Sleights of Privacy: Framing, Disclosures, and the Limits of Transparency

Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, George Loewenstein
Carnegie Mellon University
{iadjerid, acquisti, lbrandim, gl20}@andrew.cmu.edu

ABSTRACT
In an effort to address persistent consumer privacy concerns, policy makers and the data industry seem to have found common grounds in proposals that aim at making online privacy more "transparent." Such self-regulatory approaches rely on, among other things, providing more and better information to users of Internet services about how their data is used. However, we illustrate in a series of experiments that even simple privacy notices do not consistently impact disclosure behavior, and may in fact be used to nudge individuals to disclose variable amounts of personal information. In a first experiment, we demonstrate that the impact of privacy notices on disclosure is sensitive to relative judgments, even when the objective risks of disclosure actually stay constant. In a second experiment, we show that the impact of privacy notices on disclosure can be muted by introducing simple misdirections that do not alter the objective risk of disclosure. These findings cast doubts on the likelihood of initiatives predicated around notices and transparency to address, by themselves, online privacy concerns.

Categories and Subject Descriptors
K.4.1 [**Computers and Society**]: Public Policy Issues, Privacy

General Terms
Design, Human Factors

Keywords
Privacy, Privacy Policies, Notice and Choice, Behavioral Economics, Biases

1. INTRODUCTION
In response to persistent consumer privacy concerns and high profile privacy incidents [16], US policy makers have primarily resorted to two strategies. One strategy has consisted of imposing fines on organizations that used consumer data in manners deemed invasive. The other strategy has consisted of self-regulatory efforts to increase transparency about firms' data handling practices (for instance, through simple, accessible privacy policies or notices), as well as to increase consumer control over their personal information [8]. Such "transparency and control" solutions (or choice and notification regimes, as they are also called) have recently become the object of a surprisingly broad consensus between policy makers, industry, and privacy advocates. Both the FTC white paper on consumer privacy and the White House Consumer Bill of Rights [8,29] presented transparency and notice as central tenants to consumer privacy protection. Industry leaders, such as Facebook and Google, broadly concurred with the approaches outlined by policy makers. In comments on the FTC privacy framework, Facebook stated that "…companies should provide a combination of greater transparency and meaningful choice…" for consumers, and Google stated that making the "collection of personal information transparent" and giving "users meaningful choices to protect their privacy" are two of their guiding privacy principles [24]. Some privacy advocates have also embraced these approaches [23]. While researchers have highlighted the limitations of *current* privacy policies and notices [10,18], the general expectation seems to be that some *new and better* future iteration of privacy notices will solve consumers' privacy decision making issues. This manuscript presents experimental evidence that this approach, alone, may not be sufficient.

In principle, both notices and choice can certainly improve consumer disclosure decisions, while avoiding potentially burdensome regulation of firms. In normative terms, giving individuals more control over, and more information about, how their personal data is used seems an unarguable improvement over a situation in which consumers are left in the dark. In particular, policy makers posit that improved transparency will counter the status quo in which privacy concerns are secondary in online decision making, and most consumers do not read overly complex and lengthy privacy notices. Unfortunately, the ability of even improved transparency solutions or additional control tools to better align consumer attitudes towards privacy with actual behavior and reduce regret from over sharing is ultimately questionable.

As for control, the literature has started pointing out the actual impact of providing consumers with more choice and oversight over their personal information. While additional control, in principle, may allow consumers to better manage the flow and publication of their data, Brandimarte, Acquisti, and Loewenstein found in a recent paper that, in practice, an increased feeling of control over the publication of personal data can paradoxically result in increased, and riskier, disclosures [4].

As for the role of transparency and notification, this manuscript investigates the hypothesis that even simple, straightforward, and easily accessible privacy notices may not always be effective aids to privacy and disclosure decisions. Specifically, this manuscript argues that well documented and systematic biases or limitations in decision making (such as relative judgments and bounded attention) can hinder the propensity of privacy notices to achieve

Symposium on Usable Privacy and Security (SOUPS) 2013, July 24–26, 2013, Newcastle, UK.

the desired effect of supporting consumers in navigating disclosure related choices.

In a series of experiments, we find that while simple privacy notices communicating lower privacy protection can, under some conditions, result in less disclosure from participants (in line with the policy aims for increased transparency), simple and common changes in the framing of those same notices, that exploit individual heuristics and biases, can result in the effect of even straightforward and accessible privacy notices being predictably manipulated (Experiment 1) or entirely thwarted (Experiment 2).

In Experiment 1, we demonstrate that the impact of privacy notices on disclosure is sensitive to whether notices are framed as increasing or decreasing in protection, even when the objective risks of disclosure stay constant. Of particular interest is that users may effectively be led to disclose more than the level justified by objective privacy protection, and therefore face higher objective risks, if online providers put a strong emphasis on increases in privacy protection. Also, we find evidence of a diminishing propensity of privacy notices to impact disclosure over time, suggesting that notice may have an initial impact but that users may settle back into familiar disclosure habits in a short period of time. In Experiment 2, we demonstrate that the propensity of privacy notices to impact disclosure can be muted by a number of simple and minimal misdirections (such as a mere 15 second delay between notices and disclosure decisions) that do not alter the objective risk of disclosure. We argue that the sort of manipulations captured by the experimental design mimic (if anything, conservatively so) the sort of hurdles that consumers face when making real privacy decisions online. It follows that privacy notices can – on the one hand – be easily marginalized to no longer impact disclosure, or – on the other hand – be used to influence consumers to share varying amounts of personal information. Transparency may, therefore, become a "sleight" of privacy.

Such findings cast doubts on the ability of policy initiatives and design solutions built around transparency and control to, alone, address consumer privacy concerns. Note that the main implication of the results presented in this manuscript is not that notice and consent mechanisms should be avoided, or are entirely ineffective. In fact, notice and consent may be necessary conditions for meaningful privacy protection. Instead, our results suggest that, disjointed from the rest of the OECD privacy principles [20] of which they were originally part (such as purpose specification, use limitation, and accountability), transparency and choice may not be *sufficient* conditions for privacy protection. Worse, they may reduce to a case of "responsibilization" – a situation where individuals are "rendered responsible for a task which previously would have been the duty of another […] or would not have been recognized as a responsibility at all" [31].

2. THEORETICAL BACKGROUND AND HYPOTHESES

Extant privacy research [30] has highlighted that hurdles and inconsistencies in privacy decision making may be due, at least in part, to problems of asymmetric information: consumers who face privacy sensitive decisions may often be unaware of how their data is collected and used, and with what consequences. This challenge has been primarily attributed to privacy policies ineffectively communicate privacy risks to consumers. For instance, prior work has found that many privacy policies are not readable, with many policies beyond the grasp of the average internet user [10] and that privacy policies may be excessively costly to navigate [18]. To address this issue, researchers have attempted to improve the readability and "usability" of privacy policies. For example, Kelley et al. (2009) developed a "nutrition-label" style presentation of privacy policies that outperformed standard formats in readability, recall, and comprehension [15].

While evidence suggests that improved privacy notices can better inform consumers about the way their data is used, their ability to actually engage in "better" privacy decision making (that is, decisions that the consumer is less likely to later regret, or that better reflect stated preferences) is unclear. Under rational accounts of privacy decision making [22,27], predicated on the implicit premise that people can estimate stable trade-offs between privacy and other concerns, increasing the availability and comprehensibility of information should result in some increased consistency in privacy decision making. However, substantial literature in behavioral economics and decision research documents systematic inconsistencies in individuals' choices. That research shows that choice is sensitive to how choice alternatives are framed, and the salience of available information to consumers. For example, Kahneman and Tversky (1979) find 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 [14]. Moreover, Kahneman, Knetsch and Thaler (1990) find significant differences in the amount individuals are willing to pay for an item compared to individuals' willingness to accept for the same item [13]. Given that privacy's tangible and intangible consequences are often difficult to estimate, numerous heuristics and biases can influence and distort the way individuals value data protection and act on privacy concerns [1,2]. A growing body of empirical research has started highlighting the role of such systematic inconsistencies in privacy decision making. In a similar manner, heuristics may affect how consumers read, and react to, privacy notices. In this manuscript, building on the existing body of behavioral and decision research, we use two experiments to evaluate the impact of framing and bounded rationality on the propensity of privacy notices to impact disclosure.

