy decision making, but also how changes in relative perceptions of privacy risk, in the absence of changes in objective risk, can still affect consumer privacy decision making.

Considerable empirical evidence supports the notion of reference-dependent decision making, and rules out alternative rational explanations of reference dependence (e.g., lack of information or consumer inexperience with a choice context). For example, Kahneman and Tversky (1979) found that individuals are much more likely to accept a gamble when the choice is framed as avoiding a loss compared to when the objectively equivalent choice is framed as obtaining a gain. Moreover, seminal work on the endowment effect (e.g., Kahneman et al. 1991) highlights significant differences in the amount buyers are willing to pay (WTP) for an item compared to the amount sellers are willing to accept (WTA) for the same item. Such a WTA–WTP gap has been attributed to the difference between buyers' and sellers' reference points: whereas buyers consider the purchase of a new item as a gain, sellers consider it as a loss (e.g., Novemsky and Kahneman 2005). A similar WTA–WTP gap has also been found to operate in the context of disclosure decisions (Acquisti et al. 2013). In fact, Kőszegi and Rabin (2006) incorporate reference dependence in classical models of consumer utility, allowing for consumer utility to be derived from both objective features of a choice set and also deviations

from a reference point. Simply put, PT offers a theoretically and empirically validated framework that offers defensible deviations from normative models of decision making, including those that relate to normative perspectives of privacy decision making.

We use insights from PT and the empirical literature on reference dependence to evaluate the impact of relative changes in privacy protection on privacy decision making. Under normative perspectives, identical privacy notices should result, on average, in comparable levels of disclosure irrespective of relative changes in privacy notices. However, under an alternative account of decision making that incorporates reference dependence, consumers would evaluate privacy notices relative to their deviation from a reference point, such as the level of protection they had in the recent past or the one they currently use (i.e., the status quo). We capture the role of behavioral factors (relative changes in risks and benefits) on privacy behavior through the following hypothesis:

> *H2: One's relative perception of the level of privacy protection will influence individual privacy decision making: levels of privacy protection perceived to be higher relative to a reference point will result in higher levels of disclosure of personal information.*

Privacy Behavior in Actual Versus Hypothetical Choice Contexts

Given that compelling accounts and empirical evidence exist for both normative and behavioral perspectives on privacy decision making, a consideration of which factors may moderate the effect of objective and relative privacy protection on privacy decision making is useful. One such factor could be whether privacy decision making is being studied in hypothetical settings (and captured in the form of attitudes or behavioral intentions) or in behavioral settings (and captured in the form of actual behaviors and choices). Smith et al. Xu (2011, p. X) note that "it is quite common for researchers to measure stated intentions instead of actual behaviors" in the extant literature. On the other hand, the behavioral literature cited previously predominantly uses actual choice as the outcome of interest (see Acquisti et al. 2012; Brandimarte et al. 2013; Egelman et al. 2013; John et al. 2011). Of course, one possibility is that both normative and behavioral factors influence behavior to the same degree in hypothetical and actual choice settings, rendering this empirical distinction between the literatures inconsequential. Alternatively, normative and behavioral factors may play different roles when moving from hypothetical to actual choice settings. These differing effects can emerge if consumers either misjudge the impact of normative factors between hypothetical and actual choice settings, if they misjudge the impact of behavioral factors between hypothetical and actual choice settings, or, of course, both.

We first consider the potential of consumers to misjudge their reaction to normative factors, moving from hypothetical to actual choice settings. In the psychology literature, *hypothetical bias* refers to a divergence between behavioral intentions in hypothetical contexts versus actual behavior in real-life settings (LaPiere 1934; Murphy et al. 2005). The initial observation of this bias is often attributed to LaPiere's (1934) classic study on racial prejudice, in which 92% of more than 250 surveyed service establishments responded that they would not accept members of the Chinese race, whereas in practice, 95% of the same establishments admitted and provided service to a Chinese couple without hesitation.[5] Decades after LaPiere's seminal study, the literature is still investigating this bias. Ajzen et al. (2004) find that individuals significantly overstate their propensity to donate to a scholarship fund in hypothetical versus actual choice settings. Prior work also finds that individuals overstate their propensity to use condoms, undergo a cancer screening, or exercise in hypothetical relative to actual choice settings (Sheeran 2002). FeldmanHall et al. (2012) find that subjects say they will give up more money to spare others from mild electrical shocks than they actually do when the shocks are real. This body of work suggests that hypothetical choice settings can be substantively different from their actual choice counterparts. LaPiere's study, for instance, demonstrates that these behaviors measured in hypothetical contexts need not even correlate with their real-world instantiations.

Researchers have explained the intention–behavior gap by the activation of more favorable beliefs and attitudes in hypothetical versus actual choice settings. For example, Ajzen et al. found that participants in the hypothetical choice setting were significantly more likely to indicate that donating to the fund was a social norm (e.g., more likely to indicate they *should* contribute and that those close to them would do the same). They also found that hypothetical settings elicited a more favorable attitude toward donating (e.g., more likely to indicate that the behavior was good rather than bad, right rather than wrong, etc.). If choice contexts that involve protecting one's privacy carry similar dynamics, consumers' hypothetical evaluations of high versus low levels of privacy protection (i.e., our normative manipulations) may be affected by overly positive attitudes related to protecting one's privacy, as well as elevated perceptions of it being a social norm. If

[5] At the time of this study, denying service to Chinese customers would have been viewed as the socially desirable response from these establishments.

these positive attitudes and perceptions of others' behavior do not carry over to actual choice settings (as the literature would suggest), we may expect consumers to overstate their response to objective changes in privacy protection in hypothetical settings relative to actual ones. This line of argument is formalized in the following hypothesis:

> *H3: The impact of normative factors (i.e., objective changes in privacy protection) will be stronger on hypothetical intentions to disclose compared to actual disclosures.*

Simultaneously, the impact of behavioral factors may vary when moving from hypothetical to actual behavior. The behavioral economics literature evaluates behavioral factors across hypothetical and actual choice settings and finds that, at a minimum, we can expect behavioral factors to have an impact in actual choice settings. Knetsch et al. (2001) found that the endowment effect (people's tendency to give higher valuations when they are selling rather than buying a good) is robust to repeated trials in actual choice settings. Lichtenstein and Slovic (1971) found evidence for preference reversals that are consistent with behavioral models of choice using actual behavior in a casino (e.g., people preferring gamble A to B when asked to choose one, but at the same time being willing to pay more for gamble B). Pommerehne et al. (1982, p. X) conclude that even "when the subjects are exposed to strong incentives to make motivated, rational decisions, the phenomenon of preference reversal does not vanish."[6] Van den Assem et al. (2012) confirm this finding and show that behaviors found in the lab, such as cooperation in a prisoner's dilemma game, are reproduced when very high stakes (e.g., 20,000 GBP) are involved, as was found among contestants of a game show called "Split or Steal."

Several works have, in fact, suggested that the impact of behavioral factors may be even more pronounced in actual choice settings relative to hypothetical ones. For instance, prior work has documented a cold–hot empathy gap where consumers in a "cold state" (when they are not influenced by a visceral driver, such as anger or hunger) have difficulty anticipating their behavior when they are in a "hot state" (when they are impacted by the same visceral driver). This results in consumers being unable to anticipate the actual choices they will make in future hot states when considering the same choice context hypothetically (Kang and Camerer 2012; Loewenstein 2000). More so, Loewenstein and Adler (1995) find that participants consistently underestimate the impact of being given an item on their subsequent valuation of that item (i.e., the endowment effect); O'Donoghue and Rabin (2000) find consumers can be naïve in their estimation of their own susceptibility to an immediate gratification bias (i.e., time-inconsistent discounting); and Liberman et al. (2004) find that participants grossly underestimate the impact of subtle framing changes to the labels of choice contexts on their subsequent behavior. Kühberger et al. (2002) examine differential effects of behavioral factors between hypothetical and actual choice settings using positive versus negative framing manipulations (manipulations rooted in PT) and find an "economic anomaly" in that "real decisions with large amounts do not diminish the framing effect; it may even be stronger than is apparent from hypothetical decisions" (p. X).

Translated to privacy contexts, those findings suggest that consumers may fail to anticipate their hot state or susceptibility to behavioral factors (e.g., how privacy choice contexts are framed) when considering hypothetical disclosures relative to actual ones. In sum, this work suggests that behavioral factors may have a pronounced effect on actual relative to hypothetical choice. Formally, we posit the following hypothesis:

> *H4: The impact of behavioral factors (i.e., relative changes in privacy protection) will be weaker on hypothetical intentions to disclose compared to actual disclosures.*

Methods

In three experiments, we evaluate the role of objective changes in privacy protection and of relative judgments of privacy protections on participants' hypothetical and actual privacy-sensitive behaviors (self-disclosures and selections of privacy settings). An overview of our three experiments can be found in Table 2.

Experiment 1 is a hypothetical study in which participants are presented with questions of a personal and sensitive nature, and graphical privacy notices are used to manipulate either the objective protection or the *relative* perception of privacy protection afforded to the answers provided by subjects; participants are asked to report both their privacy concerns and their hypothetical disclosure behavior. In this experiment, we are able to test H1 and H2 for hypothetical choice settings, as well as provide a starting point for evaluating H3 and H4.

[6] Relatedly, the experimental economics literature has questioned the use of hypothetical choice settings, although on different grounds: the absence of real economic trade-offs in hypothetical settings is seen as likely to produce behavioral intentions that may not match actual behavior, because individuals have less to lose (Conlisk 1996; Smith 1991). This concern has led some authors to suggest that nonnormative models of behavior may emerge in hypothetical choice settings but may wane in actual choice settings (Plott and Zeiler 2005). However, this conclusion has been critiqued and challenged in the literature (Fehr et al. 2015; Isoni et al. 2011).

