Beyond the Privacy Paradox: Objective versus Relative Risk in Privacy Decision Making

Idris Adjerid[*], Eyal Peer[¥], Alessandro Acquisti[€]

[*]University of Notre Dame

[¥]Bar-Ilan University

[€]Carnegie Mellon University

Abstract:

Privacy decision making has been investigated in the Information Systems literature using two contrasting frameworks. A first framework has largely focused on deliberative, rational processes by which individuals weigh the expected benefits of privacy allowances and disclosure against their resulting costs. Under this framework, consumer privacy decision making is broadly constructed as driven by stable, and therefore predictable, individual preferences for privacy. More recently, a second framework has leveraged theories and results from behavioral decision research to construe privacy decision making as a process in which cognitive heuristics and biases often occur, and individuals are significantly influenced by non-normative factors in choosing what to reveal or to protect about themselves. In three experiments, we combine and contrast these two perspectives by evaluating the impact of changes in objective risk of disclosure (normative factors), and the impact of changes in relative, and in particular reference-dependent, perceptions of risk (non-normative factors) on individual privacy decision making. We find that both relative and objective risks can impact individual privacy decisions. However, and surprisingly, we find that in experiments more closely modeled on real world contexts, and in experiments that capture actual privacy decisions as opposed to hypothetical choices, relative risk is a more pronounced driver of privacy decisions compared to objective risk. Our results suggest that while normative factors can influence consumers' self-predicted, hypothetical behavior, non-normative factors may sometimes be more important and consistent drivers of actual privacy choices.

1. Introduction

A significant body of work in Information Systems (IS) and related disciplines (such as economics, marketing, and human computer interaction) has explored consumer privacy decision making. Historically, much of this research has focused on how consumer privacy choices are impacted by normative factors (such as actual costs and benefits, and expected trade-offs), consistent with traditional models of economic agents. Under such models, consumers are assumed to have stable individual preferences for privacy and disclosure, and are able to consistently act on these preferences. For instance, the notion of a "privacy calculus" suggests that consumer privacy decision making is driven by considerations of the risks arising from privacy intrusions from personal data allowances and disclosure, as well as the benefits that these allowances provide (Milne and Gordon, 1993; Culnan and Armstrong, 1999). In addition, a considerable body of work has focused on identifying systematic differences in privacy concerns between consumers (e.g. Smith, Milberg, and Burke, 1996) and has suggested that elevated privacy concerns may correspond to privacy seeking behavior (Malhotra, Kim, and Agarwal, 2004). By and large, this view of consumer privacy decision making could be used to conclude, for instance, that consumers who are notified of data practices by firms will or will not alter disclosure behavior (or utilize tools to mask their online activities) as a function of their preferences for privacy, or lack thereof.

This account of consumer privacy decision making, however, faces the challenge of fully explaining surprising, yet seemingly robust, empirical phenomena observed in privacy contexts, such as the dichotomy between privacy attitudes and actual behaviors (Spiekermann, Grossklags, and Berendt, 2001; Jensen, Potts, and Jensen, 2005). As a result, a number of researchers have conjectured that privacy decision making, while influenced by the consideration of potential benefits and costs of disclosure, could also be impacted by non-normative factors—that is, factors independent of both the objective benefits and risks associated with information allowances and consumers' underlying privacy preferences. Such factors may give rise to systematic, and therefore predictable and replicable, deviations from rational accounts of consumer decision making (Acquisti and Grossklags, 2005). In contrast to traditional accounts of consumer privacy decision making, this alternative "behavioral" account casts doubts on the likelihood that consumers will consistently react to changes in the objective risks and benefits of data allowances or disclosures, since many non-normative factors – such as the framing of the options available to consumers, the default settings of an online interface, and so forth – may equally affect their behavior.

While these alternative accounts of consumer privacy decisions stem from legitimate theoretical frameworks and have stimulated considerable bodies of empirical research, privacy research to date has focused on the impact on consumer privacy decision making of either normative *or* non-normative factors, but not both. Considering that these varying and sometimes contradictory accounts of privacy decision making are likely to account for only some of the variation in observed consumer privacy choices, and given the increasing important economic consequences of individual privacy decision making for both firms and consumers, the absence of a bridge between the two streams of work represents a considerable gap in the privacy literature.

In three experiments, we attempt to address this gap by empirically contrasting the extent to which normative and non-normative factors influence privacy choice across varying privacy contexts and experimental settings. As numerous factors of each type exist and may influence consumer choice, we narrow experimental manipulations by focusing on factors that mirror prevalent features of online privacy decision making (lending our results direct applicability to our context of study) and have an extant literature (sometimes outside privacy research) on which to ground experimental hypotheses. Specifically, we manipulate normative factors by altering the degree of privacy protection afforded to participants' personal data (similar to changes in firm data practices communicated via privacy notices); we manipulate non-normative factors based on the seminal framework of Prospect Theory (Kahneman and Tversky, 1979) and the effect of people's reference dependence (i.e. comparative nature of judgment).

Mirroring the diversity of empirical approaches in the literature, we conduct our experiments using both hypothetical and actual privacy choices. Reflecting the variety of consumer privacy decision contexts, we evaluate the impact of objective and relative risk on both information disclosure behavior and choice of privacy protective options via control mechanisms (i.e. privacy settings). Furthermore, to improve the external validity of our results and evaluate the robustness of our phenomena of study, we use different sampling pools between experiments, and draw on insights from the information systems, behavioral economics, and privacy usability literatures to identify a diverse and unique set of experimental manipulations for both objective and relative risk.

We find that in a hypothetical context (Experiment 1), differences in objective privacy protection result in significant differences in participants' predicted privacy concerns and their willingness to disclose personal information. In contrast, changes in relative risk result in smaller (although significant) differences in reported privacy concerns, and no differences in predicted willingness to disclose personal information. In experiments more closely modeled on real world contexts and where participants made actual privacy choices (Experiments 2 and 3), our results

are reversed. In Experiment 2, we find that participants are twice as likely to choose the least protective option when it is presented as more protective in relative terms, even though it remains the same in absolute terms. In contrast, participants exhibit no differences in self-disclosure despite choosing (on average) objectively different privacy protections between conditions. The context in Experiment 2, however, differs from that in Experiment 1, making direct comparisons difficult. Furthermore, privacy protections in Experiment 2 were selected by participants and not exogenously manipulated. Experiment 3 addresses these concerns, mirroring what participants were asked to imagine in Experiment 1, but focusing on actual disclosure choices and exogenously manipulated levels of privacy protection. We find results consistent with Experiment 2: differences in objective privacy protection had a small or no effect on participants' self-disclosure. In contrast, we find that, holding the objective privacy protection constant, privacy protections that are subjectively perceived as relatively high increased participants' propensity for self-disclosure. We also find that, holding the objective privacy protection constant, privacy protections that are subjectively perceived as relatively low by participants decreased participants' propensity for self-disclosure. Taken jointly, our results suggest that objective privacy protections may have a more pronounced effect on privacy decision making in hypothetical settings, while relative perceptions of privacy protection may have a more pronounced effect in contexts that involve actual privacy choices from participants.

This finding contributes to the IS literature on the drivers and predictors of consumer privacy decision making and to the policy debate and literature on the appropriate mechanisms for addressing consumers' privacy concerns. We simultaneously find evidence in support of classical normative models of privacy decision making and evidence providing additional support for the emerging notion that non-normative factors can predictably impact privacy decision making, thus documenting a robust deviation from economically rational models of consumer behavior. However, we further inform the IS literature on consumer privacy decision making by highlighting the potential for a differential role of normative vs. non-normative factors in driving consumer privacy decision making. Specifically, we find that the previously discussed privacy paradox extends to the manner in which consumers react to varying degrees of privacy protection, resulting in consumers being less sensitive to changes in objective risk than they anticipate in actual choice contexts. More surprisingly, our results suggest that the non-normative factors we evaluate may be underappreciated by consumers when predicting their privacy concerns and behavior in hypothetical situations; but may actually be more influential on, and more consistent drivers of, privacy decision making in choice contexts that more closely mirror real-world privacy choice contexts and involve actual privacy behavior. Such a finding may also explain early results

from the privacy paradox literature (e.g., Spiekermann et al., 2001), suggesting that differences in privacy protection may have a limited effect on actual privacy choices. These findings are also consistent with the broader psychology and behavioral economics literature (e.g., Gilbert & Ebert, 2002; Lowenstein & Adler, 1995), in that they imply that people may overestimate the impact of normative factors on their predicted behavior while underestimating the sometimes powerful impact of decision biases on actual decision making.