2.1 Framing and Reference Dependence: Experiment 1

Research has highlighted that privacy concerns, and therefore propensity to disclose, are sensitive to relative judgments, which could be explained by "herding effects" (individuals being more willing to divulge sensitive information when told that others had also made sensitive disclosures) [3]; or by reference dependence, a concept introduced by Kahneman and Tversky in 1979 [14]. Kahneman and Tversky posited that outcomes are not only evaluated on their absolute value but also on their deviation from a reference point.

Framing, relative judgments, and reference dependence may also impact how individuals react to privacy notices. We argue (and test in a first experiment) that reference dependence may have a significant role in privacy decision making, since a space where consumers now make a considerable amount of privacy decisions – the online marketplace – consists of a constantly changing array of disclosure policies and privacy risks. For example, Facebook privacy settings have undergone several changes which have been presented to consumers as being increasingly protective of their

privacy. Thus, some consumers may perceive their privacy protection on Facebook as improving over time. Conversely, consumers that view Facebook's changes to default settings as less privacy protective, or encounter articles identifying Facebook's privacy infractions, may perceive their privacy protection as decreasing over time.

Under rational accounts of privacy decision making, if consumers are concerned about their personal data, privacy notices that offer low protection should elicit, on average, lower levels of disclosure relative to notices that offer sufficiently higher protection. Moreover, identical privacy notices should result, on average, in comparable levels of disclosure irrespective of relative changes in privacy notices (i.e. whether they have been increasing or decreasing over time in their level of protection). 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 they currently have (i.e. status quo). More specifically, consumers presented privacy notices that are framed as increasing in protection (i.e. preceded by notices that are less protective) would disclose more relative to those that experience no change in privacy protection, and the converse for those presented notices that are framed as decreasing in protection. As such, we posited the following hypotheses, which we test in Experiment 1:

H1a: The framing of privacy notices as increasing in their protection against privacy risks will result in an increased level of disclosure relative to no change in privacy notices.

H1b: The framing of privacy notices as decreasing in their protection against privacy risks will result in a decreased level of disclosure relative to no change in privacy notices.

Kahneman and Tversky also suggested that individuals are loss averse in that they perceive a greater dissatisfaction from losses as compared to the satisfaction from equivalent gains. Hence, we posited the following additional hypothesis:

H1c: changes in disclosure will be greater in magnitude for decreasing protection relative to increasing protection.

2.2 Bounded Rationality and Salience: Experiment 2

In a second experiment, we consider the impact on disclosure of bounded attention and privacy notices salience. Prior studies have demonstrated that attention is a limited resource and that the salience of stimuli can moderate their impact on behavior [5]. Economists have proposed that bounded attention may be a contributing factor to sub-optimal consumer decision making. Simon (1955) suggested that individuals may simplify complex decisions by focusing on a subset of the information provided, and DellaVigna (2007) has suggested that the propensity of costs to impact decisions is moderated by the degree of inattention by consumers [25, 6]. Similarly, Hossain and Morgan (2006) have found that, holding total cost constant, eBay auctions with lower initial prices (accessible cost) and high shipping costs (opaque cost) price significantly higher than the converse [9]. The authors argue that this difference is driven by the increased salience of product price as a cost relative to shipping. We extend this prior work to the context of an online disclosure experience, during which consumers are often multi-tasking and focusing on many different stimuli at once. We argue that privacy notices and the considerations they elicit from consumers are often disjoint from actual disclosure decisions via various "misdirections" – that is, actions or states that do not alter objective privacy risks but may distract consumers from them.

As a baseline, we initially consider the case in which privacy notices immediately precede disclosure decisions and no such misdirection is present. Given that privacy notices will be salient at the point of disclosure, we argue that notices will have an impact on disclosure. For instance, privacy notices communicating stronger privacy protection may result in higher levels of disclosure relative to privacy notices communicating weaker privacy protection. As such, we posited the following hypothesis, which we test in Experiment 2:

H2a: Absent a misdirection, presenting privacy notices immediately preceding disclosure decisions will have an impact on disclosure behavior, with notices presenting low protection resulting in lower levels of disclosure, on average, relative to notices offering higher protection.

We then consider misdirections that have no relevance to the risks communicated in privacy notices, but simply have the propensity to distract consumers from them. For example, consider a brief delay between the presentation of privacy notices and disclosure decisions: it may allow consumers' attention to drift away from privacy notices to other items – such as other websites or their email accounts. This distracted state may lead to a diminished impact on disclosure of privacy notices communicating privacy risks to consumers. However, we note that the directional changes in disclosure in the presence of a misdirection may be ambiguous and likely dependent on the misdirection itself. For example, a misdirection that positively impacts goodwill (e.g. reading an article on charitable organizations) may result in all affected consumers disclosing at some commonly high level, irrespective of risks communicated in privacy notices. Conversely, a misdirection that negatively impacts trust (e.g. reading an article on phishing emails) may result in all affected consumers disclosing at some commonly low level despite privacy notices. The common feature, however, is that privacy notices are no longer the main factor driving disclosure and the differences they elicit absent a misdirection should be diminished. As a result, we posited the following hypothesis:

H2b: Introducing a privacy irrelevant misdirection following the presentation of privacy notices and before disclosure decisions diminishes the propensity of privacy notices to impact disclosure behavior.

Finally, we consider privacy "relevant" misdirections. These are misdirections that relate to the privacy risks communicated in the privacy notices but only focus consumers' attention on a subset of risks. In effect, they do not alter objective privacy risk but may potentially distract from some dimensions of risk communicated in the privacy notice. An example of this type of misdirection is commonly found on online social networks, when consumers are provided granular notice and control over some dimensions of sharing and privacy preferences (e.g. access to one's personal information by other consumers of the service), but fairly minimal and less salient notice and controls (if any) over the collection and use of personal information by the service providers (e.g. Google+ or Facebook). We posit that this may result in disproportionate focus on the privacy risks from other consumers of a service and lessened focus on the providers of these services. As a result, we posited the following hypothesis:

H2c: Introducing a privacy relevant misdirection focusing on a subset of privacy risks communicated in a notice diminishes the impact on disclosure of other dimensions of risk communicated in the notice.

Following an approach that is pervasive in the experimental literature on privacy and disclosure [12,21,32], we tested our hypotheses using two survey-based experiments with random assignment, employing as main dependent variable the propensity of participants to answer personal questions in the surveys as a proxy for privacy concerns [7,26].

3. EXPERIMENT 1

Experiment 1 was a four condition between-subjects design in which we manipulated changes in privacy notices as increasing or decreasing in protection, and examined the effect of such changes on disclosure relative to conditions in which privacy notices did not change.

3.1 Procedure

Participants were recruited through Mechanical Turk, an online service that connects researchers with potential participants and is becoming increasingly popular among social scientists conducting online experiments. Participants were invited to take two online studies on ethical behavior each of which paid $.20. At the end of the first study (Survey 1), they were asked to confirm that they wished to continue to the second (ostensibly unrelated but consecutive) study (Survey 2) for an additional $.20 (all participants chose to continue to the second study but three were prohibited from completing the study because they failed our attention check). In Survey 1, participants were first asked demographic questions, which included email as a required field. Participants were told that we would check the validity of their email addresses prior to approving payment for the study (even though we did not actually store their email addresses). Then, they were provided with a simple (i.e. brief, using mundane language, and dealing only with anonymity of responses) privacy notice about the way their answers to the questionnaire would be used. Finally, participants were presented with six questions related to ethically questionable activities (See Appendix B). If they decided to take it, participants would then start Survey 2, which followed the same structure as Survey 1 (see Figure 1 for the flow of the experiment) but had a different aesthetic design (see Appendix C) to help convince participants they were participating in a separate study. In exit questions, participants confirmed that they felt they had participated in two separate studies and that it was unlikely that their responses from the first study could be linked to their responses from the second study (see Appendix H). Participants were again asked for their emails and demographic information; then, they were provided a privacy notice about the way their answers to the ensuing questions would be used; finally, they were presented with six new questions about other ethically questionable behaviors (See Appendix B). Participants were debriefed at the end of the study communicating the true nature of the experiment including the fact that we did not actually store the email addresses provided to us.