Table 2. Overview of Experiments

Experiment	Dependent Variable	Participant Pool	Number of Participants	Purpose
1	Hypothetical Disclosure	Amazon Mechanical Turk	221	Evaluate the effect of objective and relative changes in privacy protection on hypothetical disclosure choices.
2	Actual Disclosure	Amazon Mechanical Turk	415	Evaluate the effect of objective and relative changes in privacy protection on actual disclosure choices.
3	Hypothetical+Actual Disclosure	Prolific Academic	739	Simultaneously evaluate the impact of objective and relative changes in privacy protection on both actual and hypothetical disclosure choices.

For example, if objective changes in protection have no effects on hypothetical behavior, diminishing effects of objective changes in protection (H3) cannot be supported. Similarly, if relative changes in protection have strong effects on hypothetical choices, observing even stronger effects in actual choice settings (H4) will be unlikely. Experiment 2 complements Experiment 1 by using a similar experimental context but evaluating H1 and H2 in an actual choice context. Thus Experiment 2 allows us to further evaluate H3 and H4 by analyzing whether support for H1 and H2 differs in actual choice settings relative to what we observed in the hypothetical choice setting of Experiment 1. Although this comparison is useful, it is still imperfect, because other factors are altered between Experiments 1 and 2 (e.g., form of the privacy notices, time of data collection, etc.). Experiment 3 addresses this issue by manipulating simultaneously changes in the objective and relative levels of privacy protection (using the same type of privacy notice), and capturing both hypothetical and actual behavior. As a result, Experiment 3 allows us to more conclusively evaluate how support for H1 and H2 differs between hypothetical and actual choice contexts (H3 and H4).

In addition to how the experiments complement each other in the evaluation of our hypotheses, they also include various design refinements that bolster our results. For example, we vary the study population between experiments, to ensure the findings are not idiosyncratic to a specific sampling population. In addition, changes in the design of Experiments 2 and 3 allow us to separately identify the effect of relative decreases and relative increases in privacy protection. Finally, in Experiment 3, we address potential confounds associated with previous levels of self-disclosure by having participants perform a task unrelated to privacy in lieu of providing initial disclosures. Including a non-privacy filler task better reflects real work scenarios where individuals may not face two different privacy situations in rapid succession.

We measure self-disclosure by capturing participants' answers to questions of a personal and sensitive nature. Prior research has successfully used this approach to examine privacy-sensitive behaviors (e.g., Acquisti et al. 2012; Moon 2000). Although the ostensible goal of the studies is to investigate participants' engagements in various behaviors (e.g., "Have you ever looked at pornographic material?"; see examples in Appendix B), we are not interested in those behaviors per se, but rather in whether participants are willing to disclose information about engaging in them. Because all of our experiments use random assignments to the different conditions, we can assume the distribution of participants' actual past engagement in these activities to be similar across conditions. Thus, higher or lower admission rates across conditions signal an impact of the manipulation on self-disclosure levels.

A strong rationale for this method is its ability to circumvent significant obstacles in obtaining actual self-disclosures from participants. First, it avoids divulging that the goal of the study relates to privacy and self-disclosure, allowing participants to act more naturally without being primed by experimenter demand effects. Second, simply asking participants to disclose sensitive information like SSNs or health information may have legal or ethical implications and is difficult to interpret, because we cannot ascertain when consumers choose to mask information or disclose it (i.e., we cannot tell the difference between a fake and an actual response if they provide a nine-digit number for their SSN). Alternatively, participants with privacy concerns are unlikely to admit to a behavior when they haven't engaged in it (i.e., mask their lack of engagement in a sensitive activity). This fact leaves admissions in our context as the complement of the sum of (1) people who did not engage in the behavior (which we assume to be similar across conditions) and (2) those who did not admit to the behavior although they engaged in it. John et al. (2011) note that this approach is conservative because the impact of

a given experimental manipulation "has to rise above the noise (error variance) produced by differences in true rates of engaging in the behavior across conditions" (p. X).

Our participants were sampled from two different but complementary online participant pools, allowing us to test the robustness of the findings. We recruited participants for Experiments 1 and 2 from Amazon Mechanical Turk (AMT). Prior research has validated AMT samples as at least as representative as other Internet samples, and significantly more representative than student samples (Buhrmester et al. 2011); furthermore, central findings in IS and the decision sciences have been replicated using AMT samples (Goodman et al. 2013; Steelman et al. 2014). Moreover, AMT offers an effective payment and reputation management system that offers researchers the ability to only sample participants of higher quality (research has shown that targeting these participants ensures high data quality; see Peer et al. 2014). In Experiment 3, we use another crowdsourcing platform called Prolific Academic, which is similar in most respects to AMT, except that participants on that platform only participate in academic research (whereas AMT also offers commercial uses). Woods et al. (2015), in comparing Prolific Academic to AMT, found that the two platforms offered many of the same features but that Prolific Academic had a smaller, but growing (1,000 new participants a month at the time of their study), user base compared to AMT. Peer et al. (2017) found, in two large-scale studies, that Prolific Academic produces high data quality (in terms of participants' attention, reliability, and reproducibility) comparable to that of AMT, and that Prolific Academic participants are more naïve, honest, and represent more diverse populations.

Estimation Approach

Across the three experiments in this manuscript, we evaluate the impact of randomized manipulations on non-repeating dependent variables (e.g., measures of privacy concerns) and repeated measures of information disclosure where a single participant is asked to make a series of hypothetical or actual disclosure decisions. For non-repeated measures, we evaluate the impact of our randomized treatments using the appropriate statistical tests for our variable of interest (e.g., t test, chi-square test, etc.). To be conservative, we use two-sided tests to conduct all evaluations of significance. Our evaluation of participants' disclosure behavior relies on comparably more complex tests. Because participants across all experiments were presented a series of questions (asking them either to predict their propensity to make, or to actually make, sensitive disclosures), we observe multiple, correlated responses from each single participant. As a result, we use a random-effects linear-regression model to evaluate differences in average disclosure between conditions.[7] This model accounts for the correlation between responses from a given participant when estimating the variance-covariance matrix of the coefficients, assuming constant correlation ρ between any two answers within a participant (exchangeable correlation structure: Liang and Zeger 1986). Specifically, we estimate the following general model:

$$Disclosure_{ij} = \beta * Treatment_i + \delta * \boldsymbol{X}_j + \alpha * \boldsymbol{Y}_i + \theta_i + u_{ij}$$

$Disclosure_{ij}$ measures a participant's predicted or actual propensity to disclose sensitive information or admit to sensitive behavior, i = (1,…,N participants), and j = (1,..,k questions). In some specifications, we also include $\boldsymbol{X}_j$, a vector of controls for different features of the questions participants are asked; for instance, $Intrusive_j$ controls for questions that differ in their intrusiveness. $\boldsymbol{Y}_i$ is a vector with controls for participant-specific characteristics (e.g., age and gender). θ_i is the participant-specific random effect, and u_{ij} is the error term. Estimates on randomly assigned treatments ($Treatment_i$) are unbiased because they should be uncorrelated with observed ($\boldsymbol{X}_j$, $\boldsymbol{Y}_i$) and unobserved (θ_i) individual differences and the error term u_{ij}. Although our controls are not necessary for the unbiased estimation of the effect of our treatments on disclosure behavior, we include them in some specifications to rule out any breaks in randomization, and to account for some of the variation in disclosure behavior between participants.

Experiment 1

In Experiment 1, we manipulated, between subjects, either changes in objective levels of privacy protection, or changes in perceived levels of privacy protection (increase or decrease over time) while actually holding objective privacy levels constant. We used hypothetical willingness to disclose as our key dependent variable.

Participants

A total of 221 participants from AMT (M_{female} = 37.56%; M_{age} = 29.16, SD_{age} = 9.76) completed the study and were each paid $0.30. We recruited participants from AMT who resided

[7] We use a linear probability model estimation in lieu of a nonlinear estimation approach (e.g., logit) for the straightforward interperation of regression coefficients and the flexibility of OLS in analyzing both Likert-scale-dependent variables and binary outcomes. Angrist and Pischke (2008) have shown little qualitative difference between the logit and linear probability specification.

in the United States and had completed at least 500 previous tasks (HITs) on AMT with an approval rate of at least 95%. These criteria are in line with what the prior literature has suggested for generating high quality data from AMT (Peer et al. 2014). The study was advertised as taking 2–5 minutes (the average participant completed the study in about 4 minutes).

Design and Procedure

Participants were asked to provide their personal opinions regarding two surveys our research group was ostensibly planning to conduct. Participants were told our research group conducts surveys that include sensitive questions on ethical behavior, and that the confidentiality protections for these surveys can vary depending on the study. Specifically, participants were asked for their opinions regarding two surveys called Survey A and Survey B. First, participants were given a description of Survey A, including its level of privacy protection. Protection levels were described using a figure that showed, on five parameters, the degree to which participants' privacy would be protected during the study—for instance, whether certain identifying information would or would not be required (and possibly linked to the answers provided by the participants), or whether the survey offered a particular protection to the answers the subjects provided. In one condition, the first survey (Survey A) provided a low overall privacy-protection level with the "Less Protective" option for four of the five parameters described (see Figure 2a), and in the other condition, the survey provided a high privacy-protection level with the "More Protective" option for four of the five parameters described (see Figure 2b). All other details of the survey (length, purpose, and payment) were the same in both conditions. Participants were then asked a set of questions that confirmed they had evaluated and understood each dimension of the notice provided (e.g., "Are responses kept after the study ends?"—see Appendix A). As a manipulation check, participants were then asked to report their satisfaction with the protections provided in each survey, their perception of potential harm from disclosure in the study, and their concerns about their privacy (see Appendix A). Finally, participants were asked questions gauging their hypothetical willingness to disclose descriptive but sensitive information (e.g., address or phone number),[8] and how often had they engaged in a set of potentially sensitive or even unethical behaviors (see Appendix B). Similar to the extant literature using hypothetical or intended behavior (e.g., Dinev and Hart 2006; Xu et al. 2009), we used five-point scales ranging from very likely to very unlikely in order to measure participants' behavioral intention to disclose this information.