Finally, the findings contribute to the literature on reference dependence and relative judgment. Prior literature has proposed models of reference-dependent utility that account for both the utility from absolute levels of consumption and deviations from a reference point (Koszegi and Rabin, 2007; Kahnmen and Tversky, 1979). In this paper, 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 over time, compared to absolute or objective level of protection provided. This is in line with recent work theorizing that changes from reference points might affect attention to information provided, which can subsequently alter choice (Bhatia, 2013).

2. Conceptual Background

An accumulating body of empirical and theoretical economic research has provided behavioral accounts of consumer decision making (e.g., Camerer, Lowenstein, & Rabin, 2011; Ho, Lim, & Camerer, 2006), leading to an increased need to understand the conditions under which consumers may act "rationally" (i.e., react to normative factors) and the conditions under which they may not. This sentiment has been echoed in the IS literature. Goes (2013), for instance, highlights the need to incorporate insights from behavioral economics into theoretical and empirical IS research. This need has become particularly evident for research exploring the drivers of consumer privacy behavior, in light of a growing stream of empirical studies that uncovered evidence of biases and heuristics in consumer privacy decision making (e.g., Acquisti & Grossklags, 2005). The implications of this stream of work would extend to outside academia: for instance, proposed notice and choice mechanisms, which are widely advocated by industry and policy makers (FTC 2012), 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).

Hoofnagle and Urban (2014) argue that the broad support for notice and choice policy approaches is partially rooted in the work of early and highly influential privacy scholars that focused on identifying stable privacy preferences for consumers, and considered consumer

behavior to be largely consistent with rational choice theory. Specifically, they note that Westin (2000) posits 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. A considerable body of economic, marketing, and IS research has been predicated around the notion that privacy decision making is, at least partially, a rational process driven by normative factors, such as the benefits and costs of information disclosure and stable individual differences in privacy preferences. In fact, a considerable portion of privacy research identified in a recent interdisciplinary review of privacy research (Smith, Dinev, and Xu, 2011) is, in our interpretation, normatively focused and generally mirrors these theorized features of privacy decision making.

First, a *privacy calculus* view of consumer decision posits that privacy is subject to interpretation in "economic terms" (Klopfer and Rubenstein, 1977; Dinev and Hart, 2006) and that consumer privacy choices are driven by a systematic weighing of the benefits of information disclosures against the perceived privacy risks from such disclosures (Milne and Gordon, 1993). The potential benefits of information disclosures can be numerous and vary by context: information disclosure can lead to an improved experience in retail via customization of products, promotions, and even user interfaces (Ansari and Mela, 2003), enable users to derive personal and economic value from social networks (Ellison, Steinfield, and Lampe, 2007), and underlies business models for online services providing free content and applications (Leontiadis et al., 2012). A privacy calculus paradigm suggests that consumers weigh these benefits against the potential risks of loss due to these information disclosures. These losses could include those stemming from the misuse of disclosed data (Featherman and Pavlou, 2003), sharing of personal information with third parties, or price discrimination as a result of information disclosures (Viswanathan et al., 2007). A parallel body of work has focused on individuals' "privacy concerns"— an individual's beliefs and attitudes towards information disclosure. Specifically, this body of work has focused on identifying stable differences in in privacy attitudes and concerns between consumers (e.g., Milberg et al., 1995) and evaluates the impact of such concerns on intentions to either disclose information or engage in commercial personal transactions that introduce privacy risks (Smith et al., 1996; Malhotra et al., 2004).

Within this literature, a number of scholars using these differing perspectives on consumer privacy decision making have evaluated the potential for changes in privacy protection to influence consumer privacy choices. Using a privacy calculus lens, scholars argue that increased privacy protections should diminish consumer perceptions of risk from transactions involving personal information and increase consumer intentions to disclose personal information (Milne and Gordon, 1993). A similar logic stems from the literature focused on consumer privacy

concerns, with scholars suggesting that factors that reduce consumer privacy concerns will translate into an increased willingness to disclose personal information (Smith et. al 1996; Malhotra et al., 2004). The literature finds considerable evidence in support of this view. 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; Miyazaki and Krishnamurthy (2002) and Hui, Teo, and Lee (2006) find a significant effect of privacy seals on consumer perception of firm privacy practices and their willingness to disclose personal information; 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; and Xu et al. (2012) find that industry self-regulation and government regulation reduce consumer privacy concerns. Based on this literature, one may evaluate the impact of changes in privacy protections (e.g. the breadth of access to personal information and the anonymity of responses) afforded to participants via a "privacy notice" (similar to firm privacy policies), and also changes in privacy protection as a result of privacy levels chosen by participants in our experiments (similar to consumer privacy settings) and hypothesize that:

H1: Changes in objective levels of privacy protection will impact individual disclosure: lower privacy protection will lead to lower levels of disclosure of personal information.

In addition to the body of work focusing on normative factors impacting consumer privacy concerns and subsequent behavior, a rising theme in the literature is that factors which ostensibly have little or no impact on objective risk and benefits of disclosure can considerably impact people's privacy concerns and personal preferences for self-disclosure (e.g., Moon, 2000). For example, 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, Acquisti, and Loewenstein, 2011). Or, holding objective risk constant, perceived control over who can access and use online personal information can result in an increased likelihood to make sensitive disclosures (Brandimarte, Acquisti, & Lowenstein, 2013). Also, 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); and are also more inclined to divulge information online than when they communicate face-to-face (e.g., Harper & Harper, 2006). In our experiments, we consider the potential of the significant heterogeneity of online privacy practices both between firms and across time (Stutzman, Gross, and Acquisti, 2013) to introduce the comparative nature

of judgment into consumer privacy decision making. Examples of relevant heterogeneity in firm data practices abound in privacy contexts. For instance, firms may notify consumers of their privacy protections in a manner that highlights the relative privacy gains from their services compared to those of their competitors.[1] Moreover, firms that aggregate consumer privacy information often highlight improvements (i.e. relative changes) to consumer privacy over time (e.g., Cox, 2012). We consider the potential for this heterogeneity to introduce changes in individual privacy decision-making that are removed from the objective privacy protections provided.

More formally, we consider the impact on consumer privacy choices of reference dependent judgment as introduced by Kahneman and Tversky (1979) within the framework of Prospect Theory. Specifically, Kahneman and Tversky posit that individuals 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. Prospect Theory is a particularly useful framework for studying the impact of normative vs. non-normative factors for two reasons. First, it allows for both the impact of objective features of a particular choice context that should influence choice (e.g. price of a product) and the non-normative 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 choice to influence behavior. Second, considerable empirical evidence exists in support of reference dependent decision making, and to also rule out alternate, 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, Knetsch & Thaler, 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 point: whereas buyers consider the purchase of a new item as a gain, sellers consider it as a loss (e.g., Novemsky & Kahneman, 2005). The similar WTA-WTP gap has also been shown in relation to disclosure decisions (Acquisti, John, & Lowenstein, 2013). Additionally, Knetsch, Tang, and Thaler (2001) find that the endowment effect is robust to repeat trials, suggesting that is it not a side effect of consumer learning or experience. In fact, more recent

[1] Microsoft's "Scroogled" Ad campaign sought to highlight the privacy protectiveness of their services (e.g. search, email, etc.) relative to those of Google.

literature (e.g., Koszegi and Rabin, 2007) has simply incorporated 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. In our experiments, we evaluate whether insights from Prospect Theory and the empirical literature on reference dependence hold in the context of privacy decision making. For instance, we consider the role of reference dependence in terms of how individuals react to privacy notices communicating privacy protection. Under normative accounts of privacy decision making, 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 also consider the role of reference dependence in consumer choice of settings that restrict the use and access to their personal information. Namely, we argue that choice sets can vary significantly between services and across time and that this variation may lead to differences in the perceived relative protectiveness of objectively identical privacy levels. As a result, we hypothesize that:

H2: Individuals' relative perception of privacy protection will influence individual privacy decision making: the more protective privacy levels will be perceived to be, the more desirable they will be; leading to higher levels of disclosure of personal information.