Figure 1. Flow of Experiment One

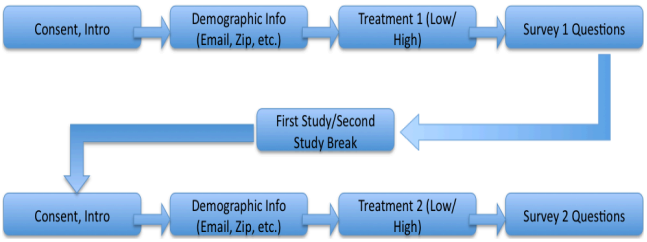

The questions used in both studies were the ones that were rated most intrusive in a 2012 paper by Acquisti, John and Loewenstein [3]. The number and type of questions were kept constant across conditions, but the order of questions was randomized within each survey.

3.2 Design

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 (randomly assigned) 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). Figure 2 provides an overview of the experimental design. The different levels of protection depended on the degree of possible linkage between the participants' emails and their responses, and therefore different levels of protection on their responses. Specifically, participants offered "low" protection to their responses were informed in the privacy notice that their answers would be linked to their email accounts. Conversely, those offered "high" protection to their responses were informed in the privacy notice that their answers would not be linked to their email accounts (See Appendix A for text of notices). We find large and significant differences between participants in high and low protection conditions in terms of their perceived identifiably of their responses (See appendix H for exit questions), with participants in the low protection condition reporting a significantly higher perception that their responses are linked back to them (78/100 vs. 14/100, P<.0001).

Figure 2. Design of Experiment 1

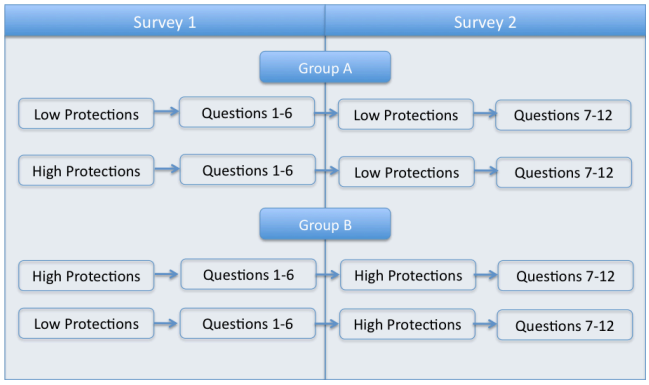

This design allowed us to evaluate both the baseline impact of protection using responses in Survey 1 and how the impact of protection may change over time, as some participants were provided identical protection in Survey 2. However, and importantly, the key feature of this design is that in Survey 2 for both the decreasing and increasing protection conditions,

participants actually faced identical privacy notices as their respective comparative conditions – thereby allowing us to evaluate the impact of the change in privacy notices on disclosure in Survey 2.

3.3 Analytical model

We used a panel random effects Probit estimation approach to evaluate the overall differences in the propensity to admit to unethical behavior across conditions. We estimated the following model:

$$Admit_{ij} = \beta_0 + \beta_1 * Treatment_i + \beta_2 * Survey1Sharing_i + \beta_3 * Intrusive_j + \beta_4 * Intrusive_i * Treatment_i + \beta_5 * Age_i + \beta_6 * Male_i + \beta_7 * Design1_i + u_{ij}$$

$Admit_{ij}$ measures the propensity to disclose, with a value of 1 if the participant admitted to the behavior and 0 if she denied or skipped, i={1,…,N participants per interaction set}, and j={1,..,12 questions}. *Treatment* is a binary indicator of the presence of our treatment. For example, in the case of decreasing protection, 1 represents a participant that was assigned to a decrease in protection, while 0 represents participants that were assigned to a condition of no change from Survey 1 to Survey 2. *Survey1Sharing* is a measure of participant sharing levels in Round 1 and 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. This was included to control for the possible impact of disclosing more in the first survey on disclosure in the second survey. *Intrusive* is a binary measure of whether a question is highly intrusive or not. *Intrusive*Treatment* captures any interaction between our treatment and highly intrusive questions. Namely, it captures whether our increases or decreases in protection had a differential impact for more intrusive questions. Lastly, *Design1* controls for which survey aesthetic design participants viewed. The model assumes serial correlation between observations within a panel unit. We allow for the correlation between responses from a single participant when we estimate the variance-covariance matrix of the coefficients, assuming constant correlation between any two answers by the same individual [17].

3.4 Results

Four-hundred and thirty-six participants (M_{Age} = 30, SD=13.5; M_{Female} = .43) completed the study (Survey 1 and 2) There were no significant differences in age or gender across conditions.

All respondents were presented an attention check question similar to those in a 2009 paper by Oppenheimer, Meyvis, and Davidenko [19] to ensure participants were carefully reading directions (See Appendix C). Lastly, participants responded to exit questions that gauged both their perception of whether privacy protections increased, decreased, or stayed the same (depending on the condition) and their recall of privacy notices in both surveys (see Appendix H). A minority (12%) weren't able to accurately recall privacy notices and, thus, disagreed that protections had increased, decreased, or stayed the same. These participants were excluded from our study, leaving 386 usable survey responses, with 97 responses in the Low to Low (LL) condition, 88 responses in the High to Low (HL) condition, 108 responses in the High to High (HH) condition, and 93 responses in the Low to High (LH) condition. [1]

3.4.1 High vs. Low Protection

We first evaluated the disclosure rates of participants in Survey 1, where participants were randomized into conditions in which they were either presented high or low protection. At this point in the experiment, no participants had been presented our central manipulation of either changing or constant levels of protection from Survey 1 to Survey 2. Figure 3 shows that participants in Survey 1 were more likely to disclose for 5 of the 6 questions when they were provided high protection ($p < .05$). Normalizing for base rates of disclosure between questions, this translates to a 14% average increase in the propensity to disclose when participants were afforded high protections in Survey 1, with some questions exhibiting more than a 30% increase in the propensity to disclose (see Table 1, Column 1 for the estimated coefficient on *HighProtection*).

Figure 3: Differences in Survey 1 Disclosure

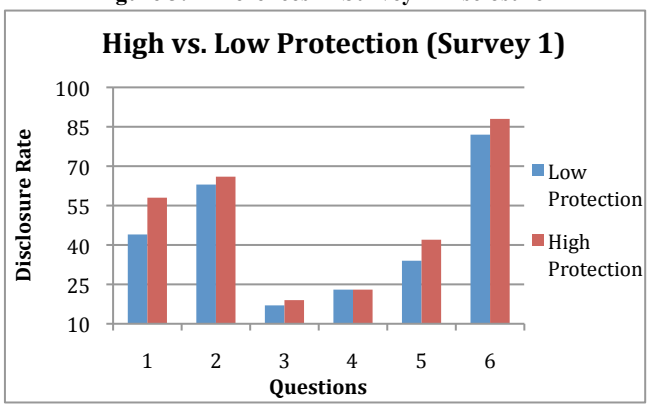

Figure 4: Differences in Survey 2 Disclosure

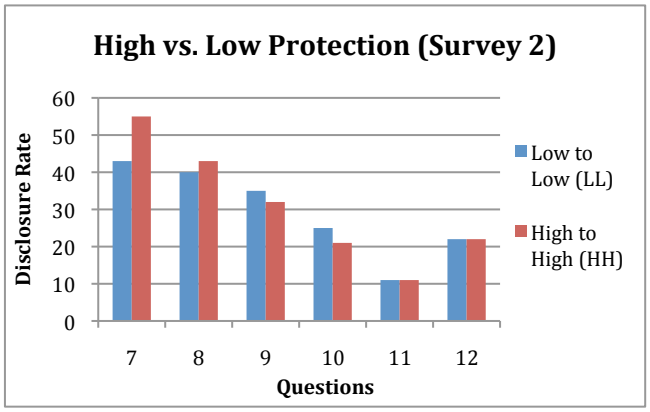

Next, we evaluated the impact of high protection relative to low protection when presented in Survey 2 . Specifically, we compare participants that had high protection in both surveys to participants (HH) that had low protection in both surveys (LL). Figure 4 shows that the impact of high protection does not extend to Survey 2 (see also Table 1, Column 2 for the estimated

[1] Inclusion of these observations results in qualitatively similar results but effects are less significant (P<.1) for participants that perceived increased protections.

coefficient on *HighProtection*, which is not significant), suggesting potentially some habituation to privacy protection and that users fall into some default mode of disclosure over time. Disclosure in Survey 2 was not systematically impacted by having high protection with only 2 of the 6 questions, demonstrating an increase in the propensity to disclose.