Next, all participants proceeded to review a second survey (Survey B), which provided a relatively medium privacy level in both conditions (see Figure 2c). Participants were asked to evaluate Survey B using the same questions as used for Survey A. The level of protections afforded in Survey B was designed so that participants in the first condition would perceive an *increase* in the privacy level from Survey A to Survey B (low to medium), whereas participants in the second condition would perceive a decrease in the privacy level from Survey A to Survey B (from high to medium). We used participants' perception of privacy concerns, potential harm to them, and satisfaction with protections in Survey A and B to evaluate whether they indeed perceived a decrease or increase in protections between conditions. Notably, the actual privacy notice provided in Survey B was identical for both conditions, although the subjective perception of the level of that survey's privacy might have changed.

Results

We found that participants, by and large, were able to accurately understand the notices provided in the study. For Surveys A and B, 91.85% and 94.57% correctly recalled at least four of the five dimensions. We also found that our manipulation of objective risk using a high and low protection notice (Figure 2a and 2b) was effective in influencing the perception of privacy protection in the first survey (Survey A): participants who were provided high protections reported being significantly more satisfied with those protections ($M_{High} = 3.36$, $M_{Low} = 1.56$), $t\,(219) = 12.15$, $p < .001$, $d = 1.64$, significantly less concerned about privacy ($M_{High} = 2.39$, $M_{Low} = 3.87$), $t\,(219) = -12.15$, $p < .001$, $d = -1.64$, and significantly less concerned that harm would come to them as a result of disclosing personal information ($M_{High} = 2.86$, $M_{Low} = 4.02$), $t\,(219) = -7.46$, $p < .001$, $d = -1$ (see Table 3).

We used a random-effects linear-regression estimation approach to evaluate the impact on participants' predicted disclosure behavior of differences in objective privacy protection. Participants reported their likelihood of disclosure for a given question on a five-item scale (1 = "Very Unlikely" to disclose, 5 = "Very Likely" to disclose). We found that the objective differences in privacy levels in Survey A had a significant effect on participants' predicted behavior. Participants given the low privacy level predicted being significantly less likely ($\beta_{Low} = -.67$, $p < .01$) to disclose personal informa-

[8] These questions were not used in actual disclosure settings since, as we note in our "Methods" section, validation that responses were truthful is not possible.

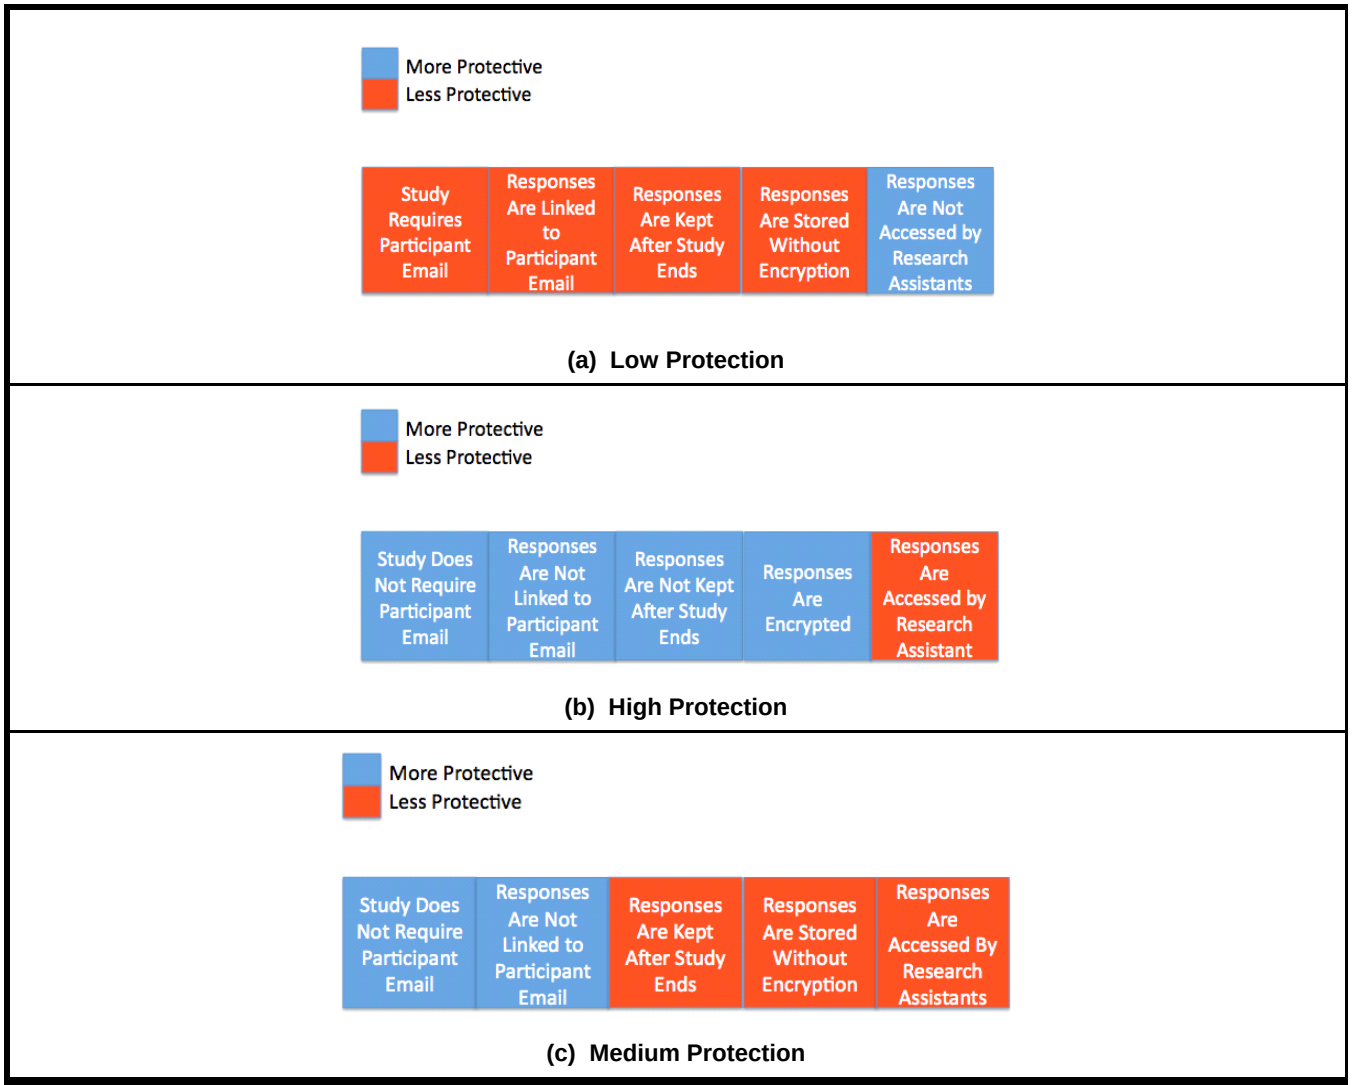

Figure 2. Levels of Privacy Protection

Table 3. Experiment 1 Summary Results

	Survey A			Survey B		
Conditions	**High Protection**	**Low Protection**	***p*-value**	**Increasing**	**Decreasing**	***p*-value**
Privacy Concern	2.39	3.87	$p < .001$	2.76	3.29	$p < .01$
Protection Satisfaction	3.36	1.56	$p < .001$	2.86	2.41	$p < .01$
Harm Perception	2.86	4.02	$p < .001$	3.37	3.68	$p = .04$

Table 4. Experiment 1 Regression Results

Variables	Admission (1 Very Unlikely – 5 Very Likely)			
	(1)	(2)	(3)	(4)
Low Protection	-0.669**	-0.650**		
	(0.120)	(0.118)		
Increasing			0.0925	0.109
			(0.123)	(0.120)
Descriptive		-0.494**		-0.565**
		(0.0607)		(0.0601)
Age		-0.0132*		-0.0100
		(0.00651)		(0.00680)
Gender		0.130		0.196
		(0.124)		(0.129)
Constant	3.631**	4.173**	3.328**	3.772**
	(0.0701)	(0.229)	(0.0784)	(0.249)
Observations	2,210	2,210	2,210	2,210
Number of id	221	221	221	221

Robust standard errors in parentheses; **p < 0.01, *p < 0.05, +p < 0.1.

tion (Table 4, column 1*)*. Moreover, we found consistent results (β_{Low} = -.65, p < .001) when including controls for question type (descriptive versus ethical) and participants' age and gender (Table 4, column 2). Broadly, these results provide strong support for the hypothesis that objective risk will affect consumer privacy choices (H1 supported).