2.1 Differential Effects of Objective vs. Relative Privacy Protection

Given that compelling accounts exist for both normative and non-normative factors driving privacy decision making, we supplement our formal hypotheses by exploring whether factors exist that moderate the effect of objective and relative privacy protection on privacy choice. Specifically, we focus on an important distinction noted by IS privacy scholars (e.g., Bélanger and Crossler, 2011; Smith et al., 2011) between classical, normatively-focused work on consumer privacy decision making, and a more recent empirical and behavioral economics privacy literature: whether privacy choice is observed in hypothetical settings, where stated privacy concern and predicted behavior is measured, or in settings modeled more closely on real world contexts and actual privacy choices are observed.

This distinction is relevant to the impact of normative vs. non-normative factors in light of the literature documenting a privacy paradox where consumers' stated privacy concerns (and

associated intended or predicted behavior) can be disjoint from actual observed behavior (Spiekermann et al., 2001; Jensen et al., 2005). Given that this observed phenomenon belies some limitation in the ability of consumers to act on stated privacy concerns, we conjecture that this inconsistency in choice may extend to the manner in which consumers react to objective differences in privacy protection. Specifically, we evaluate whether consumers may be sensitive to changes in privacy protection in hypothetical contexts but have a diminished sensitivity to these changes in actual choice contexts more closely modeled on real world contexts. Interestingly, the literature evaluating actual consumer behavior is seemingly mixed in this respect. For instance, while Jensen et al. (2005) document an inconsistency between privacy concerns and choice, they find that the existence of trust seals and privacy notices do impact consumer privacy behavior; and, in later work, Tsai et al. (2011) find that privacy seals can increase consumers' willingness to engage in commercial transactions and that some consumers are willing to pay a premium to shop at privacy preserving retailers. However, this work has evaluated whether the *existence* of privacy protection has an impact on behavior, but not necessarily whether consumers will react to variation in levels of privacy protection. In fact, Jensen et al. (2005) note that while consumers reacted to privacy notices, hardly any of the participants actually read them, suggesting that the specific protections therein were not the source of differences in observed behavior. In addition, Spiekermann et al. (2001) finds that disclosure behavior was not impacted by "harsh" vs. "protective" privacy notices. However, Spiekermann et al. (2001) did not measure how participants' stated privacy concerns or their predicted behavior were impacted by the two notices used in the study, making it unclear whether consumers would anticipate this lack of sensitivity to privacy notices.

While the privacy paradox literature suggests that changes in normative factors may have an inconsistent impact on individual decision making, the behavioral economics and psychology literatures offer insights into when the impact of non-normative factors is likely to be pronounced. Specifically, these literatures suggest that individuals tend to overestimate their propensity to act rationally and to underestimate the influence of decision biases on their own behavior. For instance, Liberman, Samuels, and Ross (2004) found that participants grossly underestimate subtle changes to the labels of choice contexts on their subsequent behavior; Loewenstein and Adler (1995) find that participants consistently underestimated the impact of being given an item on their subsequent valuation of that item (i.e. the endowment effect); and O'Donoghue and Rabin (2000) find that individuals can be naïve in their estimation of their rational future behavior, such as their susceptibility to an immediate gratification bias (i.e. time inconsistent discounting).

Taken together, these arguments suggest a novel account of the influence of normative vs. non-normative factors on consumer privacy decision making. First, our formal hypotheses seek to contrast works from the various bodies of work in the literature and allow both normative and non-normative factors to have an influence on privacy decision making. However, and distinct from what has been proposed and shown in the literature, we suggest that while normative factors are likely to strongly influence consumers' predicted behavior, non-normative factors may be more important and consistent drivers of behavior in actual choice contexts.

3. Methods

In three experiments, we evaluated the role of objective changes in privacy protection and relative judgments of privacy protections on consumers' self-disclosure behavior and consumers' choice between privacy levels. In all experiments, we manipulated a factor expected to affect the relative perception of the protectiveness of the privacy level. In Experiment 1, we employed the context of a hypothetical study on (un)ethical behavior where graphical privacy notices were used to manipulate both the objective protection and the *relative* perception of privacy protection, and participants reported their privacy concerns and predicted their disclosure behavior. In contrast, Experiments 2 and 3 involved actual privacy decision making by experimental participants while using contexts and privacy protections more closely modeled on real world privacy choice. Experiment 2 used a social networking context in which the relative position of privacy protection options in privacy settings was used to manipulate the relative perception of privacy protection. Subsequently, participants were asked to complete a social media profile that involved actual, sometimes sensitive, information disclosures. Finally, in Experiment 3 we again used the context of a study on ethical behavior; but unlike Experiment 1, provided text-based notices similar (although simplified) to those used by online services, and asked participants to make actual disclosures of sensitive personal information.

3.1. 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 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.). Our evaluation of participant disclosure behavior is comparatively more involved.

Because participants in all experiments were presented a series of questions asking them to predict their propensity to make or to actually make sensitive disclosures, we observe multiple, correlated responses from a single participant. As a result, we use a random effects linear regression model to evaluate differences in average disclosure between conditions.[2] This model accounts for the correlation between responses from a single 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 estimated 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 specification, we also include $\boldsymbol{X}_j$*:* a vector of controls for different features of the questions asked to participants. For instance, $Intrusive_j$ controls for questions that differ in their intrusiveness. $\boldsymbol{Y}_i$ is a vector with controls for participant specific controls (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 as they should be uncorrelated with observed ($\boldsymbol{X}_j$, $\boldsymbol{Y}_i$) and unobserved (θ_i) individual differences and the error term u_{ij}. While our controls are not necessary for the unbiased estimation of the effect of our treatments on disclosure behavior, they are included in some specifications to rule out any breaks in randomization, and account for some of the variation in disclosure behavior between participants.

4. Experiment 1

In Experiment 1, we manipulated, between subjects, both differences in objective privacy levels and whether participants perceived a relative increase or decrease in the privacy level *over time* (while holding objective privacy levels constant). We used self-reported perceptions of the privacy levels (e.g., satisfaction with the protections provided and overall concerns for privacy) and hypothetical willingness to disclose as key dependent variables and indicators for privacy judgments.

[2] we opted for a linear probability model estimation in lieu of a non-linear 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.

4.1. Participants

Two hundred and twenty one participants from Amazon Mechanical Turk[3] (M_{female}=37.56%; M_{age} = 29.16, SD_{age} = 9.76) completed the study and were paid $0.30 for completing the study.

4.2. Design and procedure

Participants were asked to provide their personal opinions regarding two surveys that our research group is planning to conduct. Participants were told that our research group conducts surveys which 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 the first survey (Survey A), including the confidentiality protection of the survey. This was described using a figure that showed, on five parameters, whether the survey is offering a particular protection or not. For example, a survey may (or may not) require participants' email addresses. In one condition, the first survey (Survey A) provided a low overall privacy level with the "Less Protective" option for four of the five parameters described (see Figure 3a), and in the other condition, the survey provided a high privacy level with the "More Protective" option for four of the five parameters described (see Figure 3b). 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. "Does the survey require a valid email address?"). They were then asked to report their satisfaction with the protections provided in each survey, perception of potential harm from disclosure in the study, and concerns about their privacy (see Appendix A.1). Finally, participants were asked questions gauging their hypothetical willingness to disclose for descriptive but sensitive information (e.g. address or phone number), and how often had they engaged in a set of (un)ethical behaviors (see Appendix A.2).

Next, all participants proceeded to review the second survey (Survey B) which provided a medium privacy level in both conditions (see Figure 3c). Participants were asked to evaluate Survey B using the same questions as used for Survey A. In the resulting design, participants in the first condition perceived an increase in the privacy level from Survey A to Survey B (low to medium) whereas participants in the second condition perceived a decrease in the privacy level from Survey A to Survey B (from high to medium). Notably, the actual level of privacy for

[3] Prior research has validated AMT samples as equally representative relative to other internet samples and much more representative than student samples (Buhrmester, Kwang, and Gosling, 2011) and central findings in IS and the decision sciences have been replicated using AMT samples (Goodman, Cryder, and Cheema, 2013; Steelman, Hammer, and Limayem, 2014).