Table 1: Regression Results - High vs. Low Protection

	(1) – Survey 1	(2) – Survey 2
	Admit	Admit
HighProtection	0.044	-0.005
	(0.023)**	(0.029)
Intrusive	0.041	-0.111
	(0.019)**	(0.026)***
Age	-0.003	0.002
	(0.001)***	(0.001)**
Male	0.008	0.019
	(0.020)	(0.024)
Survey1Sharing		0.108
		(0.010)***
Design1	0.051	-0.004
	(0.023)**	(0.029)
Constant	-0.029	-0.003
	(0.27)	(0.062)
Observations	2634	1146
** significant at 5%; *** significant at 1%		

3.4.2 Changes in Protection

Next, we evaluated the impact of changes in protection on disclosure. Figures 5 and 6 show relative disclosure rates for each question in Survey 2 of the experiment. Figure 5 displays a trend of higher propensity to disclose (4 of the 6 questions) when participants were presented increasing protection relative to no change. Conversely, Figure 6 displays a trend of a lower propensity to disclose (4 of the 6 questions) when participants were presented decreasing protection relative to no change. In both cases, differences for questions that did not exhibit the trend were not significant.

However, base rates of disclosure for questions varied, so we also considered relative differences in the propensity to admit to a particular behavior. We found that, in the increasing protections conditions, participants were, on average, 10% more likely to disclose, with some questions having as high as 30% increase in the propensity to disclose. Similarly, for decreasing protections conditions, we found that participants were, on average, 14% less likely to disclose, with some questions having as high as a 40% reduction in the propensity to disclose.

Figure 5. Response Rates for Increasing Protection

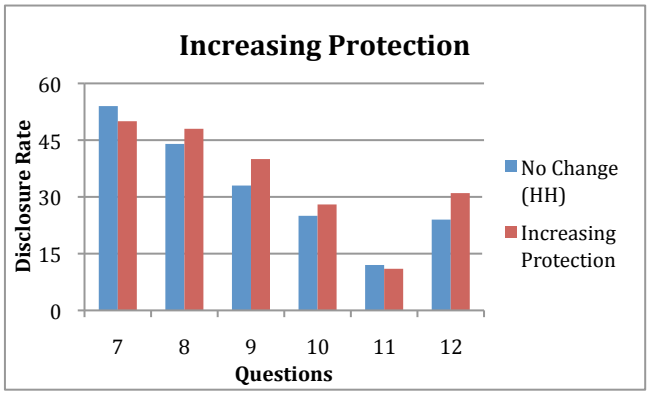

Figure 6. Response Rates for Decreasing Protection

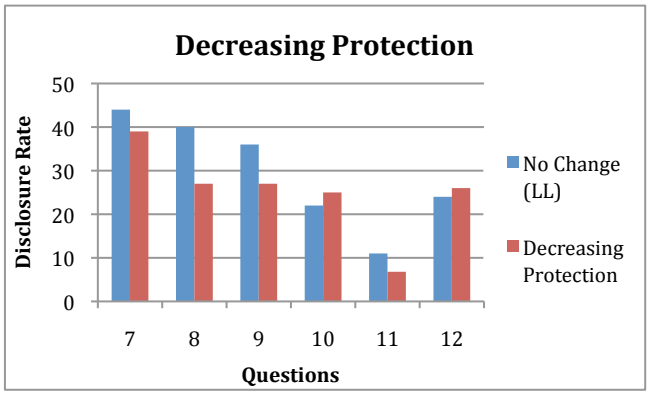

Table 2 provides estimates of the model described in the prior section. We evaluated differences in disclosure in Survey 2, where participants were provided identical privacy notices, but those in the treatment condition were presented a decrease in protection from the prior round (HL), while those in the control condition (LL) were assigned to no change in protection. Also, we evaluated differences in disclosure in Survey 2 where, again, participants were presented identical privacy notices, but those in the treatment condition were presented an increase in protection from the prior round (LH), with those in the control condition (HH) being assigned to no change in protection. Specifications (1) and (3) estimate a baseline model with only our measure of Round 1 sharing and a dummy variable for the aesthetic design viewed by participants. Specifications (2) and (4) use the full specification described above, with additional controls for the intrusiveness of questions, age and gender. We provide the estimates of the marginal effects from a random effects probit model and their associated standard errors. These estimates can be interpreted as the change in the probability of disclosure due to the random assignment to the treatment condition, holding all other covariates constant (at their mean).

Table 2: Regression Results – Changing Protection

	Decreasing Assurances		Increasing Assurances	
	(1)	**(2)**	**(3)**	**(4)**
	Admit	Admit	Admit	Admit
Treatment	-0.073**	-0.137***	0.069**	0.114**
	(0.029)	(0.047)	(0.033)	(0.051)
Round11 Sharing	0.120***	0.123***	0.106***	0.112***
	(0.011)	(0.011)	(0.012)	(0.012)
Intrusive	--	-0.152***	--	-0.100**
	--	(0.044)	--	(0.043)
Intrusive* Treatment	--	0.115	--	-0.063
	--	(0.062)	--	(0.058)
Age	--	0.002	--	0.004***
	--	(0.001)	--	(0.001)
Male	--	0.048	--	0.044
	--	(0.026)	--	(0.028)
Design1	-0.021	-0.026	-0.005	0.013
	(0.029)	(0.030)	(0.033)	(0.034)
Sigma_u_i	0.18	0.17	0.33	0.31
	(0.11)	(0.12)	(0.07)	(.08)
Rho	(0.03)	(0.03)	0.09	0.09
	(0.04)	(0.04)	(0.04)	(0.04)
Obs	1038	1038	1146	1146
** significant at 5%; *** significant at 1%				

We found that, in our basic specification, participants presented decreasing protection disclosed 7% less (P<.05) than participants that were presented no change in privacy notices, supporting H1a. Moreover, participants that were presented increasing privacy protection shared 7% more (P<.05) than participants that were presented no change in privacy notices, in support of H1b. In our baseline specification, we did not find support for the loss aversion hypothesis (H1c). In the extended specification, where we teased out the impact of the treatment on the non-intrusive questions separate from intrusive questions, we found a larger baseline effect of the treatment, with a 14% (P<.01) decrease in disclosure for participants presented with decreasing protection and a 11% increase in disclosure (P<.05) for participants presented with increasing protection (directionally consistent with H1c, although the difference is not significant).

To put our results in perspective, consider that, according to some sources, Facebook users posted 1.85 million status updates every 20 minutes, or approximately 49 trillion status updates in 2011.[2] Other disclosures (uploading a photo, posting a comment, tagging another user) on Facebook happen at comparable rates. Moreover, Facebook generally advertises changes in privacy settings and practices, often to highlight improvement to user privacy protections.[3] Now: in our experiment, we found effects that range from 10% to 14% in terms of influencing disclosure; however, it may be the case that our results are only applicable to a subset of user disclosures (e.g. disclosures with sensitive information). If we take this into account, and assume that our effects apply to only 1% of all status updates on Facebook, a 10% increase in these disclosures translates to an increase of 49 million status updates in 2011. Moreover, if the effects we identified are most applicable to sensitive information, these may in fact be the subset of disclosures most concerning for consumers.