For the second survey (Survey B), which had an objectively identical medium privacy level (Figure 2c) for both conditions, we found that participants in the increasing-protection condition reported being significantly more satisfied with the protections provided (M_{Inc} = 2.86, M_{Dec} = 2.41), t (219) = 2.97, p < .01, d = 0.40, less concerned about privacy (M_{Inc} = 2.76, M_{Dec} = 3.29), t (219) = -3.48, p < .01, d = -0.47, and less concerned that their responses might be used in ways that could harm them (M_{Inc} = 3.37, M_{Dec} = 3.68), t (219) = -2.04, p = .04, d = -0.28. However, the relative change in privacy protection in Survey B did not have a significant effect on participants' predicted disclosure behavior. Specifically, we found that increasing privacy protection did not have a significant effect ($\beta_{Increasing}$ = .09, p =.451) on overall predicted disclosure levels (Table 4, column 3). This result is robust ($\beta_{Increasing}$ = .11, p = .363) to including controls for question type and participant age and gender (Table 4, column 4). In this hypothetical disclosure setting, our results suggest a lack of support for the hypothesis that the relative perception of privacy protection will influence behavior (H2 not supported).

Discussion

The results of Experiment 1 suggest that differences in both objective and relative risk had some effect on participant perceptions of protection in the study, but only objective changes in risk influenced predicted levels of self-disclosure decisions (H1 supported). By contrast, we did not find differences in predicted levels of self-disclosure after changes in relative risk (H2 not supported).

Experiment 1 focused on hypothetical elicitation of privacy choices, which, as discussed previously, may activate choice processes distinct from actual behavior. Although this study captured only hypothetical choice, it provided some initial evidence consistent with our conjecture in H3 and H4 that the impact of normative factors may be pronounced in hypothetical settings, whereas the impact of behavioral factors may be diminished in hypothetical settings. The evidence provided by Experiment 1 in support of H3 and H4 is, of course, limited, because we cannot determine whether impacts of objective and relative risk will change or stay constant when shifting to actual self-disclosures. Moreover, Experiment 1 used a graphical representation of privacy-protection levels, including a key that alerted participants to riskier uses of their personal information. However, privacy protections online are often communicated in text-based notices, where changes in protection may not be as salient. Finally, the design of

Experiment 1 did not allow us to identify the distinct effect of relative increases and decreases in the privacy level. For example, our results might have been driven by decreases in the privacy level, and increases in the privacy level may not have had an impact (or vice versa). We address these issues in Experiment 2.

Experiment 2

Experiment 1 focused on how objective and relative privacy protection influence participants' hypothetical disclosures. In Experiment 2, we examined the role of objective and relative changes in privacy protection on actual disclosures. In addition to evaluating H1 and H2 in actual choice settings, Experiment 2 complements Experiment 1 in other ways. First, we use text-based privacy notices. Second, the experimental design allows us to evaluate the unique impact of relative increases and decreases in privacy protection as well as objective changes. Specifically, we asked participants in Experiment 2 to take part in two separate surveys about their personal behaviors. Similar to Experiment 1, each survey provided different stated levels of privacy protection to participants. Between participants, we kept the objective level of privacy offered by the surveys the same (and used a simple text-based privacy notice), and manipulated whether participants experienced a relative increase or decrease in privacy-protection levels. We examined the effects of such changes on actual disclosure behavior. By including accompanying control conditions in which protections did not change, we were able to isolate the specific impact of increases and decreases in privacy levels.[9]

Participants

A total of 415 participants from AMT (51.61% females, M_{age} = 31.27, SD_{age} = 10.72) completed the study online. The experiment was advertised to participants as two ostensibly unrelated surveys on (un)ethical behavior.[10] We recruited participants from AMT who resided in the United States and had completed at least 500 previous tasks (HITs) on AMT with an approval rate of at least 95%. We used a custom code to ensure participants who completed Experiment 1 could not take part in this experiment (Peer et al. 2012). Participants were paid $.25 cents for each survey (resulting in $.50 payment if both surveys were completed) and each survey was advertised to take 2–5 minutes. The average participant took about 3 minutes to complete each survey.

Design and Procedure

Experiment 2 consists of a 2 (high versus low protection in the first survey) × 2 (high versus low protection in the second survey) between-subjects design. Thus, the study consisted of four groups of participants whose privacy protection *increased* from the first to the second survey (low protection to high protection: LH), *decreased* (high protection to low protection: HL), or stayed the same (low to low protection: LL, or high to high protection: HH).

In the first survey, we asked participants demographic questions, including their email address. Participants were told we would check the validity of their email address prior to approving payment for the study (we did not actually store email addresses). Then, we provided participants with a privacy notice concerning the way their answers to the questionnaire would be stored. To more closely model privacy protections in real-world contexts, we presented participants with text notices (as opposed to the graphical notices presented in Experiment 1) focusing on whether their responses would be identified or anonymous (see Appendix C for full text of notices provided). Specifically, we informed participants offered "low" protections that their answers would be linked to their email address. Conversely, we informed participants offered "high" protection that their answers would not be linked to their email address.[11] We then presented participants with six questions relating to ethically questionable activities (see Appendix E for full set of questions). The questions included a subset of the questions that had been judged in Acquisti et al. (2012) as highly intrusive (e.g., "Have you ever had sexual desires for a minor?").

Thereafter, participants were asked to complete an additional survey that followed the same structure as the first survey but had a different visual design, consistent with the idea that participants were asked to participate in two separate studies (see Appendix D). Also, we provided participants two separate confirmation codes to submit in order to receive payment for completing both surveys.[12] In the second survey, we

[9]Early analysis of Experiment 2 was included in a short paper focused on the effect of privacy notices, published as part of the ACM proceedings from the 2013 Symposium on Usable Privacy and Security (SOUPS).

[10] Participants in Experiment 1 were not able to participate in Experiment 2.

[11]In Experiment 1, participants commented in a free-text exit question that they were most concerned about the propensity of a study to require them to provide email addresses and to link their responses via their email address.

[12]Of the participants who completed the exit questions, 99.5% indicated they had participated in more than one study and 96.59% indicated differences existed between the two studies. Results do not differ if we exclude these participants.

again asked participants for their email and for demographic information. Then, they received a privacy notice concerning the way their answers to the questions would be stored. As in the first survey, the privacy notice signaled either high protection (not linking responses to emails) or low protection (responses linked to emails). Next, we presented participants with six questions, different from those in the first survey, about other ethically questionable behaviors (see Appendix E). Finally, participants responded to some exit questions that gauged their perception of whether privacy protections changed in each study (e.g., whether they increased, decreased, or stayed the same, depending on the condition) and their recall of privacy notices in both surveys (see Appendix A for text of exit questions).

Results

We found that our manipulations of high and low protection elicited the hypothesized effect, with participants in the low-protection conditions reporting significantly higher beliefs that their responses would be linked back to them (M_{Low} = .79, M_{High} = .14, t (411) = 18.81, $p < .001$, d = 1.86), relative to participants in the high-protection condition. We first evaluated the disclosure rates of participants in the first survey. We found that participants were statistically more likely to disclose (β_{High} = .05 p = .04) when they were provided with high protection in the first survey (Table 5, column 1). However, our results were not significant (β_{High} = .04 p = .07), with the inclusion of controls for question intrusiveness, the survey's visual design, and participant demographics (Table 5, column 2). In the first round, we find initial evidence in support of the hypothesis that objective risk will influence participant behavior.

We then evaluated disclosure behavior in the second survey of our experiment, in which participants were presented with increasing, decreasing, or identical protection compared to the first survey. A few participants (11%) were unable to accurately recall the privacy notices (whether protections had increased, decreased, or stayed the same from the first to the second survey) and were excluded from our second-survey analysis, leaving 368 usable responses.[13] First, we compared participants who had high protection in both surveys with participants who had low protection in both surveys. For analysis in the second round, we control for the possible impact disclosing more in the first survey has on second-survey disclosures, using *Survey1Sharing*, which ranges from a value of 0 (for participants who did not admit to any of the behaviors in Survey 1) to a value of 6 (for participants who admitted to all behaviors in Survey 1). In contrast to our results for the first survey, we found no effect of high protection versus low protection on disclosure (β_{High} = -.003, p = .9) in the second survey (Table 5, column 3). This result is robust (β_{High} = .0001, p = .99) to including controls for question intrusiveness, the survey's visual design, and participant demographics (Table 5, column 4). This finding suggests that participant sensitivity to different levels of privacy protection diminished over a fairly short period of time (i.e., in the time it took to complete the first survey), and results in mixed support for H1.

Second, we evaluated the impact of changing protection on disclosure relative to conditions in which participants did not perceive an increase or decrease (participants received objectively equivalent privacy notices). We found an increase in the propensity to disclose ($\beta_{Increasing}$ = .06, p = .04) for participants who perceived an increase in protection relative to those whose protections stayed constant. This result is robust to including controls for question intrusiveness, the survey's visual design, and participant demographics (Table 5, columns 5–6). Conversely, we found a decrease in the overall propensity to disclose ($\beta_{Decreasing}$ = -.08, p = .006) for participants who perceived a decrease in protection relative to those whose protections stayed constant (Table 5, column 7). Again, this result was robust to including controls for question intrusiveness, the survey's visual design, and participant demographics (Table 5, column 8). These results suggest that participants' relative perceptions of privacy protection had a consistent impact on disclosure behavior (H2 supported).

Discussion

As in Experiment 1, the results of Experiment 2 suggest that both objective and relative changes in privacy protection can influence participants' self-disclosure behavior—even in actual choice settings. However, compared with the results we obtained in Experiment 1 (where we used a hypothetical choice context), we observe in Experiment 2 a reversal in the prominence of effects of relative versus objective changes in protection. Specifically, we found that objective changes in the levels of protection had only a weak initial effect on disclosure (only significant at the 10% level with controls) and no effect of objective differences in privacy protection on disclosure behavior in the second survey (H1 mixed support). By contrast, *relative* changes in privacy protection had a significant impact on disclosure behavior (H2 supported). These findings suggest that seemingly minor factors such as the relative, instead of the absolute, value of privacy protection can influence participants' propensity to disclose personal information, whereas the impact of objective changes in protection on actual behavior may be more limited. Com-

[13] The pattern and significance of the results remain similar when we include these participants.