Survey B remained the same for both conditions, although the subjective level of that survey's privacy might have changed.

[Figure 1a: Low Protection]

[Figure 1b: High Protection]

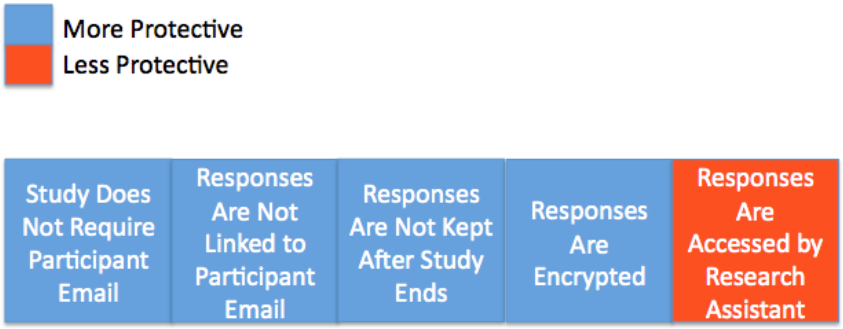

[Figure 1c: Medium Protection]

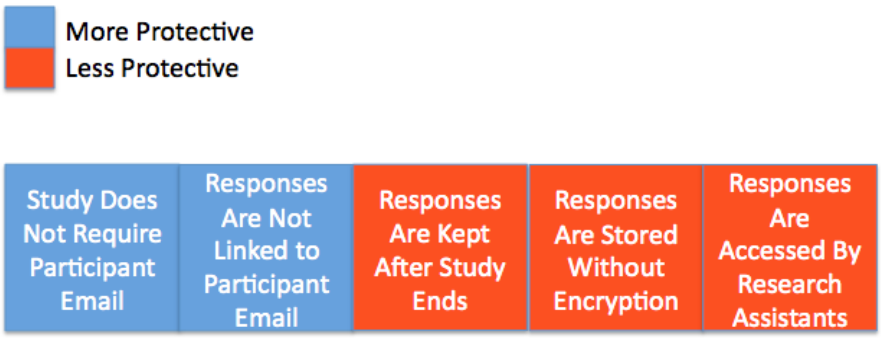

4.3. Results

First, 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 3a and 3b) was effective at influencing the perception of privacy protection in the first survey (Survey A): Participants provided high protections reported being significantly more satisfied with the protections provided (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 them as a result of disclosing personal information (M_{High}=2.86, M_{Low}=4.02), t (219) = -7.46, p <.001, d=1. We also evaluated the impact on participants' predicted disclosure behavior of differences in objective privacy protection using a random effects linear regression estimation approach. In this experiment, we asked participants to report on a five item scale (1 indicating being "Very Unlikely" to disclose and 5 being "Very Likely" to disclose) their likelihood of disclosure for a given question. We found that the objective differences in privacy levels in the Survey A continued to have a significant effect on participants' predicted behavior. Participants provided the low privacy level predicted being significantly less likely (β_{Low}= -.67, p=<.001) to disclose personal information (Table 2, Column 1*)*. Moreover, we find consistent results (β_{Low}= -.65, p<.001) when including controls for question type (descriptive vs. ethical) and participants' age and gender (Table 2, Column 2). Broadly, these results provide strong support for the hypothesis that objective risk will impact consumer privacy choice (H1 supported).

[Table 1: 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	$p = .04$

For the second survey (Survey B), which had an objectively identical medium privacy level (Figure 3c) 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 may be used in way that may harm them (M_{Inc} = 3.37, M_{Dec}= 3.68), t (219) = -2.04, $p = .04$, d = 0.28. However, these effect sizes were considerably smaller than those for objective risk. Moreover, 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 2, Column 3), that this was consistent ($\beta_{Increasing}$= .11, p=.363) when including controls for question type and participant age and gender (Table 2, Column 4). These results suggest mixed support for the

hypothesis that relative perception of privacy protection will impact privacy choice (H2 mixed support).

[Table 2: Experiment One Regression Results]

VARIABLES	(1) Disclosure	(2) Disclosure	(3) Disclosure	(4) Disclosure
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

4.4. Discussion

The results of Experiment 2 suggest that differences in both objective and relative risk have some effect on individual perceptions of protection in the study, including satisfaction with privacy protections and privacy concerns. However, these differences were considerably more pronounced for differences in objective risk, and we find no differences in predicted levels of disclosure for changes in the relative perception of provided privacy protection. Initially, this study suggests that, while relative changes in privacy levels can impact perceptions of privacy levels, they may not necessarily impact actual consumer choices.

However, this study focused on hypothetical elicitation of privacy choices, which, as discussed previously, may not necessarily reflect actual participant behavior. We address this concern in Experiments 2 and 3 where we measure actual participant privacy decision making. Moreover, the presentation of privacy levels in this experiment may have been somewhat heavy-handed when compared to real world contexts. First, we used a graphical representation of privacy levels, including a key that alerted participants to riskier uses of their personal information. However, privacy levels online are most often communicated in text based privacy notices where changes in protection may not be as salient. Moreover, the focus was largely on the

privacy levels provided as participants were explicitly asked to carefully scrutinize and interpret the described privacy levels—prior research finds that consumers usually don't pay much attention to online privacy notices (Vila, Greenstadt, and Molnar, 2004). We address these issues in Experiment 3.

Finally, the design for this experiment did not allow us to identify the distinct effect of relative increases and decreases in the privacy level. For example, it may be the case that our results were purely driven by decreases in the privacy levels, and that increases in the privacy level did not have an impact. This issue is also addressed in Experiment 3. These concerns notwithstanding, this study still offers some evidence in support of the notion that both objective and relative risk can impact people's perception of privacy protections, and potentially their privacy and disclosure choices.

5. Experiment 2

In Experiment 1 we focused on how objective and relative privacy protection could impact consumers' satisfaction from privacy protections and subsequent hypothetical disclosure. In Experiment 2, we examined the role of reference dependence on choice of privacy levels from participants, as well as on actual subsequent disclosure. To increase the ecological validity of our findings, we used a real-life scenario, and invited students to join a new online social network that is being formed in their university. This scenario enabled us to examine consumers' decisions in what appeared to them to be a high-stake situation involving the actual disclosure of their personal information to others.

5.1. Participants

Participants were 177 (50.85% males, M_{age} = 22.73, SD=12.95[4]) newly admitted students in a large North-Eastern University in U.S. Participants were given a small candy bar as a token of appreciation for contributing to the research.

5.2. Design and procedure.

Students were approached during the two orientation weeks before their first year in the university. The students were asked to take part in a "research project about a new online social network" in their university. First, participants chose who would be able to see their profile in the new online social network. Participants chose between six options. The top choice was "only the

[4] Age statistics are reported based on responses from 61% of participants who chose to disclose their age or their birth year.

students I invite" (the most restrictive privacy setting). The bottom, sixth, choice was "All students, faculty and staff in the university" (the most inclusive privacy settings). Among the other four options, one of the options was always "All current students", selected as the default option. The position of this option, as well as the other options that followed or preceded it, varied between the conditions. In the "high default" condition, this option was placed as second (right after the most restrictive option), followed by four less restrictive options, and the last option of "All students, faculty and staff in the university". Conversely, in the "low default" condition, the default option of "All current students" was placed as fifth, right before the least restrictive option, preceded by four more restrictive options between the default option and the top, most restrictive, option (see Figure 2). Participants were asked to choose one of the options as their personal privacy setting in the new online social network.