3.5 Limitations

The results in this study rely on a self-selected sample of individuals as we mandated that users provide their email as a condition of participation in the study. This requirement may have likely excluded potential participants with high sensitivity towards the sharing of their emails. However, those individuals may have reacted even more drastically to changes in privacy notices, and their exclusion may in fact bias our effects downward. Moreover, while anonymity of responses is strongly related to the level of protection of participant responses and anonymization of sensitive data is a common mechanism for ensuring personal privacy, moving from anonymous to identified responses may also involve other changes, besides changes in privacy notices (e.g. impact of disclosures on self-perception of ethicality). Future studies may look at other variants of privacy notice that include other privacy dimensions (e.g. the length of retention of responses and breadth of access to responses).

4. EXPERIMENT 2

For Experiment 2, participants were invited to create a profile (via an online survey) on a new networking service exclusive to their university. This study was IRB approved but since we did not collect identifiable information from participants and the confidentiality protections was our central manipulation, we did not have an official consent form at the beginning of the study. Participants were debriefed after the study and told that no online social network would be created and their responses are anonymous and would only be accessible by the researchers. Participants were asked exit questions about how they thought their responses were going to be used and what the survey was about (see Appendix H). Participants that answered the question did not question the validity of the experimental context and generally reiterated the context provided in the study.

In experiment 2, we manipulated three main independent variables or factors. First, we manipulated, between subjects, the breadth of access to information disclosures communicated via a privacy notice (profile accessible only to students for the Students Only conditions, or to students and faculty for the Students and Faculty conditions). Second, we manipulated again between subjects, whether participants were presented no misdirection (i.e.

[2] See summary of Facebook usage statistics. (http://www.onlineschools.org/visual-academy/facebook-obsession).

[3] See, for instance, Facebook announcement of new privacy options in 2008 (http://blog.facebook.com/blog.php?post=11519877130); Zuckerberg's 2009 open letter about Facebook eliminating "networks" (http://blog.facebook.com/blog.php?post=190423927130); and his 2010 announcement of further privacy changes (http://blog.facebook.com/blog.php?post=391922327130).

they immediately made disclosure decisions after they were presented the notice) or were presented one of four misdirections: delay, department information pages, student committee, student committee and choice. In the context of our experiment, misdirections are actions or states that do not alter objective privacy risks but may distract consumers from them. This allows us to evaluate whether the presence of such a misdirection moderates the impact of privacy notices. Finally, all participants were asked a total of 37 questions on demographics, housing, academics, and social activities, among which were nine questions that, ex ante, we considered disproportionately sensitive to disclose to faculty relative to students. These questions asked students for their opinions on faculty, departments, and courses (e.g., "Who was your least favorite faculty member?"), dealt with the witnessing and reporting of cheating, and student effort put into academics. For free text answers (e.g. least favorite faculty member), participant responses were validated to ensure that they referred to a valid university faculty member, course, or department. The full list of questions in Experiment 2 is provided in Appendix E. The resulting experiment was a 10 condition, 2 (access) x 2 (intrusiveness) x 5 (misdirections) mixed design (see Figure 7 for flow of experiment).

Figure 7. Flow of Experiment 2

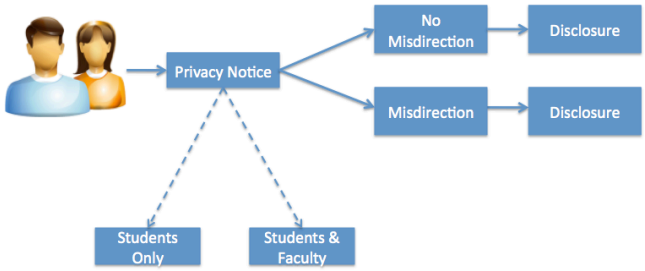

4.1 Procedure

Participants were recruited at the student center of a major North American university. They were told that they were participating in a research study which would create an online social network exclusive to the university. They were asked to create profiles using an online survey instrument which was pre-loaded on three laptops in a public space at the university student center. Participants were instructed verbally and again in the written instructions that all disclosures in the study were optional (See Appendix D for introductory text). They were compensated with a candy bar (approximate value $1.00) for their participation. Across all conditions, participants were initially presented with a privacy notice. Depending on the condition, they were either informed that their profile would only be accessible by the university students only (Students condition) or by both faculty and students (Students and Faculty condition, see Table 3). The text for each privacy notice can be found in Appendix D. Thereafter, participants in the no misdirection conditions proceeded immediately to disclosure decisions where they filled out various fields on their profile, while participants in the other conditions were presented with one of four different misdirections before proceeding to fill out the same profile fields. Participants did not have the option to go back and change their responses but were provided a link on each page which provided a pop-up box describing how their responses would be used. This link was labeled "remind me who will be able to access my responses". Twenty three participants clicked this link with no significant differences between conditions.

We first considered two misdirections without privacy relevance, in that they did not refer to the information provided in the notices nor did they deal with access to participant profiles. The first of these misdirections was a simple 15 second delay between the privacy notice and participant disclosure decisions. The second misdirection presented participants a page where they were asked if they wished to sign up for departmental information pages. We next considered two additional misdirections which were privacy relevant. The third misdirection informed participants that a student planning committee would be using their profile in order to plan upcoming activities. The fourth misdirection utilized the same student planning committee context, but provided participants control over whether this committee may access their profile. We considered these treatments as privacy-relevant in that they refocus participants' attention on the student access to their profiles. We considered them misdirections as profiles were already accessible by students under all conditions. Table 3 provides an overview of the experimental design and Appendix F provides screenshots of each individual misdirection.

Table 3. Overview of Conditions in Experiment 2

Notice	Control	Treatment	
Students	No Misdirection	Privacy Irrelevant (Delay, Dept Pages)	Privacy Relevant (Student Comm, Student Comm+Choice)
Students & Faculty	No Misdirection	Privacy Irrelevant (Delay, Dept Pages)	Privacy Relevant (Student Comm, Student Comm+Choice)

4.2 Analytical model

Similarly to the analysis in the prior experiment, we used a panel random effects Probit estimation. The results are summarized in Table 4. Again, this model assumes that responses from a single participant are serially correlated:

$$Admit_{ij} = \beta_0 + \beta_1 * Student\&Faculty_i + \beta_2 * Academic_j + \beta_3 * Misdirection_i + \beta_4 * Student\&Faculty_i * Academic_j + \beta_5 * Student\&Faculty_i * Misdirection_i + \beta_6 * Academic_j * Treatment_i + \beta_7 * Student\&Faculty_i * Academic_j * Misdirection_i + \beta_8 * Identified_i + u_{ij}$$

$Admit_{ij}$ measures the propensity to disclose, with a value of 1 if the participant answered the question and 0 if she denied or skipped, i={1,…,N participants per interaction set}, and j={1,..,37 questions}. $Student\&Faculty_i$ is a binary variable that indicates whether participants were either presented the Student and Faculty (1) or the Students Only privacy notice (0). $Academic_j$ is a binary variable indicating whether a question dealt with sensitive academic issues. $Misdirection_i$ is a binary variable that indicates whether the participant was presented a misdirection. To evaluate H2a (the impact of different privacy notices when a misdirection

is absent) we are interested in the estimates on β_4. β_4 captures the impact of the Student and Faculty notice on the disclosure of sensitive academic questions relative to non-academic questions when a misdirection was absent (i.e. our control conditions). To evaluate H2b and H2c (the impact of privacy notices when a misdirection is present) we are interested in the estimates on β_7. β_7 evaluates whether the impact of the Student and Faculty notice on the disclosure of sensitive academic questions differed when a misdirection was in place (i.e. treatment conditions). If the Student and Faculty notice impacts disclosure in the same manner irrespective of the presence of a misdirection, the estimate on β_7 will be insignificant and near zero. The other covariates are of lesser interest in our analysis but are necessary in order to correctly estimate our coefficients of interest. For example, β_1 and β_5 have analogous interpretations as β_4 and β_7 (respectively) but for non-sensitive questions. Finally, we introduced a variable, $Identified_i$, which captures whether participants chose to identify themselves, and takes a value of 1 if participants shared both their first and last names or they shared their email, and 0 otherwise. In contrast to our prior experiment, identifying information was optional in this study. We included this measure to adjust for any potential differences in participant propensity to identify themselves.