Table 5. Experiment 2 Regression Results

	Probability of Admission							
Variables	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
High Protection	0.0499*	0.0423+	-0.00336	0.0001				
	(0.0240)	(0.0231)	(0.0278)	(0.0278)				
Increasing					0.0605*	0.0604*		
					(0.0295)	(0.0292)		
Decreasing							-0.075**	-0.071**
							(0.0269)	(0.0271)
Intrusive		0.076**		-0.111**		-0.113**		-0.086**
		(0.0178)		(0.0271)		(0.0259)		(0.0288)
Age		-0.005**		0.00139		0.0032*		0.00057
		(0.0009)		(0.0016)		(0.0014)		(0.0016)
Male		0.0493*		0.0512+		0.0633*		0.0483
		(0.0232)		(0.0301)		(0.0303)		(0.0307)
Survey Design		0.0379		-0.0120		0.00729		-0.0234
		(0.0231)		(0.0300)		(0.0305)		(0.0276)
Survey 1 Sharing			0.105**	0.105**	0.095**	0.097**	0.110**	0.109**
			(0.0093)	(0.0099)	(0.0102)	(0.0105)	(0.0099)	(0.0103)
Constant	0.444**	0.525**	0.0149	0.0206	0.0408	-0.0273	0.00092	0.0265
	(0.0176)	(0.0438)	(0.0273)	(0.0693)	(0.0302)	(0.0654)	(0.0278)	(0.0700)
Observations	2,490	2,454	1,164	1,140	1,158	1,140	1,050	1,032
Number of id	415	409	194	190	193	190	175	172

Robust standard errors in parentheses; **p < 0.01, *p < 0.05, +p < 0.1.

bining results from Experiment 1 with those from Experiment 2 provides initial evidence in support of H3 and H4. In particular, results from Experiments 1 and 2 suggest that objective and behavioral factors may have differential effects in hypothetical versus actual choice settings. Specifically, the impact of objective changes in protection seems to diminish as we shift to actual choice contexts (H3 supported), whereas the impact of relative changes in protection seems to become more pronounced as we shift to actual choice settings (H4 supported).

Although the comparison of effects between Experiments 1 and 2 is informative, the conclusions we can draw from this comparison may be limited. Although many variables were kept constant across the two experiments (e.g., the sampling population, the context of the experiments, etc.), a number of differences across experiments could limit our ability to make direct comparisons across experiments. For example, the studies used different privacy notices with different types of protection and were conducted at different times. Also, both studies were conducted on AMT, introducing some concern that the effects may only persist within that population. Experiment 3 addresses these concerns and allows us to evaluate H3 and H4 more robustly.

Experiment 3

The goal of Experiment 3 was to confirm that both behavioral and normative factors can be predictors of privacy decision making (H1 and H2), while also more directly comparing their respective impacts across hypothetical and actual choice contexts in the same experimental setting (H3 and H4). Using the same experimental setting allows us to evaluate the effect of both factors in hypothetical and actual choice contexts while significantly reducing the likelihood of other potential explanations of our effect being due to differences between the two prior experiments. Furthermore, we recruited a new cohort of participants from a new online recruiting platform (Prolific Academic), to test the robustness of the effects on populations other than AMT.

We recruited participants to take part in what was advertised as two separate studies. The first study served to set the stage in terms of privacy protection, and offered participants either a high or low level of privacy protection (as in Experiment 2), but did not involve any measure of self-disclosure. We did not ask participants to disclose information during the first study, because we did not want actual levels of disclosure participants might have made beforehand to influence disclo-

sures in the second study. The second study, as in the previous experiments, offered either high or low levels of privacy protection. The experiment therefore consisted of a four-conditions between-subjects design. Along one dimension, we manipulated both actual levels of privacy protections and the relative perceptions of those protections (compared to the level of protection participants received in the previous study). This approach enabled us to examine how both actual and relative changes in privacy protections affected self-disclosure. Along the other dimension, we manipulated whether participants were asked to disclose personal information (actual choice setting), or whether they were asked to self-report how likely they would be to answer the self-disclosure questions (hypothetical choice setting). With this design, we can test the robustness of the finding that both objective and behavioral factors influence behavior, and we can directly evaluate the conjecture that objective and relative changes in privacy protection may have differential effects across hypothetical and actual choice settings.

Participants

We recruited 739 participants (51.7% males, M_{age} = 29.67, SD = 10.1) from Prolific Academic (https://www.prolific.ac) to complete the study for 1.5 GBP. We recruited participants from various countries,[14] but limited our sampling to participants who indicated their first language is English, and who had an approval rating of at least 90% on their previous assignments.[15] The study was advertised as taking 10–15 minutes, and the average participant took about 12 minutes to finish the study.

Design and Procedure

Participants were invited to complete two studies (similarly to Experiment 2) that were administered one after the other but had unrelated contexts. In the first study, participants first received instructions regarding the privacy protections afforded by the study. Participants received information about the settings of the study that signaled either a high or a low protection level for the answers provided in the study, as in Experiment 1 (see Figures 2a and 2b, respectively). Participants were then asked to rate (on a five-point Likert scale) how high or low they considered the protection offered in the study. In the low-protection condition, participants were required to provide their email address, to increase the perception that responses could be linked to identity. All participants then engaged in a filler task that separated the first survey from the second (the filler task consisted of viewing a 5-minute video clip and answering open-ended questions about it). Including a non-privacy filler task better reflects real work scenarios, where individuals may not face two different privacy situations in rapid succession.

The second study used a different look and feel (type of font, background color, etc.) than the first study. Participants first received the information about the level of privacy protection provided in the new study, which was either high or low (see Figures 2a and 2b), and were asked to rate their view of the level of protection from "very low" to "very high." Participants in the low-protection condition were again required to provide their email address. Next, participants were randomly assigned to either the "actual" or the "hypothetical" disclosure condition. In the actual-disclosure condition, participants were asked to answer five personal and sensitive questions used in Acquisti et al. (2012)—see Appendix E. Participants were asked to provide their answers on a four-point scale that ranged from "never" to "many times," with the fifth option being "I prefer not to say." In the hypothetical-disclosure condition, the questions remained the same, but participants were asked to imagine taking part in a study with a certain level of protection afforded to the answers. Similar to Experiment 1 (and again in line with measures used in extant literature), hypothetical-behavior participants were told they would be presented with a set of questions relating to (un)ethical behaviors and were asked to indicate on a five-point scale (ranging from "definitely no" to "definitely yes") their likelihood of admitting to such behaviors. Finally, participants indicated their age and gender to complete the study.

Results

In the first study, participants in the high-protection condition rated the study as offering higher privacy protection (M = 3.95 versis 2.87, SD = 0.93, 1.19, t (737) = 13.75, $p < .001$). We found consistent results for the ratings of privacy protections in the second study (M = 4.00 versus 2.57, SD = 0.81, 1.28, t (737) = 18.32, $p < .001$). We thus conclude that our manipulation worked as expected, and turn to examining the effects on actual and hypothetical levels of self-disclosure. We first focus our analysis on the participants in the hypothetical settings, where we considered participants as admitting to the behavior if they responded with either "strongly

[14] As of September 2016, about 40% of participants on Prolific Academic were from the United States and Canada, and another 30% were from the United Kingdom. For updated data, see https://www.prolific.ac/demographics?metric=54bef0fafdf99b15608c504e.

[15] We could not guarantee that participants who completed this study did not also complete the prior studies on AMT. However, these studies were conducted almost a year apart, and if we have any repeat takers, they would be randomly distributed across conditions.

Table 6. Experiment 3 Hypothetical Choice Results

	Probability of Admission		
Variables	**(1)**	**(2)**	**(3)**
High Protection	0.0878*		
	(0.0441)		
Decreasing		-0.0305	
		(0.0459)	
Increasing			0.00185
			(0.0433)
Age	-0.00154	-0.000852	-0.000954
	(0.00251)	(0.00231)	(0.00276)
Male	-0.105*	-0.107*	-0.0441
	(0.0453)	(0.0467)	(0.0441)
Constant	0.737**	0.720**	0.718**
	(0.0915)	(0.0910)	(0.0967)
Observations	950	910	915
Number of id	190	182	183

Robust standard errors in parentheses; p < 0.01, *p < 0.05, +p < 0.1.

agree" or "agree" to the question of whether they would admit to a particular behavior. We find statistically significant differences in hypothetical admission rates between those with objectively different (high vs. low) levels of protection (63% versus 53%, $t(188) = 2.01$, $p = .046$). Conversely, we do not find any significant differences in hypothetical admissions when protections are held objectively constant but decrease in relative terms (53% versus 50%, $t(180) = .669$, $p = .50$) or increase in relative terms (63% in both conditions, $t(181) = .006$, $p = .99$). These results are robust to alternative measurements for hypothetical admission, including a continuous measure (i.e., 1–5 on the Likert scale) and considering those that report being uncertain (neither agree nor disagree) as also admitting to the behavior (see Appendix G).

We confirm these results in a random-effects regression (Table 6). We find that objective differences in protection (high protection) have a significant effect in the hypothetical context (column 1), whereas the relative changes have no effect (columns 2 and 3).