[Figure 2: Privacy setting choices in the high and low default conditions]

High default

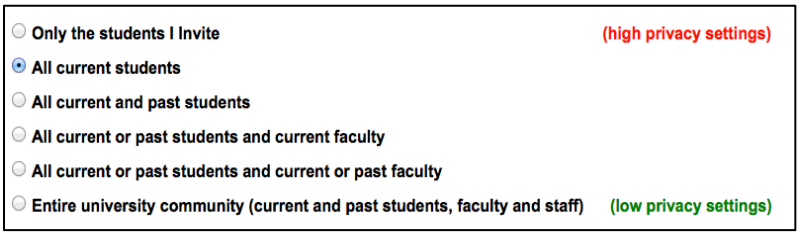

Low default

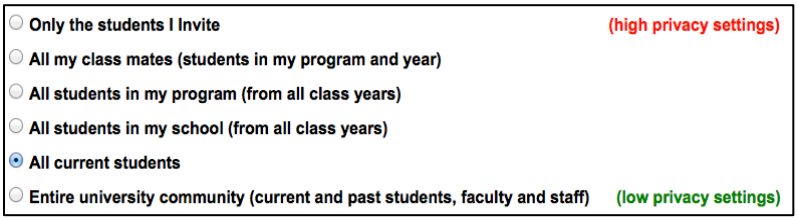

After choosing their desired privacy setting, participants were asked to rate, on a scale of 1 (*very low)* to 5 (*very high),* how high or low the privacy level that they had chosen was. Participants then proceeded to what appeared as "profile building pages", in which they answered several personal questions about themselves that would be used to create their personal profiles in the new online social network. These questions included various personal items (such as name, gender, address, school and department, past GPA, relationship status, and sexual orientation). Lastly, participants answered several questions that gauged their suspicion of the study's cover story (the creation of a new online social network). None of the participants expressed any disbelief regarding the stated objective of the study.

5.3. Results

In response to the manipulation check question ("how low or high do you consider the privacy setting you chose to be?"), we found that participants in the high default condition rated their chosen settings as subjectively higher than those in the low default condition, and that these differences were statistically significant, F (1,171) = 5.06, p = .03.

[Table 3: Experiment 2 Summary Results]

Privacy setting	High default	Low default
Only the students I Invite	34.83%	30.68%
All my class mates (students in my program and year)		10.23%
All students in my program (from all class years)		6.82%
All students in my school (from all class years)		3.41%
All current students	34.83%	34.09%
Current and past students	12.36%	
Current and past students and current faculty	6.74%	
Current and past students and faculty	3.37%	
Entire university community	7.87%	14.77%
Summary		
Above the default	34.83%	51.14%
Default	34.83%	34.09%
Below the default	30.34%	14.77%

Next, we observed the distribution of participants' privacy setting choices (see Table 3). Considering the choices common to both conditions (the top, the bottom and the "default" choice), we found a) that those in the low default condition were twice as likely (14.77% vs. 7.87) to choose the least restrictive option; b) similar percentages of participants chose the default setting (34.83% vs. 34.09%); and c) participants in the low default condition were somewhat less likely (30.68% vs. 34.83%) to choose the most protective option. Moreover, we found that the percent of participants who chose a setting that was higher (more restrictive) than the default was higher in the low vs. high default condition (51.14% vs. 34.83%), while the percent of participants who chose a setting that was lower (less restrictive) than the default was higher in the high vs. the low default condition (30.34% vs. 14.77%). These differences were statistically significant, χ^2 (2) = 7.49, p = .02. Together, these results provide support for the hypothesis that relative perception of privacy protection can influence choice of objectively identical options (H2 supported).

We also examined how the manipulation of framing privacy options affected participants' subsequent self-disclosure as they answered the questions that were (ostensibly) used to create their profile on the new online social network. Although we did not exogenously manipulate participants' objective risk (i.e. they chose their own privacy levels), we did observe significant differences in privacy levels between our exogenous treatment conditions, with participants in the high default condition choosing (on average) less restrictive settings. A normative view of privacy decision making may thus posit that participants in the high default condition should disclose less than their counterparts in the low default condition. Evaluating the impact of our randomized treatment on the full set of profile questions asked to participants (Table 2, Column 1) showed no significant effect of on disclosure for those in the high default condition ($\beta_{\text{HighDefault}}$ = -.016, p=.56). Focusing on the three specific questions that were found, in our pre-tests, to be regarded as especially sensitive by students from the same university (past GPA scores, current relationship status and sexual orientation), we still failed to find a significant difference in disclosure for those in the high default condition (Table 4, Column 2; $\beta_{\text{HighDefault}}$ =-.039, p=.42). Finally, when we restrict our sample to sensitive questions and the participants that chose privacy levels that were not common to both conditions (Table 4, Column 3), we still find insignificant differences in disclosure for participants in the high default condition ($\beta_{\text{HighDefault}}$ =-.056, p=.29). In other words, although participants in the two conditions made different choices of privacy settings, they did not "correct" for their chosen privacy setting by disclosing more or less. In fact, it appears that *despite* having selected different privacy settings, participants subsequent disclosure remained the same (H1 Not Supported).

[Table 4: Experiment 2 Regression Results]

VARIABLES	(1) Disclosure	(2) Disclosure	(3) Disclosure
High Default	-0.0161	-0.0388	-0.0561
	(0.0275)	(0.0483)	(0.0524)
Constant	0.783**	0.788**	0.795**
	(0.0192)	(0.0341)	(0.0360)
Observations	3,009	531	417
Number of id	177	177	139

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

5.4. Discussion

The results of this experiment suggest that the different positioning of the default choice ("All current students") affected participants' reference point which altered their perceptions of the

restrictiveness of the common extreme options ("only students I invite" vs. "all students, faculty and staff") between conditions, leading to differences in the chosen privacy setting. Specifically, the high positioning of the default choice seemed to have made participants perceive the choice to share with the entire university community (the least restrictive option) relatively more risky, making it less attractive. Conversely, the low positioning of the default seemed to have made the objectively identical choice seem, in relative terms, less risky, thus increasing participants' choice of that option. In other words, the default setting served as a reference point above or below which options were considered more or less protective. However, and in contrast to the predicted behavior of participants in Experiment 1, we do not find differences in disclosure behavior between experimental conditions that chose ostensibly different levels of privacy protection. These results suggest that when participants are asked to make actual privacy choices in a context modeled on a common privacy choice context, participant behavior is influenced considerably by the relative perception of privacy protection but not by objective differences in the chosen privacy protection.

However, there are some alternative explanations that might hinder these conclusions. First, it is possible that there was an objective difference in the *attractiveness* of the choices given in one condition that were not given in the other condition. For example, it is possible that the options below the default in the high default condition were objectively more attractive than the options above the default in the low default condition. A closer look at Table 3 reveals that this concern is unsupported. About 22.5% of the participants in the high default condition chose the options that did not appear on the low-default condition compared to 20.4% of participants in the low-default condition that chose the options that did not appear on the high-default condition. This lack of difference suggests that it was not the attractiveness of "unseen" options that was responsible for the effect of the reference point (the default choice) on choices of privacy settings.

Another alternative explanation could be that the effect we observer was simply the differences in choice sets between conditions, and the subsequent granularity of the options provided between the default choice and the extreme options that were common to both conditions. That is, participants in the low default condition with *true* preferences for an intermediate setting between the default setting and the least restrictive option (i.e. the settings available in the high default condition) may have been forced to choose either the default or the least restrictive option. This presents a potential confound as this phenomena may yield results similar to those found in our experiment. Alas, this confound is very hard (or impossible) to disentangle without creating other, possibly more problematic, confounds. For example, if the choices were to remain the same in both conditions, and the default option would be the second

one in one condition and the fifth one in the other condition, that would have to entail that the default choice be different between the conditions, meaning that the position of the default would have been confounded with the objective level of privacy that default offers. Thus, this confound seems to be inherent to the method employed in this experiment.

Finally, while we focus on differences in disclosure behavior between exogenously manipulated conditions, this is not an ideal for evaluating the impact of different privacy protection on behavior since participants self-select into different levels of privacy protection—a choice likely correlated with other unobserved factors that could confound our results. Moreover, while this context demonstrates the range of choice contexts where reference dependence can play a role, the change in contexts makes comparisons between Experiment 1 and 2 problematic.

Experiment 3 addresses these limitations by mirroring Experiment 1 while recording actual privacy choices and focusing on a disclosure context in which objective privacy levels presented to participants are exogenously manipulated while still manipulating the relative protectiveness of provided privacy levels.