4.3 Results

Two hundred and eighty participants completed the experiment (M_{Age} = 21.5 , SD=3.1; M_{Female} =.37), with about 26 to 30 participants per condition. Figure 7 presents disclosure rates for the control conditions (i.e. without a misdirection) and for all conditions with misdirections in aggregate (this pattern is consistent for each individual misdirection as well). In the control condition, participants presented with the "Student Only" notice were 26% more likely to disclose (P<.05) relative to participants presented the "Student and Faculty" notice. In the conditions with a misdirection, we see near zero and insignificant differences in disclosure between participants presented "Student Only" and "Student and Faculty" notices. This trend is similar if we look at each misdirections individually (see Appendix E).

Figure 8. Control Conditions relative to Aggregated Misdirection Conditions

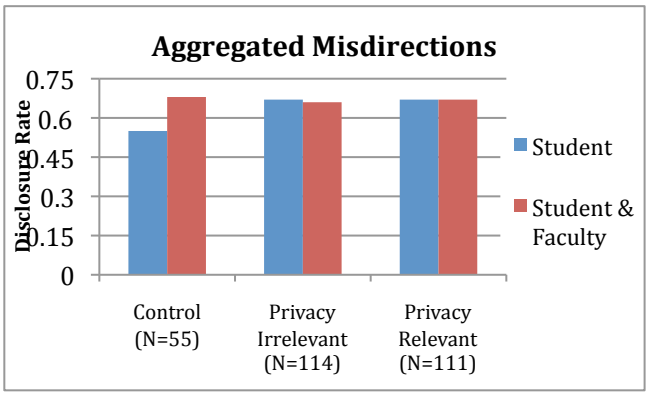

Table 4 below provides estimates of our model. We provide the estimates of the marginal effects from our random effects probit model and their associated standard errors. These can be interpreted as the impact of each covariate on the probability of disclosure holding all other covariates constant (at their mean). Estimates of the aggregate effect of all misdirections are found in column (1). Estimates distinguishing privacy-irrelevant and privacy-relevant misdirections are found in columns (2) and (3) respectively. For clarity, less relevant covariates have been excluded (estimates of the full model are available in Appendix G). Incidentally, the estimates on the other covariates were not significant. First, we found a negative and significant coefficient (P<.05) on the interaction *Student & Faculty*Academic*, indicating that, absent a misdirection (i.e. in the control condition), the Student and Faculty privacy notice had a negative impact on disclosure for sensitive academic questions relative to Students Only (H2a supported). To evaluate whether this effect is robust to the inclusion of a misdirection, we focus on the interaction *Student&Faculty* Misdirection* Academic*. As a reminder, the estimate on this measure would be zero and insignificant if the Student and Faculty notice impacted disclosure in a similar manner when a misdirection was present. Consistent with our summary results, we found a positive and significant coefficient on this interaction, indicating that the effect of the Student and Faculty notice was less pronounced when either a privacy irrelevant or privacy relevant misdirection was present (Columns 2 and 3 respectively). In fact, in support of H2b and H2c, the estimate is about equal in magnitude but in the opposite direction for both types of misdirections, suggesting that both of them muted the effect of the notice.

Table 4. Experiment 2 – Regression Results

	All Misdirection	**All Privacy Irrelevant**	**All Privacy Relevant**
	(1)	(2)	(3)
	Admit	Admit	Admit
Student&Fac* Academic	-0.109**	-0.110**	-0.111**
	(0.050)	(0.050)	(0.050)
Student&Fac* Academic* Misdirection	0.097**	0.105**	0.092**
	(0.041)	(0.045)	(0.047)
Observations	10073	6112	5959
* significant at 10%; ** significant at 5%; *** significant at 1%			

4.4 Limitations

The results in Experiment 2 rely on the control condition not being a false positive. We have since replicated this experiment and found consistent results. We also observe that, in the treatment conditions, disclosure rises disproportionally in the Student & Faculty condition relative to the Students Only. However, the current work does not evaluate underlying processes that may be driving the effect we observe. Initial analysis suggests that misdirections distracted participants from privacy concern and thus these concerns were less primed as they made disclosure decisions. Future work will evaluate this claim more directly via process-oriented studies.

5. DISCUSSION AND CONCLUSIONS

The findings we presented in this paper provide evidence of two potential inconsistencies in the impact of privacy notices on

disclosure. In our first experiment, we demonstrated that the impact of privacy notices is sensitive to reference dependence, with notices framed as increasing in protection eliciting increased disclosure and notices framed as decreasing in protection eliciting decreased disclosure. In our second experiment, we found that the downward impact of riskier privacy notices on disclosure can be muted or significantly reduced by a slight misdirection which does not alter the objective risk of disclosure.

Our results have the most applicability to online services in which user disclosure is a central function (e.g. online social networks), but also have implications for technology settings that attempt to address consumer privacy through privacy notices (e.g. online retailers). For example, Experiment 1 mimics the evolutions of Facebook's and Google's notices with respect to presenting to consumer improvements in privacy protections; Experiment 2 mimics the delays that in real life separate the reading of a privacy notice and later privacy decisions.

Policy makers and firms that deal with the exchange of consumer personal information have advocated the increased readability and usability of privacy policies as improved privacy decision aides for consumers. While these measures may provide some incremental improvements in privacy decision making, inconsistencies in decision making may result in continued disparity in consumer concerns and disclosure behavior, potentially increasing regretful disclosures by users. Our results suggest that current policy and design approaches focusing just on transparency may be limited in their ability to improve consumer privacy decision making.

The broad support for self-regulatory approaches focusing on making privacy transparent will likely make privacy notices simpler and more accessible, providing consumers certain benefits: attentive consumers concerned about their privacy may be able to better utilize short, simple, and well-formatted privacy notices to inform disclosure decision. However, the attentiveness of consumers to privacy issues may be sporadic and limited, inhibiting the usefulness of even simple and clear privacy notices. Even worse, attention paid towards self-regulatory approaches with dubious effectiveness may come at the cost of focusing on solutions that get at the heart of the privacy problem. In this regard, the experiments we presented in this manuscript illustrate the need to expand the concept of transparency to not only include clarity and ease of comprehension, but also making information communicating privacy risks salient and readily available to consumers when they most require them, at the point of disclosure.

For instance, behavioral economists and decision researchers have identified various strategies to aid consumers in improved decision making that may also be useful for privacy decision making. The 2008 book by Thaler and Sunstein [28], describes how policy makers can use soft paternalistic interventions or "nudges" to counter-act known limitations in decision making that may inhibit consumers' ability to make optimal decisions. A nudge utilizes or counteracts a known decision bias (e.g., a default effect) to urge consumers that exhibit limitations in decision making (e.g. limited attention or immediate gratification bias) towards improved decision making, while allowing rational and cognizant consumers to make informed, willful decisions. One example of a nudge is a default choice or setting, be that in the form of savings plans or organ donations. In a similar manner, research may consider providing default settings for consumers that are more protective of consumer privacy, or consider counteracting consumers' limited attention by intelligently providing relevant parts of privacy notices to consumers at the points of disclosure. For example, research may intelligently identify disclosures that consumers are likely to regret (e.g. disclosures with vulgarity or mentions of their bosses and coworkers) and remind them at that instant of the various entities that may view this particular disclosure. Similarly, prior to a consumer accepting privacy invasive terms and conditions of a particular application or service, the actual choice may be briefly delayed to allow the time for an evaluation of the trade-offs associated with such a decision, or for a reminder of intrusive uses and exchanges of their data that this particular application may engage in.

Finally, the findings presented in this manuscript may have implications for firms that collect and use personal data, particularly those with consumer personal information at the core of their business models (e.g. online advertising). Firms are likely to have significant long-term ramifications from inadequately communicating information practices to consumers, particularly if failing to do so results in high profile misuses of consumer personal information or breaches of consumers' expectations of privacy. One ramification with major implications for these firms is that consumers' propensity to disclose information may significantly change over time.