Next, we consider participants in the actual-disclosure condition, where participants were shown the same privacy protections and asked the same questions as their counterparts in the hypothetical-disclosure condition. For these participants, we considered an admission as any response to our questions that indicated the participant engaged in a particular behavior at least once (similar affirmative admission rates were used in prior work; John et al. 2011). We find that objective differences (high versus low) in protection, unlike the hypothetical context, did not have a significant effect on disclosure behavior (65% versus 59%, $t(178) = 1.47$, $p = .15$). Conversely, and again in contrast to the hypothetical condition, we find that those who perceived a relative decrease in protection disclosed significantly less than those who did not perceive a change (49% versus 59%, $t(161) = -2.09$, $p = .038$). Recall that participants were provided objectively identical protections between these conditions. Finally, similar to the hypothetical context, we find the relative increases in protection did not have a significant effect on disclosure (64% versus 65%, $t(201) = -0.35$, $p = .73$), suggesting that the effect of relative increases in privacy protection increasing disclosure (identified in Experiment 2) may not be robust. These results are confirmed in our random-effects regression (Table 7). We find that the objective difference in protection (high protection) does not have a significant effect on actual behavior (column 1), whereas relative changes (specifically, a relative decrease in protection) have a strong observable effect (column 2).

Discussion

In Experiment 3, we controlled for a number of factors that varied across Experiments 1 and 2, and continued to find that H1 is supported in hypothetical choice but not in actual choice settings, whereas H2 is not supported in hypothetical settings but is supported in actual choice settings. Thus, the results bolster the findings from the previous experiments for simultaneous effects of normative (H1) and behavioral (H2)

Table 7. Experiment 3 Actual Choice Results

	Probability of Admission		
Variables	**(1)**	**(2)**	**(3)**
High Protection	0.0552		
	(0.0410)		
Decreasing		-0.108*	
		(0.0476)	
Increasing			-0.0126
			(0.0354)
Age	0.00143	0.00248	-1.24e-05
	(0.00277)	(0.00232)	(0.00192)
Male	-0.0296	-0.0826+	-0.0321
	(0.0408)	(0.0480)	(0.0366)
Constant	0.594**	0.646**	0.695**
	(0.0959)	(0.0987)	(0.0694)
Observations	895	810	1,010
Number of id	179	162	202

Robust standard errors in parentheses; ** p < 0.01, * p < 0.05, + p < 0.1.

factors on privacy decision making. They also provide more robust evidence supporting the conjecture of a diminished effect of normative factors in hypothetical relative to actual choice contexts (H3 supported), and a pronounced impact of behavioral factors in hypothetical relative to actual choice contexts (H4 supported). Comparing the Hedge's g (a bias-corrected and normalized measure of effect size across experiments; Hedges 1981) for identical treatments in hypothetical versus actual choice settings supports the insight: identical relative decreases in privacy protection in hypothetical choice contexts had a treatment effect of only .09, relative to a .33 treatment effect in analogous actual choice settings. Conversely, identical objective decreases in protection had a treatment effect of .32 in hypothetical choice settings, but a diminished treatment effect of .21 in actual choice settings. Summarizing results across our three experiments highlights our main findings (see Table 8). We consistently find evidence that both normative and behavioral factors can simultaneously influence consumer perceptions of privacy risk and actual privacy choices, but these effects may emerge differentially across hypothetical versus actual choice contexts.

General Discussion and Conclusions ■

Our work builds on the IS literature on consumer privacy decision making and the behavioral economics literature on reference dependence and relative judgment. Leaning on proposed models of reference-dependent utility, which account for both the utility from absolute levels of consumption and deviations from a reference point, we present some evidence suggesting that, in the context of privacy decision making, relative changes may have an increasingly important impact on decision making, particularly in actual choice contexts, relative to absolute or objective levels of protection provided. Our findings extend and complement the growing literature at the intersection of behavioral economics and privacy decision making. With respect to PT specifically, our results bolster findings in Keith et al. (2012) concerning the effect of relative changes in perceived privacy risks on privacy behavior. Specifically, we extend this prior work by exogenously assigning different privacy reference points, allowing us to disentangle the relative effects of perceived increases and decreases in privacy protection. This approach also allows us to address issues of endogenous selection by individuals into low or high initial levels of privacy risk. Moreover, this work complements work by Keith et al. (2014) that focuses on the role of relative changes in the perception of the benefit of disclosure on consumer privacy behavior. In our work, we hold the benefits of disclosure constant, and find that relative changes in the perception of risk can also have significant impacts on privacy choices.

Our work also has some important limitations. First, we evaluate specific manipulations of the normative and behavioral perspective, which introduces the question of whether these effects would extend to other such manipulations. In particu-

Table 8. Overview of Results

	Experiment 1	Experiment 2	Experiment 3	
	Hypothetical Choice	***Actual Choice***	***Hypothetical Choice***	***Actual Choice***
H1: Objective Privacy Protection	Supported	Mixed Support	Supported	Not Supported
H2: Relative Privacy Protection	Not Supported	Supported	Not Supported	Supported
	H3, H4 Supported		**H3, H4 Supported**	

lar, privacy decision making appears prone to a number of biases, including hyperbolic time discounting, framing effects, and a control paradox (Acquisti et al. 2015; Adjerid et al. 2016; Keith et al. 2012). We alleviate some of these concerns by using experimental treatments that vary in how they manipulate these factors (e.g., visual versus text notice) and focus on manipulations that are well rooted in their respective literatures—prospect theory, for example, is a seminal theory in the behavioral economics literature and informs numerous behavioral phenomena. Moreover, we do not claim that normative factors are always stronger predictors of behavior in a given context. Clearly, this claim is problematic because it is possible to arbitrarily alter the strength of any normative or behavioral manipulation such that one dominates in a given setting. Rather, we suggest that the effects of similar or identical manipulation of normative *or* behavioral factors differ as they are assessed in hypothetical versus actual choice contexts.

Another limitation relates to the experimental nature of the work. For example, self-selection bias may arise when participants choose not to participate when prompted to provide their email address in some of our experiments. This form of self-selection, however, would be more likely to affect individuals with *higher* privacy sensitivities, which might make the findings more *conservative* (e.g., privacy-conscious individuals might react more strongly to changes to privacy protection). Furthermore, although experimental work may have limited external validity, the tension between decision making in hypothetical settings and decision making in more realistic actual choice settings is, in some sense, part of our empirical testing. In any case, to address residual concerns in this regard, we leverage diverse online samples that provide more diverse participant pools. Nevertheless, these online samples should not be considered as completely representative, and like other experimental work, the characteristics of the countries (mostly the United States and the United Kingdom) from which we have sampled limit our results.

These limitations notwithstanding, our results also have important implications for theories of consumer privacy behavior and future privacy research. Centrally, our results suggest that the behavioral factors we evaluate may be underappreciated by consumers when, in hypothetical situations, they anticipate their future privacy concerns and actual behaviors. On the other hand, those factors may actually be more influential on, and more consistent drivers of, behaviors in contexts that involve actual privacy choices. These findings are consistent with the broader psychology and behavioral economics literature (e.g., Gilbert and Ebert 2002; Lowenstein and Adler 1995), in that they imply people might overestimate the impact of normative factors on their hypothetical behavior while underestimating the sometimes-powerful impact of decision biases on actual decision making. Our results are also consistent with and extend the growing literature on how privacy decision making may be particularly susceptible to deviations from economically rational models of decision making, by not only presenting additional evidence of these deviations, but also by starting to identify the conditions under which these effects are most likely to materialize.

Our results may start to reconcile some of the dissonance in the privacy literature and help explain early results from the privacy-paradox literature (e.g., Spiekermann et al. 2001) by substantiating a critical link in how limitations in consumer (ir)rationality may be driving the observed dissonance between consumer concerns and hypothetical behavior and actual decision making. We clarify, however, that pragmatic considerations can lead to the decision to use hypothetical versus actual choice (e.g., whether observing actual behavior in the context of interest is feasible). Therefore, our focus is not to suggest that one approach is always preferable, but simply that this distinction could be relevant to the tension between normative and behavioral models of privacy decision making. In fact, our results substantiate that both perspectives are predictive of consumer behavior across both choice settings, and may simply emerge differently between them. That said, our results suggest scholars should carefully consider research goals when designing experiments or conducting empirical evaluations in privacy settings. In particular, the assumption of constancy of either normative or

behavioral effects between hypothetical and actual choice settings is problematic. More generally, our results suggest further evaluation of behavioral factors in privacy research is warranted, particularly given the increased importance of evaluating actual privacy choices. In particular, our results have implications for recent privacy research focused on designing tools and mechanisms to nudge consumers toward more protective privacy choices (Almuhimedi et al. 2015). Specifically, our findings suggest such tools may need to be designed in very different ways if they are intended to be used prospectively (e.g., tools that set up rules for future disclosures) or if they are intended to be used when actual choices are being made (e.g., just in time notices or control mechanisms).

Disentangling these diverging perspectives has implications beyond the academic discourse on consumer privacy behaviors. For policy makers, a range of consumer privacy protections intended to aid consumers in making self-interested privacy decisions are predicated on the notion that consumers are able to consistently and predictably react to changes in normative factors within privacy contexts (e.g., the objective benefits and costs of data allowances and disclosures). These protections include what is provided via regulation that covers some subset of personal information (e.g., HIPAA in the United States, or the EU Data Directive in the European Union), and self-regulatory mechanisms (e.g., companies' privacy notices and choice mechanisms) that allow consumers to select their desired degree of privacy protection. If consumers' judgments of privacy protections in actual choice contexts are relative rather than absolute, the validity of assumptions about consumers' behavior may not hold. Specifically, policy makers' goal of protecting consumer privacy through transparency and control mechanisms may not be realized if firms choose to highlight gains and downplay losses to privacy protection over time and among their competitors. At the same time, the relative nature of consumer privacy choices could also present an opportunity for policy makers to bring attention to high-relevance privacy contexts by mandating that firms clearly highlight changes in data practices over time, including decreases in protection. This approach may be particularly effective given that, over time, relative changes in protection in our experiments influenced privacy decisions more than the objective risk participants faced.