6. Experiment 3

In addition to addressing the abovementioned limitation of Experiment 2, Experiment 3 also addressed two limitations that arose in the Experiment 1: a) the overly "heavy-handedness" of the presentation of privacy protection levels (by using, in Experiment 3, text-based privacy notices) and b) the inability of Experiment 1 to manipulate both objective and relative privacy risk and evaluate the unique impact of each of them. To overcome these limitations, in Experiment 3 we asked participants to take part in two separate studies on (un)ethical behavior. Similarly to Experiment 1, each survey provided different stated levels of privacy protections to participants. Between participants, we kept the objective level of privacy offered by the surveys at the same level (and used as a simple text-based privacy notice), and manipulated whether participants experienced a relative increase or decrease in privacy levels. We examined the effects of such changes on actual disclosure behavior by participants. 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.[5]

[5] Early analysis of Experiment 3 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).

6.1. Participants

Four hundred and fifteen participants from Amazon Mechanical Turk (51.61% females, M_{age} = 31.27, SD_{age} = 10.72) completed our online study about (un)ethical behavior. The experiment was advertised to participants as two, ostensibly unrelated, surveys on (un)ethical behavior.[6]

6.2. Design and Procedure

The design was a 2 (high vs. low protection in the first survey) X 2 (high vs. low protection in the second survey). Thus, our study consisted of four groups in which privacy either *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, participants were asked demographic questions, including email address as a mandatory question. Participants were told that we would check the validity of their email addresses prior to approving payment for the study (we did not actually store email addresses). Then, participants were provided with a privacy notice about the way their answers to the questionnaire would be stored in either a low or high protection condition. 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 A.3 for full text of notices provided). Specifically, participants offered "low" protections were informed that their answers would be linked to their email addresses. Conversely, those offered "high" protection were informed that their answers would not be linked to their email addresses.[7] Participants were then presented with six questions relating to ethically questionable activities (see Appendix A.5 for full set of questions). The questions included a subset of the questions that were judged in Acquisti et al., (2012) as highly intrusive (e.g. "Have you ever had sexual desires for a minor?"). Participants were then asked to complete an additional survey that followed the same structure as the first survey but had a different visual design to help convince participants they were participating in two separate studies (see Appendix A.4). Also, participants were provided two separate confirmation codes to submit in order to receive payment for completing both surveys.[8] In the second survey, participants were again asked for their emails and demographic information; then,

[6] Participants in experiment 1 were not able to participate in experiment 3.

[7] In Experiment 1, participants commented in exist questions 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.

[8] 99.50% of the participants that completed the exit questions indicated they had participated in more than one study and 96.59% of participants indicated that there were differences between the two studies.

they were given a privacy notice about the way their answers to the questions would be stored. As in the first survey, the privacy notice was either high (not linking responses to emails) or low (responses linked to emails). Then, participants were presented with six questions, different from those presented to them in the first survey about other ethically questionable behaviors (see Appendix A.5). Lastly, participants responded to some exit questions that gauged both their perception of whether privacy protections changed in each study (e.g. whether the increased, decreased, or stayed the same, depending on the condition) and their recall of privacy notices in both surveys.

6.3. Results

We found that our manipulations of high and low protection elicited the desired effect with participants in the high protection conditions reporting significantly higher beliefs that their responses would be linked back to them (M_{High}=.79, M_{Low} =.14), t (411) = 18.81, $p < .001$, d = 1.86, relative to participants in the low 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 only marginally 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 impact participant behavior.

Secondly, we evaluated disclosure behavior in the second survey of our experiment where participants were either presented an increasing, decreasing, or identical protection compared to the first survey. A few (11.33%) of participants were unable to accurately recall privacy notices and thus disagreed that protections had increased, decreased, or stayed the same. These participants were excluded from our second survey analysis, leaving 368 usable responses.[9] First, we compared participants that had High Protection in both surveys to participants that had Low Protection in both surveys. For the analysis in the second round, we control for the possible impact of disclosing more in the first survey on second survey disclosures using *Survey1Sharing* which ranged from a value of zero for participants that did not admit to any of the behaviors in Survey 1, to a value of six for participants admitting to all behaviors in Survey 1.

In contrast to our results for the first survey, we found no effect of high protection vs. low protection on disclosure (β_{High} = -.003, p = .9) in the second survey (Table 5, Column 3) and this

[9] The results remained similar when including these participants.

was consistent (β_{High} = -.0001, p = .99) when including controls for question intrusiveness, the survey's visual design, and participant demographics (Table 5, Column 4). This suggests that participant sensitivity to different levels of privacy protection diminished over a fairly short period of time (i.e. between the time taken to complete the first and second survey), ultimately resulting in mixed evidence that objective risk will influence consumer privacy choice (H1 mixed support). Second, we evaluated the impact of changing protection on disclosure relative to conditions in which did not perceive an increase or decrease (participants were provided objectively equivalent privacy notices). We found an increase in the propensity to disclose ($\beta_{Increasing}$ = .06, p =.04) for participants that perceived an increase in protection relative to those whose protections stayed constant, and that this was consistent when 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 that perceived a decrease in protection relative to those whose protections stayed constant (Table 5, Column 7), and, again, that this was consistent when including controls for question intrusiveness, the survey's visual design, and participant demographics (Table 5, Column 8). These results suggest that participant relative perception of privacy protection had a consistent impact on disclosure behavior (H2 Supported).

6.4. Discussion

Similar to our conclusions from Experiment 2, our results suggest that that participants' actual self-disclosure behavior was considerably affected by the relative perception of privacy protection (i.e., from high to low or low to high) but not by consistently impacted by objective differences in protection. Specifically, we found that objective levels of disclosure had either a weak or non-effect on disclosure with estimated effect in the first survey not significant when we included our controls, and no effect of objective differences in privacy protection on disclosure behavior in second survey. However, relative changes in privacy protection had a significant and larger impact on disclosure behavior in the second survey. These findings suggest that people's propensity to disclose personal information can be influenced by seemingly irrelevant factors such as the relative, instead of the absolute, value of privacy protection while the role of objective factors driving behavior may be more limited. In particular, we note that the results in this experiment diverge considerably from the results in Experiment 1 which suggested a considerable effect of objective risk on behavior in a hypothetical study using the same context, experimental design, and sampling population.

[Table 5: Experiment 3 Regression Results]

VARIABLES	(1) Disclosure	(2) Disclosure	(3) Disclosure	(4) Disclosure	(5) Disclosure	(6) Disclosure	(7) Disclosure	(8) Disclosure
High Protection	0.0499*	0.0423+	-0.00336	0.000100				
	(0.0240)	(0.0231)	(0.0278)	(0.0278)				
Increasing					0.0605*	0.0604*		
					(0.0295)	(0.0292)		
Decreasing							-0.0745**	-0.0714**
							(0.0269)	(0.0271)
Intrusive		0.0758**		-0.111**		-0.113**		-0.0858**
		(0.0178)		(0.0271)		(0.0259)		(0.0288)
Age		-0.00538**		0.00139		0.00324*		0.000579
		(0.000964)		(0.00156)		(0.00136)		(0.00160)
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**	0.0950**	0.0973**	0.110**	0.109**
			(0.00932)	(0.00988)	(0.0102)	(0.0105)	(0.00984)	(0.0103)
Constant	0.444**	0.525**	0.0149	0.0206	0.0408	-0.0273	0.000929	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

7. Discussion: Choice Context and the Impact of Normative vs. Non-normative Factors

Across three experiments we find compelling evidence for the central thesis of this manuscript that both normative and non-normative can simultaneously influence consumer perceptions of privacy risk and actual privacy choices. This initial result bolsters the notion that incorporating non-normative factors into formal models of consumer privacy choice would likely improve the generalizability of these models and their predictive power, particularly in actual privacy choice contexts.[10] Moreover, we also find evidence for our conjecture that the impact of normative factors will be pronounced in hypothetical contexts (relative to actual choice contexts) and that the impact of non-normative factors will be pronounced in actual choice contexts (relative to hypothetical contexts). Specifically, we note that H1 was strongly supported in Experiment 1 where a hypothetical context was employed, but weakly supported or not supported at all in Experiments 2 and 3 where actual choice was observed (see Table 6). In contrast, we only find partial support for an impact of relative perception of privacy protection in Experiment 1, where a hypothetical context was employed; but we find consistent evidence in support of the impact of relative perception of privacy protection in Experiments 2 and 3, where actual choice was observed (see Table 6).