6. ACKNOWLEDGMENTS

The authors gratefully acknowledge research support from the following organizations: National Science Foundation (Award CNS-1012763), IWT SBO Project on Security and Privacy for Online Social Networks (SPION), U.S. Army Research Office under Contract DAAD190210389 through Carnegie Mellon CyLab, and TRUST (Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422).

7. REFERENCES

[1] Acquisti A. (2004). Privacy in Electronic Commerce and the Economics of Immediate Gratification. Proceedings of the ACM Conference on Electronic Commerce, New York: Association for Computing Machinery, 21–29.

[2] Acquisti A. (2009). Nudging Privacy: The Behavioral Economics of Personal Information. Security & Privacy, IEEE 7(6): 82-85.

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

[4] Brandimarte L, Acquisti A, and Loewenstein G. (2013). Misplaced Confidences: Privacy and the Control Paradox. Social Psychological and Personality Science, Volume 4, Issue: 3, 340–347. http://spp.sagepub.com/content/early/2012/08/08/1948550612455931.

[5] Broadbent DE. (1958). Perception and Communication. Elmsford, NY, US: Pergamon Press. Volume 340 pp. doi: 10.1037/10037-010.

[6] DellaVigna S. (2007). Psychology and Economics: Evidence from the Field. NBER Working Paper No 13420.

[7] Frey JH. (1986). An Experiment with a Confidentiality Reminder in a Telephone Survey. Public Opinion Quarterly. 50, 267–269.

[8] 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.

[9] Hossain T and Morgan J. (2006). Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay. The B.E. Journals in Economic Analysis and Policy: Advances in Economic Analysis and Policy. Volume 6, Issue: 2, 1–27.

[10] Jensen C and Potts C. (2004). Privacy Policies as Decision-making Tools: an Evaluation of Online Privacy Notices. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM Press, New York, NY, 471-478.

[11] John L, Acquisti A, and Loewenstein G. (2011). Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. Journal of Consumer Research. Volume 37, Issue: 5, 858-873.

[12] Joinson AN, Woodley A, and Reips UD. (2007). Personalization, Authentication and Self-disclosure in Self-administered Internet Surveys. Computers in Human Behavior. 23, 275-285.

[13] Kahneman D, Knetsch, JL., and Thaler, RH (1990). Experimental Tests of the Endowment Effect and the Coase Theorem. Journal of political Economy, 98:6, 1325-1348.

[14] Kahneman D and Tversky A. (1979). Prospect Theory: An Analysis of Decision Under Risk. Econometrica. Volume 47, Issue: 2, 263-291.

[15] Kelley PG, Bresee J, Cranor LF, and Reeder RW. (2009). A "Nutrition Label" for Privacy. SOUPS '09: Proceedings of the 5th Symposium on Usable Privacy and Security.

[16] Krazit T. (2010). Google settles Buzz lawsuit for $8.5M. CNET. http://news.cnet.com/8301-30684_3-20015620-265.html.

[17] Liang KY and Zeger SL. (1986). Longitudinal Data Analysis Using Generalized Linear Models. Biometrika. Volume:73, Issue:1, 13-22.

[18] McDonald A and Cranor L. (2009). The Cost of Reading Privacy Policies. I/S: A J. Law and Policy Inform. Soc. Volume 4, Issue:3, 543–568.

[19] Oppenheimer DM, Meyvis T, and Davidenko N.(2009). Instructional Manipulation Checks: Detecting Satisficing to Increase Statistical Power. Journal of Experimental Social Psychology. Volume 45, Issue 4, 867–872.

[20] Organisation for Economic Cooperation and Development (OECD), OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (Sep. 23, 1980).

[21] Phelps J, Nowak G, and Ferrell E. (2000). Privacy Concerns and Consumer Willingness to Provide Personal Information. Journal of Public Policy & Marketing. 19:1, 27-41.

[22] Posner R. (1981). The Economics of Privacy. American Economic Review. Volume:71, Issue:2, 405–409.

[23] Reitman R. (2012). FTC Final Privacy Report Draws a Map to Meaningful Privacy Protection in the Online World. Electronic Frontier Foundation https://www.eff.org/deeplinks/2012/03/ftc-final-privacy-report-draws-map-meaningful-privacy-protection-online-world

[24] Santalesa R. (2011). What's Next for th`e FTC's Proposed Privacy Framework? Information Law Group. http://www.infolawgroup.com/2011/03/articles/data-privacy-law-or-regulation/whats-next-for-the-ftcs-proposed-privacy-framework.

[25] Simon HA. (1955). A Behavioral Model of Rational Choice. Quarterly Journal of Economics. Volume 69, Issue:1, 99–118.

[26] Singer E, Hippler H, and Schwarz N. (1992). Confidentiality Assurances in Surveys: Reassurance or Threat? International Journal of Public Opinion Research. 4:3, 256-268.

[27] Stigler GJ. (1980). An Introduction to Privacy in Economics and Politics. Journal of Legal Studies. Volume 9, 623–44.

[28] Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

[29] The White House. (2012). Consumer Data Privacy in a Networked World. http://www.whitehouse.gov/sites/default/files/privacy-final.pdf.

[30] Tsai J, Egelman S, Cranor L, and Acquisti A. (2011). The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. Information Systems Research. Vol. 22, Issue:2, 254–268.

[31] Wakefield A and Fleming J. 2009. The Sage International Dictionary of Policing. Sage Publications, London.

[32] Weisband S and Kiesler S. (1996). Self-disclosure on Computer Forms: Meta-analysis and Implications. Proceedings of the SIGCHI Conference on Human factors in computing systems, 3-10.

8. APPENDIX

A. Experiment 1: Notices

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.*

B. Experiment 1: Questions (Highly Intrusive in Bold)

Survey 1 Questions:

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 non consensual 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?**

Survey 2 Questions:

7. **Have you ever used drugs of any kind (e.g. weed, heroine, 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**

C. Experiment 1: Attention Checks and Study Design

Design 1 and Attention Check:

In the instructions participants were instructed to skip the question. Answering this question would result in an automatic end of the study.

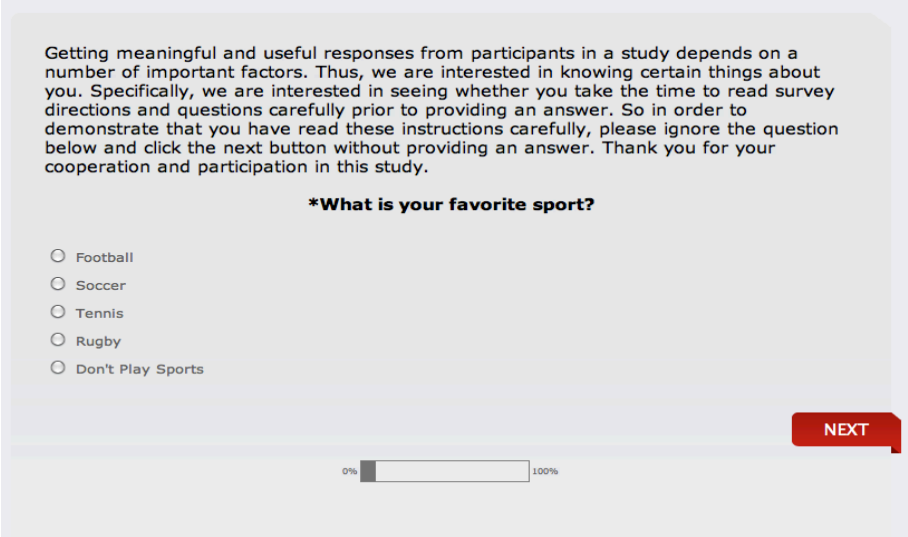

Design 2 and Attention Check:

D. Experiment 2: Introduction and Notices

Introduction:

Tartans!

We are working on a research project to create a new CMU networking website and we need your help for this launching phase!

Click on the "Next" button to participate. It'll take approximately 10 minutes.

--Page Break--

Please note that some of the questions asked (related to academics and social activities) may be somewhat sensitive. An example of a sensitive question similar to those included in these sections is:

"Have you ever made up an excuse to avoid taking an exam or handing in a term paper on time?"