Our results also have important implications for practitioners who deal with consumers' personal information. First, those practitioners might find that consumers' reaction to objectively high (or low) levels of protection may not be consistent if these protections are also evaluated in relative terms. For example, if privacy protection is increased from a very low (absolute) level of protection, consumers might consider that increase as a gain, even though the resulting privacy protection might still be low; and consumers might be more inclined to choose privacy protections that seem more protective (in relative terms) but actually may not be. Conversely, if the level of privacy protection is decreased from a high (absolute) level of protection, consumers might consider that decrease as a loss and be less willing to use the offered service or disclose personal information, even though the actual level of privacy has remained objectively high. Interestingly, Keith et al. (2014) find these dynamics may be reversed when relative gains and losses are perceived in terms of the disclosure benefit. Taken together, the findings of our research and those of Keith et al. (2014) suggest that relative increases and decreases in protection (i.e., the cost of disclosure) as well as relative increases and decreases in the benefits of disclosure can influence consumers' actual disclosure.

This possibility suggests practitioners and firms need to seriously consider the impact of relative judgments on various facets of consumer privacy decision making, and potentially employ different strategies if they are concerned about influencing actual or hypothetical judgments. For instance, privacy seals and privacy notices may be more effective in the hypothetical disclosure stages (e.g., advertisements or commercials for a product or service), whereas relative comparisons to other organizations (or consumers' prior level of protection) are more effective later when consumers are making actual choices.

Finally, the implications of our results may vary based on different firms' goals. Firms seeking to compete by providing consumers with additional privacy protection (e.g., Mozilla advertises the privacy protections offered by the Firefox browser)[16] may find that simply providing consumers strong assurances may have a limited impact on their actual behavior. Rather, leveraging reference points, or other behavioral factors, is important when trying to influence consumers' actual behaviors. Similarly, firms that benefit from increased disclosure and allowances by consumers may find some short-term value in presenting notices and choices as relatively more protective. However, if actual data practices violate consumer expectations for privacy, troublesome and costly privacy incidents may persist, leading to lowered trust (and less disclosure) by consumers, and potentially decreased use in the long term. Moreover, if firms highlight the privacy-protective nature of their services relative to their competitors, consumers may have an elevated expectation for privacy, which may be inconsistent with actual firm data practices. The increasingly dynamic nature of data practices over time and the heterogeneity of data practices between

[16]https://www.mozilla.org/en-US/firefox/desktop/trust/

firms suggest that the relative perception of privacy protection will continue to be an important predictor of consumer privacy decision making, and will thus have significant implications for the effectiveness of tools mandated by policy makers and the mechanisms by which firms solicit privacy-relevant choices.

References

Acquisti, A., Brandimarte, L., and Loewenstein, G. 2015. "Privacy and Human Behavior in the Age of Information," *Science* (347:6221), pp. 509-514.

Acquisti, A., John, L. K., and Loewenstein, G. 2012. "The Impact of Relative Standards on the Propensity to Disclose," *Journal of Marketing Research* (49:2), pp. 160-174.

Acquisti, A., John, L., and Loewenstein, G. 2013. "What Is Privacy Worth?," *Journal of Legal Studies* (42:2), Article 1.

Adjerid, I., Acquisti, A., and Loewenstein, G. 2016. "Choice Architecture, Framing, and Cascaded Privacy Choices," unpublished paper (available at SSRN: http://ssrn.com/abstract=2765111).

Adjerid, I., Acquisti, A., Telang, R., Padman, R., and Adler-Milstein, J. 2015. "The Impact of Privacy Regulation and Technology Incentives: The Case of Health Information Exchanges," *Management Science* (62:4), pp. 1042-1063.

Ajzen, I., Brown, T. C., and Carvajal, F. 2004. "Explaining the Discrepancy Between Intentions and Actions: The Case of Hypothetical Bias in Contingent Valuation," *Personality and Social Psychology Bulletin* (30:9), pp. 1108-1121.

Almuhimedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., and Agarwal, Y. 2015. "Your Location Has Been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, New York: ACM Press, pp. 787-796.

Angrist, J. D., and Pischke, J. S. 2008. *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton, NJ: Princeton University Press.

Ansari, A., and Mela, C. F. 2003. "E-Customization," *Journal of Marketing Research* (40:2), pp. 131-145.

Bartling, B., Brandes, L., and Schunk, D. 2015. "Expectations as Reference Points: Field Evidence from Professional Soccer," *Management Science* (61:11), pp. 2646-2661.

Bélanger, F., and Crossler, R. E. 2011. "Privacy in the Digital Age: A Review of Information Privacy Research in Information Systems," *MIS Quarterly* (35:4), pp. 1017-1042.

Brandimarte, L., Acquisti, A., and Loewenstein, G. 2013. "Misplaced Confidences: Privacy and the Control Paradox," *Social Psychological and Personality Science* (4:3), pp. 340-347.

Buhrmester, M., Kwang, T., and Gosling, S. D. 2011. "Amazon's Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data?," *Perspectives on Psychological Science* (6:1), pp. 3-5.

Camerer, C. F., Loewenstein, G., and Rabin, M. (eds.). 2011. *Advances in Behavioral Economics*, Princeton, NJ: Princeton University Press.

Conlisk, J. 1996. "Why Bounded Rationality?," *Journal of Economic Literature* (34:2), pp. 669-700.

Culnan, M. J., and Armstrong, P. K. 1999. "Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation," *Organization Science* (10:1), pp. 104-115.

Dinev, T., and Hart, P. 2006. "An Extended Privacy Calculus Model for E-Commerce Transactions," *Information Systems Research* (17:1), pp. 61-80.

Dinev, T., McConnell, A. R., and Smith, H. J. 2015. "Research Commentary—Informing Privacy Research Through Information Systems, Psychology, and Behavioral Economics: Thinking Outside the 'APCO' Box," *Information Systems Research* (26:4), pp. 639-655.

Egelman, S., Felt, A. P., and Wagner, D. 2013. "Choice Architecture and Smartphone Privacy: There's a Price for That," in *The Economics of Information Security and Privacy*, R. Böhme (ed.), Berlin: Springer, pp. 211-236.

Ellison, N. B., Steinfield, C., and Lampe, C. 2007. "The Benefits of Facebook 'Friends': Social Capital and College Students' Use of Online Social Network Sites," *Journal of Computer Mediated Communication* (12:4), pp. 1143-1168.

Farahat, A., and Bailey, M. C. 2012. "How Effective Is Targeted Advertising?," in *Proceedings of the 21st International Conference on World Wide Web*, New York: ACM Press, pp. 111-120.

Featherman, M. S., and Pavlou, P. A. 2003. "Predicting E-Services Adoption: A Perceived Risk Facets Perspective," *International Journal of Human–Computer Studies* (59:4), pp. 451-474.

Fehr, D., Hakimov, R., and Kübler, D. 2015. "The Willingness to Pay–Willingness to Accept Gap: A Failed Replication of Plott and Zeiler," *European Economic Review* (78), pp. 120-128.

FeldmanHall, O., Mobbs, D., Evans, D., Hiscox, L., Navrady, L., and Dalgleish, T. 2012. "What We Say and What We Do: The Relationship Between Real and Hypothetical Moral Choices," *Cognition* (123:3), pp. 434-441.

FTC. 2012. *Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policy Makers*, Federal Trade Commission (available at http://www.ftc.gov/os/2012/03/120326privacyreport.pdf).

Gilbert, D. T., and Ebert, J. E. 2002. "Decisions and Revisions: The Affective Forecasting of Changeable Outcomes," *Journal of Personality and Social Psychology* (82:4), pp. 503-514.

Goes, P. B. 2013. "Editor's Comments: Information Systems Research and Behavioral Economics," *MIS Quarterly* (37:3), pp. iii-viii.

Goodman, J. K., Cryder, C. E., and Cheema, A. 2013. "Data Collection in a Flat World: The Strengths and Weaknesses of Mechanical Turk Samples," *Journal of Behavioral Decision Making* (26:3), pp. 213-224.

Harper, V. B., and Harper, E. J. 2006. "Understanding Student Self-Disclosure Typology Through Blogging," *The Qualitative Report* (11:2), pp. 251-261.

Hedges, L. V. 1981. "Distribution Theory for Glass's Estimator of Effect Size and Related Estimators," *Journal of Educational and Behavioral Statistics* (6:2), pp. 107-128.

Herrmann, P. N., Kundisch, D. O., and Rahman, M. S. 2014. "Beating Irrationality: Does Delegating to IT Alleviate the Sunk Cost Effect?," *Management Science* (61:4), pp. 831-850.

Ho, T. H., Lim, N., and Camerer, C. F. 2006. "Modeling the Psychology of Consumer and Firm Behavior with Behavioral Economics," *Journal of Marketing Research* (43:3), pp. 307-331.

Isoni, A., Loomes, G., and Sugden, R. 2011. "The Willingness to Pay–Willingness to Accept Gap, the 'Endowment Effect,' Subject Misconceptions, and Experimental Procedures for Eliciting Valuations: Comment," The American Economic Review (101:2), pp. 991-1011.

Jensen, C., Potts, C., and Jensen, C. 2005. "Privacy Practices of Internet Users: Self-Reports Versus Observed Behavior," *International Journal of Human-Computer Studies* (63:1), pp. 203-227.

John, L., Acquisti, A., and Loewenstein, G. 2011. "Strangers on a Plane: Context Dependent Willingness to Divulge Personal Information," *Journal of Consumer Research* (37:5), pp. 858-873.

Johnson, E. J., Bellman, S., and Lohse, G. L. 2002. "Defaults, Framing and Privacy: Why Opting In–Opting Out," *Marketing Letters* (13:1), pp. 5-15.

Kahneman, D., Knetsch, J. L., and Thaler, R. H. 1991. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias," *The Journal of Economic Perspectives* (5:1), pp. 193-206.