[Table 6: The Impact of Relative vs. Objective Privacy Protection]

CONDITIONS	Experiment 1	Experiment 2	Experiment 3
	Hypothetical Choice	Actual Choice	
H1: Objective Privacy Protection	Supported	Not Supported	Mixed Support
H2: Relative Privacy Protection	Mixed Support	Supported	Supported

A number of factors may be driving these effects. First, participants may underestimate the uncertainty in their preferences for disclosure in actual choice contexts and their desire to conform to experimenter requests (i.e. admit to engaging in unethical behavior). Moreover, the decision biases that have previously been shown in the context of privacy decision making may be contributing to the effects we observe. For instance, in Experiment 2, participants' perception of control may have driven participants in both conditions to disclose at equally high levels. This is consistent with prior work showing that consumers' perception of control over their personal information is a key predictor of their privacy concern (Xu et al., 2012), and that increased

[10] As we noted earlier, a considerable body of word finds support for these models on hypothetical or predicted privacy decision making.

control can result in elevated subsequent disclosure (Brandimarte et al., 2013). In Experiment 3, participants seem to be falling prey to a habituation effect over a fairly short period of time. This suggests that in contexts where consumers make repeated and similar privacy decisions, the effect of objective differences in privacy protection may stagnate over a fairly short period of time. As a result, non-normative factors (e.g. the relative change in privacy protection or contextual cues) may result in a comparatively pronounced and more consistent effect on behavior. In addition, in Experiment 1, the focus of participants was on the privacy protections provided and participants were asked to recall these protections back to us. In actual choice contexts, this is rarely the case, and privacy considerations are often secondary to the service consumed by consumers. In these cases, objective privacy protections may simply not be salient to consumers relative to deviations from expectations of privacy protection or changes in privacy protection over time.

8. Conclusions

Our work is founded on the IS literature on consumer privacy decision making and the behavioral economics literature on reference dependence and relative judgment. Namely, 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, at least 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 level of protection provided. Moreover, we find that judgmental and decisional biases impact both individual self-disclosure behavior and individual choices with regards to how information is accessed and used. The latter are pervasive online (e.g., privacy settings on online social networks such as Facebook) and are increasingly relevant as mechanisms for consumers to express their preferences for privacy in contexts where personal information is passively collected (i.e., not through explicit self-disclosure) from consumers (e.g., behavioral advertising). Our results are in-line with recent work that theorized that changes in reference point might affect attention to information provided, which can subsequently alter choice (Bhatia, 2013). More broadly, our work contributes to 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 starting to identify the conditions these effects are most likely to materialize. This work could support or inform extensions of current theoretical frameworks that attempt to model consumer privacy choice.

Moreover, the evidence in support of reference-dependent privacy decision making presented in this manuscript has, in itself, considerable implications for firms and policy makers. If consumers' judgments of privacy protections in actual choice contexts are relative rather than absolute, market choices might not necessarily capture or reflect "objective" privacy preferences. For example, if privacy protection is increased from a very low (absolute) level of protection, consumers might consider that 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 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 quite high. These results suggest that policy maker goals of consumer privacy protection 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. However, these results 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 impacted privacy decision more than the objective risk that participants faced.

The implications for firms seeking value from innovations rooted in the collection of consumer personal information is less clear. 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 less disclosure and trust by consumers and 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 firms suggests 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

1. Acquisti, A., & Grossklags, J. (2005). Privacy and rationality in individual decision making. IEEE Security & Privacy, 2, 24-30.
2. Ansari, A., & Mela, C. F. (2003). E-customization. Journal of Marketing Research, 40(2), 131-145.
3. Bélanger, F., & Crossler, R. E. (2011). Privacy in the digital age: a review of information privacy research in information systems. Mis Quarterly, 35(4), 1017-1042.
4. Bhatia, S. (2013). Associations and the accumulation of preference.Psychological review, 120(3), 522.
5. Brandimarte, L., Acquisti, A., & Loewenstein, G. (2013). Misplaced Confidences Privacy and the Control Paradox. Social Psychological and Personality Science, 4(3), 340-347.
6. Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk a new source of inexpensive, yet high-quality, data?. Perspectives on Psychological Science, 6(1), 3-5.
7. Camerer, C. F., Loewenstein, G., & Rabin, M.. (2011). Advances in behavioral economics. Princeton University Press.
8. Cox, C. (2012). Making It Easier to Share With Who You Want. Facebook. Retrieved January 28 2014 from https://www.facebook.com/notes/10150251867797131
9. Culnan, M. J., & Armstrong, P. K. (1999). Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation.Organization Science, 10(1), 104-115.
10. Dinev, T., & Hart, P. (2006). An extended privacy calculus model for e-commerce transactions. Information Systems Research, 17(1), 61-80.
11. Ellison, N. B., Steinfield, C., & 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), 1143-1168.
12. Featherman, M. S., & Pavlou, P. A. (2003). Predicting e-services adoption: a perceived risk facets perspective. International journal of human-computer studies, 59(4), 451-474.
13. Gilbert, D. T., & Ebert, J. E. (2002). Decisions and revisions: the affective forecasting of changeable outcomes. Journal of personality and social psychology, 82(4), 503.
14. Goes, P. B. (2013). Editor's comments: information systems research and behavioral economics. MIS quarterly, 37(3), 3-8.
15. Harper, V. B., & Harper, E. J. (2006). Understanding student self-disclosure typology through blogging. The Qualitative Report, 11(2), 251-261.
16. Ho, T. H., Lim, N., & Camerer, C. F. (2006). Modeling the psychology of consumer and firm behavior with behavioral economics. Journal of marketing Research, 43(3), 307-331.
17. Hoofnagle, C. J., & Urban, J. M. (2014). Alan Westin's Privacy Homo Economicus. Wake Forest Law Review, 47, 102-316.
18. Hui, K. L., Teo, H. H., & Lee, S. Y. T. (2007). The value of privacy assurance: an exploratory field experiment. Mis Quarterly, 19-33.
19. Jensen, C., Potts, C., & Jensen, C. (2005). Privacy practices of Internet users: self-reports versus observed behavior. International Journal of Human-Computer Studies, 63(1), 203-227.
20. John, L., Acquisti, A., & Loewenstein, G. (2011). Strangers on a plane: Context dependent willingness to divulge personal information. Journal of Consumer Research, 37(5), 858-873.
21. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 263-291.
22. Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. The journal of economic perspectives, 193-206.
23. Klopfer, P.H. & Rubenstein, D.I. (1977). “The Concept Privacy and its Biological Basis,” Journal of Social Issues, 33(3), 52-65.