You may skip any question in the survey you do not wish to answer.

Privacy Notice	Notice Text
Students Only	*The information you provide will appear on a profile that will be automatically created for you. The profile will be published on a new university networking website, which will only be accessible by university students.*
Students & Faculty	*The information you provide will appear on a profile that will be automatically created for you. The profile will be published on a new university networking website, which will only be accessible by university faculty and students.*

E. Experiment 2: Questions and Detailed Results

1. First name
2. Last name
3. Gender
4. Date of birth (MM/DD/YY)
5. Age in years
6. Country of birth
7. Email address
8. Home address
9. Phone number
10. Is your family in Pittsburgh?
11. How often do you see your family?
12. Are you single or married?
13. Do you have a girlfriend/boyfriend?
14. Where do you live?
15. Have you ever had troubles with your roommates?
16. Would you like to move somewhere else?
17. What program are you in? (e.g.: Undergrad Psychology, Grad Math)
18. Which courses are you taking at the moment?
19. **What was your least favorite class at this university?**
20. **What was your favorite class at this university?**

21. **Who was your least favorite professor?**
22. **Who was your favorite professor?**
23. **In your experience, which department at this university has the least likable faculty?**
24. **In your experience, which department at this university has the most likable faculty?**
25. **Have you ever seen someone cheating?**
26. **If so, did you inform the instructor?**
27. **How many hours a day do you spend studying?**
28. Are you working at the same time?
29. Do you receive financial aid from this university or some other non-profit organization?
30. Have you ever attended academic support programs (e.g. Peer Tutoring, Supplemental Instruction) in order to increase your understanding or your grades in a certain subject?
31. Are you a member of any group/community/fraternity/sorority?
32. If so, which group or groups are you a member of?
33. Do you have a Facebook profile?
34. Do you socialize/hang out at a bar at least once a month?
35. Do you have an alcoholic drink at least once a week?
36. In the last three months, have you done any volunteer service?
37. In the last three months, have you made a donation to a Non-Profit Organization?

[Disclosure by Individual Misdirection for Sensitive Academic Questions]

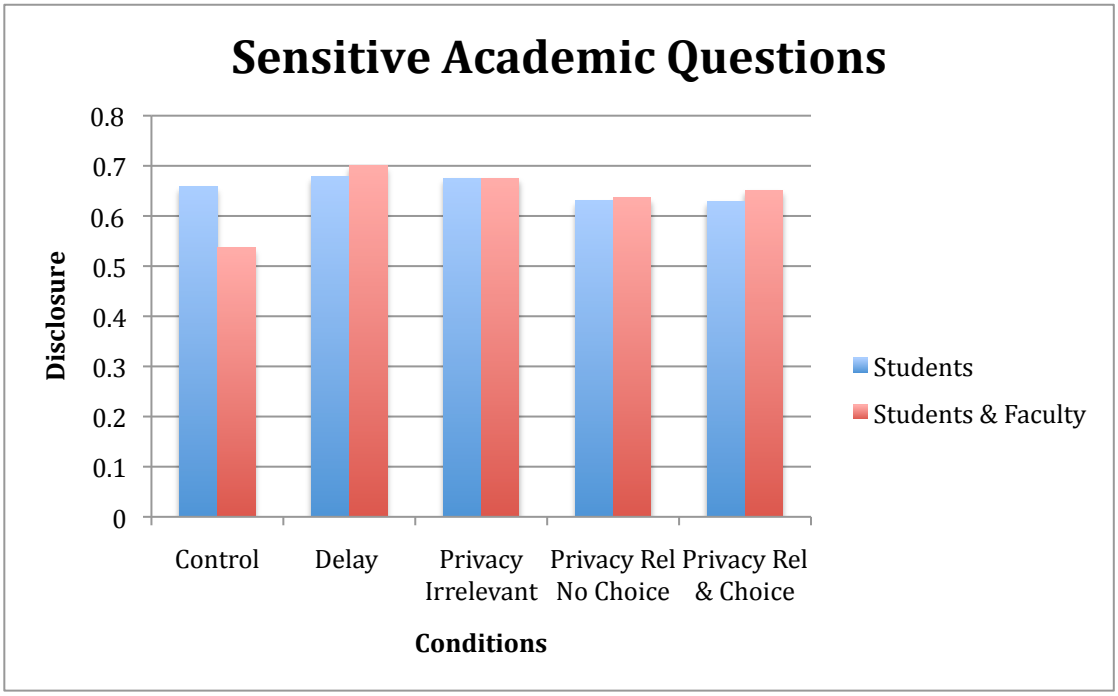

F. Experiment 2: Misdirections

Delay:

This page takes approximately 15 seconds to load, we appreciate your patience!

Department Information Pages:

A student activities planning committee, consisting of only CMU students, is requesting access to student profiles in order to better plan this year's upcoming activities. As such, they will have access to your profile.

Student Planning Committee:

We will also be creating CMU college-specific pages that will post relevant activities/lectures occurring in that school. Please select any schools below that you belong to and/or wish to be kept up-to-date on their activities.

[List of CMU Schools and Colleges]

Student Planning Committee + Choice:

A student activities planning committee, consisting of only CMU students, is requesting access to student profiles in order to better plan this year's upcoming activities.

**Do you wish to provide them access to your profile?*

G. Experiment 2: Full Regression Output

	Marginal Effects			Probit Coefficients		
	All Misdirections	Privacy Irrelevant	Privacy Relevant	All Misdirections	Privacy Irrelevant	Privacy Relevant
Student & Fac	-0.007	-0.007	-0.007	-0.022	-0.021	-0.021
	(0.034)	(0.035)	(0.033)	(0.105)	(0.106)	(0.099)
Misdirection	0.045	0.044	0.045	0.133	0.132	0.134
	(0.029)	(0.032)	(0.030)	(0.085)	(0.093)	(0.088)
Academic	-0.052	-0.052	-0.052	-0.154	-0.154	-0.153
	(0.034)	(0.035)	(0.035)	(0.100)	(0.100)	(0.100)
Stud&Fac* Acad	-0.109	-0.110	-0.111	-0.313	-0.313	-0.312
	(0.050)**	(0.050)**	(0.050)**	(0.136)**	(0.136)**	(0.136)**
Stud&Fac* Misdirection	0.016	0.007	0.026	0.049	0.021	0.078
	(0.038)	(0.043)	(0.040)	(0.117)	(0.129)	(0.122)
Misdirection* Academic	-0.050	-0.027	-0.076	-0.149	-0.081	-0.219
	(0.039)	(0.042)	(0.044)*	(0.111)	(0.122)	(0.122)*
Stud&Fac* Misdirection* Academic	0.097	0.105	0.092	0.324	0.350	0.301
	(0.041)**	(0.045)**	(0.047)**	(0.152)**	(0.167)**	(0.167)*
Identified	0.037	0.036	0.038	0.112	0.108	0.116
	(0.011)***	(0.014)***	(0.014)***	(0.032)***	(0.041)***	(0.041)***
Sigma_u_i	--	--	--	0.292	.296	0.262
	--	--	--	(0.02)	(0.03)	(0.03)
Rho	--	--	--	0.08	.08	.06
	--	--	--	(.01)	(.01)	(0.01)
Observations	10073	6112	5959	10073	6112	5959
Standard errors in parentheses						
* significant at 10%; ** significant at 5%; *** significant at 1%						

H. Select Exit Questions

Experiment 1:

1. For the CURRENT [PRIOR] study, how were your responses stored? [Random ID, Zip Code, Email Address]
2. The confidentiality protections in this study were the same as (increased relative to, decreased relative to) the confidentiality protections in the prior study.
3. For the CURRENT [PRIOR] study, to what extent could your responses be linked back to you? [Scale of 0-100 with 0 not linked at all and 100 directly linked to me]
4. How likely is it that the researchers would link your responses in the CURRENT study to your responses in the PRIOR study? [Scale of 0-100 with 0 very unlikely and 100 very likely]

Experiment 2:

1. Have you understood how your answers will be used? Please describe. [Free Response]
2. What do you think is the purpose of this survey? Please describe. [Free Response]