Kahneman, D., and Tversky, A. 1979. "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* (47:2), pp. 263-291.

Kang, M. J., and Camerer, C. 2012. "fMRI Evidence of a Hot-Cold Empathy Gap in Hypothetical and Real Aversive Choices," unpublished paper (available at SSRN 2087824).

Keith, M. J., Babb, J. S., and Lowry, P. B. 2014. "A Longitudinal Study of Information Privacy on Mobile Devices," in *Proceedings of the 47th Hawaii International Conference on System Sciences*, Los Alamitos, CA: IEEE Computer Society Press, pp. 3149-3158.

Keith, M. J., Thompson, S. C., Hale, J., and Greer, C. 2012. "Examining the Rationality of Information Disclosure through Mobile Devices," in *Proceedings of the 33rd International Conference on Information Systems*, Orlando, FL.

Kim, D. J., Ferrin, D. L., and Rao, H. R. 2008. "A Trust-Based Consumer Decision-Making Model in Electronic Commerce: The Role of Trust, Perceived Risk, and Their Antecedents," *Decision Support Systems* (44:2), pp. 544-564.

Klopfer, P. H., and Rubenstein, D. I. 1977. "The Concept Privacy and Its Biological Basis," *Journal of Social Issues* (33:3), pp. 52-65.

Knetsch, J. L., Tang, F. F., and Thaler, R. H. 2001. "The Endowment Effect and Repeated Market Trials: Is the Vickrey Auction Demand Revealing?," *Experimental Economics* (4:3), pp. 257-269.

Kőszegi, B., and Rabin, M. 2006. "A Model of Reference-Dependent Preferences," *The Quarterly Journal of Economics* (121:4), pp. 1133-1166.

Kühberger, A., Schulte-Mecklenbeck, M., and Perner, J. 2002. "Framing Decisions: Hypothetical and Real," *Organizational Behavior and Human Decision Processes* (89:2), pp. 1162-1175.

LaPiere, R. T. 1934. "Attitudes vs. Actions," *Social Forces* (13:2), pp. 230-237.

Li, H., Sarathy, R., and Xu, H. 2011. "The Role of Affect and Cognition on Online Consumers' Decision to Disclose Personal Information to Unfamiliar Online Vendors," *Decision Support Systems* (51:3), pp. 434-445.

Li, H., Sarathy, R., and Zhang, J. 2008. "The Role of Emotions in Shaping Consumers' Privacy Beliefs about Unfamiliar Online Vendors," *Journal of Information Privacy and Security* (4:3), pp. 36-62.

Liang, K. Y., and Zeger, S. L. 1986. "Longitudinal Data Analysis Using Generalized Linear Models," *Biometrika* (73:1), pp.13-22.

Liberman, V., Samuels, S. M., and Ross, L. 2004. "The Name of the Game: Predictive Power of Reputations Versus Situational Labels in Determining Prisoner's Dilemma Game Moves," *Personality and Social Psychology Bulletin* (30:9), pp. 1175-1185.

Lichtenstein, S., and Slovic, P. 1971. "Reversals of Preference Between Bids and Choices in Gambling Decisions," *Journal of Experimental Psychology* (89:1), pp. 46-55.

Loewenstein, G. 2000. "Emotions in Economic Theory and Economic Behavior," *The American Economic Review* 90(2), pp. 426-432.

Loewenstein, G., and Adler, D. 1995. "A Bias in the Prediction of Tastes," *The Economic Journal* (105:7), pp. 929-937.

Malhotra, N. K., Kim, S. S., and Agarwal, J. 2004. "Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model," *Information Systems Research* (15:4), pp. 336-355.

Marthews, A., and Tucker, C. 2014. "Government Surveillance and Internet Search Behavior," unpublished paper (available at SSRN 2412564).

Milne, G. R., and Gordon, M. E. 1993. "Direct Mail Privacy-Efficiency Trade-Offs within an Implied Social Contract Framework," *Journal of Public Policy & Marketing* (12:2), pp. 206-215.

Miyazaki, A. D., and Krishnamurthy, S. 2002. "Internet Seals of Approval: Effects on Online Privacy Policies and Consumer Perceptions," *Journal of Consumer Affairs* (36:1), pp. 28-49.

Moon, Y. 2000. "Intimate Exchanges: Using Computers to Elicit Self-Disclosure from Consumers," *Journal of Consumer Research* (26:4), pp. 323-339.

Mullainathan, S., and Thaler, R. H. 2000. "Behavioral Economics," NBER Working Paper No. 7948, National Bureau of Economic Research, Cambridge, MA.

Murphy, J. J., Allen, P. G., Stevens, T. H., and Weatherhead, D. 2005. "A Meta-Analysis of Hypothetical Bias in Stated Preference Valuation," *Environmental and Resource Economics* (30:3), pp. 313-325.

Novemsky, N., and Kahneman, D. 2005. "The Boundaries of Loss Aversion," *Journal of Marketing Research* (42:2), pp. 119-128.

O'Donoghue, T., and Rabin, M. 2000. "The Economics of Immediate Gratification," *Journal of Behavioral Decision Making* (13:2), pp. 233-250.

Peer, E., Brandimarte, L., Samat, S., and Acquisti, A. 2017. "Beyond the Turk: Alternative Platforms for Crowdsourcing Behavioral Research," *Journal of Experimental Social Psychology* (70) pp. 212-221.

Peer, E., Paolacci, G., Chandler, J., and Mueller, P. 2012. "Selectively Recruiting Participants from Amazon Mechanical Turk Using Qualtrics," unpublished paper (available at SSRN 2100631).

Peer, E., Vosgerau, J., and Acquisti, A. 2014. "Reputation as a Sufficient Condition for Data Quality on Amazon Mechanical Turk," *Behavior Research Methods* (46:4), pp. 1023-1031.

Plott, C. R., and Zeiler, K. 2005. "The Willingness to Pay–Willingness to Accept Gap, the 'Endowment Effect,' Subject Misconceptions, and Experimental Procedures for Eliciting Valuations," *American Economic Review* (95:3), pp. 530-545.

Pommerehne, W. W., Schneider, F., and Zweifel, P. 1982. "Economic Theory of Choice and the Preference Reversal Phenomenon: A Reexamination," *The American Economic Review* (72:3), pp. 569-574.

Sheeran, P. 2002. "Intention–Behavior Relations: A Conceptual and Empirical Review," *European Review of Social Psychology* (12:1), pp. 1-36.

Simon, H. A. 1959. "Theories of Decision-Making in Economics and Behavioral Science," *The American Economic Review* (49:3), pp. 253-283.

Smith, H. J., Dinev, T., and Xu, H. 2011. "Information Privacy Research: An Interdisciplinary Review," *MIS Quarterly* (35:4), pp. 989-1016.

Smith, H. J., Milberg, S. J., and Burke, S. J. 1996. "Information Privacy: Measuring Individuals' Concerns About Organizational Practices," *MIS Quarterly* (20:2), pp. 167-196.

Smith, V. L. 1991. "Rational Choice: The Contrast Between Economics and Psychology," *Journal of Political Economy* (99:4), pp. 877-897.

Spiekermann, S., Grossklags, J., and Berendt, B. 2001. "E-Privacy in 2nd Generation E-Commerce: Privacy Preferences Versus Actual Behavior," in *Proceedings of the 3rd Conference on Electronic Commerce*, New York: ACM Press, pp. 38-47.

Steelman, Z. R., Hammer, B. I., and Limayem, M. 2014. "Data Collection in the Digital Age: Innovative Alternatives to Student Samples," *MIS Quarterly* (38:2), pp. 355-378.

Tourangeau, R. 2004. "Survey Research and Societal Change," *Annual Review of Psychology* (55), pp. 775-801.

Tucker, C. 2012. "The Economics of Advertising and Privacy," *International Journal of Industrial Organization* (30:3), pp. 326-329.

Van den Assem, M. J., Van Dolder, D., and Thaler, R. H. 2012. "Split or Steal? Cooperative Behavior When the Stakes Are Large," *Management Science* (58:1), pp. 2-20.

Vance, A., Elie-Dit-Cosaque, C., and Straub, D. W. 2008. "Examining Trust in Information Technology Artifacts: The Effects of System Quality and Culture," *Journal of Management Information Systems* (24:4), pp. 73-100.

Viswanathan, S., Kuruzovich, J., Gosain, S., and Agarwal, R. 2007. "Online Infomediaries and Price Discrimination: Evidence from the Automotive Retailing Sector," *Journal of Marketing* (71:3), pp. 89-107.

Von Neumann, J., and Morgenstern, O. 1944. *Games and Economic Behavior*, Princeton, NJ: Princeton University Press.

Westin, A. F. 2000. "Intrusions: Privacy Tradeoffs in a Free Society," *Public Perspective* (11:6), pp. 8-11.

Woods, A. T., Velasco, C., Levitan, C. A., Wan, X., and Spence, C. 2015. "Conducting Perception Research Over the Internet: A Tutorial Review," *PeerJ* (https://peerj.com/articles/1058/).

Xu, H., Teo, H. H., Tan, B. C., and Agarwal, R. 2009. "The Role of Push-Pull Technology in Privacy Calculus: The Case of Location-Based Services," *Journal of Management Information Systems* (26:3), pp. 135-174.

Xu, H., Teo, H. H., Tan, B. C., and Agarwal, R. 2012. "Research Note: Effects of Individual Self-Protection, Industry Self-Regulation, and Government Regulation on Privacy Concerns: A Study of Location-Based Services," *Information Systems Research* (23:4), pp. 1342-1363.

About the Authors

PLEASE PROVIDE ONE-PARAGRAPH BIOS.