24. Knetsch, J. L., Tang, F. F., & Thaler, R. H. (2001). The endowment effect and repeated market trials: Is the Vickrey auction demand revealing?. Experimental Economics, 4(3), 257-269.
25. Kőszegi, B., & Rabin, M. (2006). A model of reference-dependent preferences.The Quarterly Journal of Economics, 1133-1165.
26. Leontiadis, I., Efstratiou, C., Picone, M., & Mascolo, C. (2012). Don't kill my ads!: balancing privacy in an ad-supported mobile application market. In Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications (p. 2). ACM.
27. Liberman, V., Samuels, S. M., & 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), 1175-1185.
28. Loewenstein, G., & Adler, D. (1995). A bias in the prediction of tastes. The Economic Journal, 929-937.
29. Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' information privacy concerns (IUIPC): the construct, the scale, and a causal model. Information Systems Research, 15(4), 336-355.
30. Milberg, S. J., Burke, S. J., Smith, H. J., & Kallman, E. A. (1995). Values, personal information privacy, and regulatory approaches. Communications of the ACM, 38(12), 65-74.
31. Milne, G. R., & Gordon, M. E. (1993). Direct mail privacy-efficiency trade-offs within an implied social contract framework. Journal of Public Policy & Marketing, 206-215.
32. Miyazaki, A. D., & Krishnamurthy, S. (2002). Internet seals of approval: Effects on online privacy policies and consumer perceptions. Journal of Consumer Affairs, 36(1), 28-49.
33. Moon, Y. (2000). Intimate exchanges: Using computers to elicit self-disclosure from consumers. Journal of consumer research, 26(4), 323-339.
34. Novemsky, N., & Kahneman, D. (2005). The boundaries of loss aversion.Journal of Marketing Research, 42(2), 119-128.
35. O'Donoghue, T., & Rabin, M. (2000). The economics of immediate gratification.Journal of Behavioral Decision Making, 13(2), 233-250.
36. Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: an interdisciplinary review. MIS quarterly, 35(4), 989-1016.
37. Smith, H. J., Milberg, S. J., & Burke, S. J. (1996). Information Privacy: Measuring Individuals' Concerns About Organizational Practices. MIS quarterly,20(2).
38. Spiekermann, S., Grossklags, J., & Berendt, B. (2001). E-privacy in 2nd generation E-commerce: privacy preferences versus actual behavior. In Proceedings of the 3rd ACM conference on Electronic Commerce (pp. 38-47). ACM.
39. Steelman, Z. R., Hammer, B. I., & Limayem, M. (2014). Data collection in the digital age: innovative alterantives to student samples. Mis Quarterly, 38(2), 355-378.
40. Stutzman, F., Gross, R., & Acquisti, A. (2013). Silent listeners: The evolution of privacy and disclosure on facebook. Journal of Privacy and Confidentiality,4(2), 2.
41. The Federal Trade Commission (FTC). (2012). Protecting consumer privacy in an era of rapid change: recommendations for businesses and policy makers. http://www.ftc.gov/os/2012/03/120326privacyreport.pdf.
42. Tourangeau, R. (2004). Survey research and societal change. Annu. Rev. Psychol., 55, 775-801.
43. Tsai, J. Y., Egelman, S., Cranor, L., & Acquisti, A. (2011). The effect of online privacy information on purchasing behavior: An experimental study. Information Systems Research, 22(2), 254-268.
44. Viswanathan, S., Kuruzovich, J., Gosain, S., & Agarwal, R. (2007). Online infomediaries and price discrimination: Evidence from the automotive retailing sector. Journal of Marketing, 71(3), 89-107.
45. Westin, A. F. (2000). Intrusions: Privacy tradeoffs in a free society. Public Perspective, 11(6), 8-11.

46. Xu, H., Teo, H. H., Tan, B. C., & 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), 135-174.
47. Xu, H., Teo, H. H., Tan, B. C., & 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), 1342-1363.
48. Zeger, S. L., & Liang, K. Y. (1986). Longitudinal data analysis for discrete and continuous outcomes. Biometrics, 121-130.

APPENDICES

A.1 Experiment 1- Summary Questions

Measure	Description
Privacy Concern	I would be concerned about my privacy if I was a participant in this upcoming study.
Protection Satisfaction	I am satisfied with the protections provided in this upcoming study.
Harm Perception	I would be concerned that my responses in this upcoming study could be used to harm me.

A.2 Experiment 1- Disclosure Questions

Question	Text	Category
1	What is your annual income?	Descriptive
2	What is your sexual orientation?	Descriptive
3	What is your address?	Descriptive
4	What is your phone number?	Descriptive
5	What is your view on gay rights?	Descriptive
6	Have you every downloaded a pirated song?	Ethical
7	Have you ever flirted with someone other than your partner or spouse?	Ethical
8	Have you ever used drugs of any kind (e.g. weed, heroin, crack)?	Ethical
9	Have you ever looked at pornographic material?	Ethical
10	Have you ever made up a serious excuse, such as a grave illness or death in the family, to get out of doing something?	Ethical

A.3 Experiment 3 Privacy Notice

Privacy Notice	Notice Text
High Protection	*The analysis for this study requires that your responses are stored using a randomly assigned ID. All other information that could potentially be used to identify you (email, zip code, etc.) will be stored separately from your responses. As such, your responses to the following set of questions cannot be directly linked back to you.*
Low Protection	*The analysis for this study requires that your responses are stored using your email. As such, your responses to the following set of questions may be directly linked back to you.*

A.4 Attention Check and Study Designs

Design 1 and Attention Check:

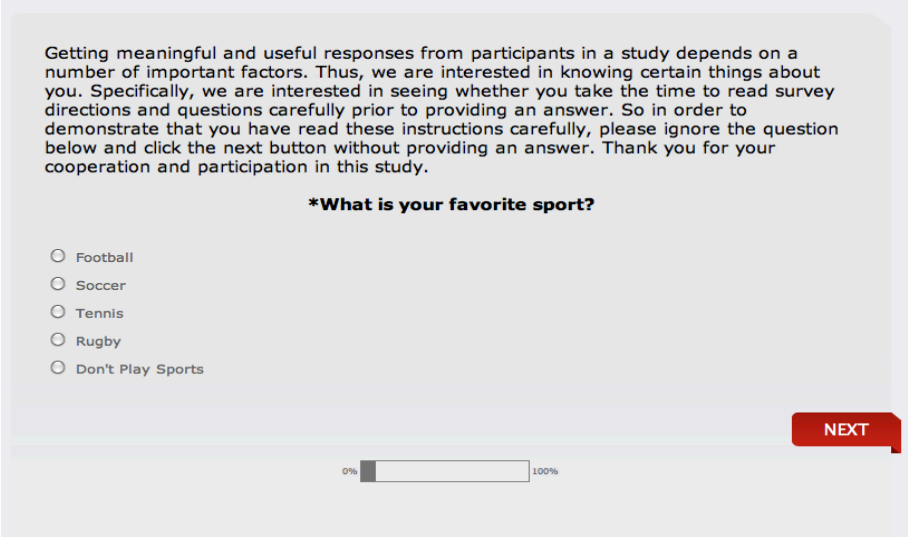

Design 2 and Attention Check:

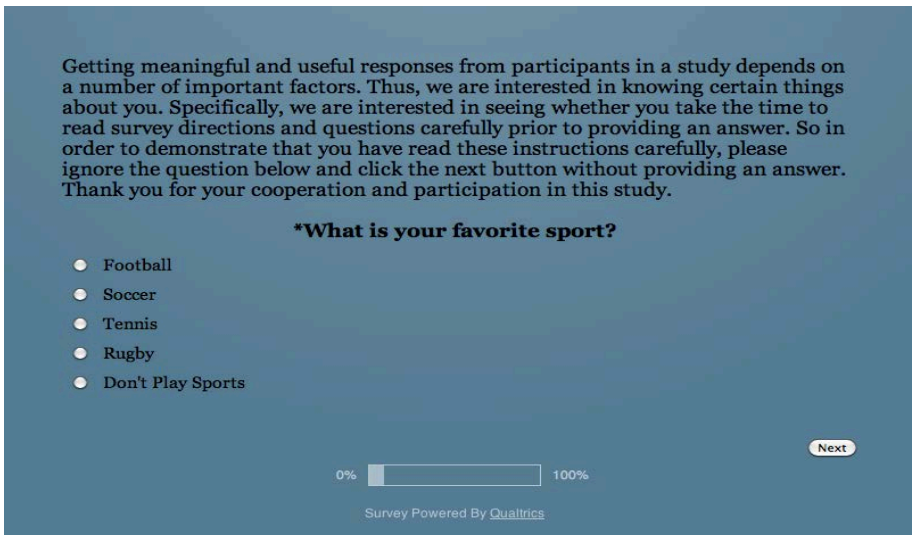

A.5 Experiment 3 Disclosure Questions (Highly Intrusive in Bold)

Question	Text
1	Have you ever downloaded a pirated song from the internet?
2	While in a relationship, have you ever flirted with somebody other than your partner?
3	**Have you ever masturbated at work or in a public restroom?**
4	**Have you ever fantasized about having violent nonconsensual sex with someone?**
5	**Have you ever tried to gain access to some else's (e.g. a partner, friend, or colleague's) email account?**
6	**Have you ever looked at pornographic material?**
7	**Have you ever used drugs of any kind (e.g. weed, heroin, crack)?**
8	Have you ever let a friend drive after you thought he or she had had too much to drink?
9	Have you ever made up a serious excuse, such as grave illness or death in the family, to get out of doing something?
10	**Have you ever had sex in a public venue (e.g. restroom of a club, airplane)?**
11	**Have you ever while an adult, had sexual desires for a minor?**
12	**Have you ever had a fantasy of doing something terrible (e.g. torture) to someone?**