



Format vs. Content: The Impact of Risk and Presentation on Disclosure Decisions

Sonam Samat and Alessandro Acquisti, *Carnegie Mellon University*

<https://www.usenix.org/conference/soups2017/technical-sessions/presentation/samat-disclosure>

This paper is included in the Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017).

July 12–14, 2017 • Santa Clara, CA, USA

ISBN 978-1-931971-39-3

Open access to the Proceedings of the Thirteenth Symposium on Usable Privacy and Security is sponsored by USENIX.

Format vs. Content: The Impact of Risk and Presentation on Disclosure Decisions

Sonam Samat

Alessandro Acquisti

Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15232
{sonamsam, acquisti}@andrew.cmu.edu

ABSTRACT

Although the importance of format and presentation of privacy notices has been extensively studied in the privacy literature, less explored is the interplay of presentation and content in influencing users' disclosure decisions. In two experiments, we manipulate the content as well as the format of privacy notices shown to participants who were asked to choose whether they would like to disclose personal information. We manipulate content by changing the objective privacy risk that participants face from disclosing personal information. We manipulate format by changing the manner in which these notices are presented. We find that participants are significantly less likely to share their personal information when the privacy notice is presented under a 'Prohibit [disclosure]' frame, as compared to an 'Allow [disclosure]' frame. However, and importantly, we find that the effect of changes in framing on disclosure decisions is small when the objective privacy risk from disclosure is low, but the effect of framing becomes larger when the risk is increased—that is, for potentially more sensitive decisions. Our results highlight the nuanced interaction effects between the objective content of privacy notices and the manner in which they are presented, on disclosure behavior.

1. INTRODUCTION

Online companies collect many different types of information about their users, such as browsing behavior, search queries, purchase history, location information, and demographic information. Typically, users agree to share this information when they register for an online service and accept the service's privacy policy. In some cases, the permission to collect specific types of information is obtained after the registration process is complete, while the user is using the service. For example, a mobile app may obtain consent for collecting browsing behavior and purchase history in its privacy policy, but may later display a prompt asking for permission to collect location information while the user is employing a feature that specifically requires the use of location information. In both cases, the service provider designs the interface where the user makes his or her choice to disclose personal information—thus, it can act as a “choice architect” [21],

and influence users' decisions and behaviors. Substantial behavioral research in the privacy field has, in fact, suggested that the interface itself, and not just the content of the policy, may affect individuals' propensity to disclose personal information [e.g., 2]. Much less studied, however, is how the effect of changes in the presentation of privacy-relevant information interacts with the effect of changes in the objective privacy risk from disclosure, on individuals' propensity to disclose personal information.

In two experiments, we manipulate the content as well as the format of privacy notices shown to participants. We manipulate content by changing the privacy risk that participants face from disclosing information (for example, by varying the entity with which the information is to be shared). We manipulate format by changing the frame under which these notices are presented to the subjects. We find that participants are significantly less likely to share their personal information when the privacy notice is presented under a 'Prohibit [disclosure]' frame, as compared to an 'Allow [disclosure]' frame. However, and importantly, we also find that the effect is small when the objective privacy risk from disclosure is low, but becomes larger when the risk is increased to moderate levels. The results highlight the nuanced interactions between the actual content of privacy notices and the way they are presented, in influencing consumer behavior.

The implications of the results are twofold. First, these results highlight the challenges of relying solely on providing notice and choice to consumers to achieve a policy maker's goal of consumer privacy protection. The manner in which notices are framed can have a significant effect on behavior. As long as firms are the choice architects of their own privacy notices, they may implement framing nudges that influence consumers' choices, for instance to affect the rate of disclosure of personal information. Second, these results provide insights into identifying specific situations where framing effects matter the most (when objective risks are moderate), thus helping organizations, individuals, and policy makers direct their attention to notices that may lead to strong framing effects.

2. THEORETICAL BACKGROUND

Framing effects refer to the phenomena whereby “simple and unspectacular changes” in the presentation of decision problems lead to changes in choice [11]. These simple changes do not alter the objective factors of the decision. Evidence from behavioral decision research shows that such seemingly insignificant changes can have a significant impact on individuals' choices. In other words, they can act as “nudges.” In 1981, Tversky and Kahneman presented participants with the choice between a certain treatment that can save 200 of 600 people affected by a disease, and a

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

Symposium on Usable Privacy and Security (SOUPS) 2017, July 12 -- 14, 2017, Santa Clara, California.

probabilistic option that has the same expected value (1/3rd probability that 600 people will be saved). They found that participants prefer the certain choice to the probabilistic one, but when the same choice is framed in terms of the number of lives lost (i.e., 400 people die out of 600 affected) then participants preferred the probabilistic option to the certain choice [22]. The mere framing of the choices (in terms of lives saved or lost) had a significant impact on individuals' decisions. In 1988, Levin and Gaeth showed that labeling a pack of ground beef as "75% lean" instead of "25% fat" significantly impacted participants' perception of the quality of the beef [14]. Along similar lines, other researchers have found differences in perceptions when situations are described in terms of success rates versus failure rates [6].

Even though framing effects have been shown to significantly impact behavior in many cases, there is also no dearth of examples where framing effects have failed to change behavior ([15] provides a review of when framing effects have been shown to appear and disappear). For example, Druckman's 2001 work showed that framing effects can drastically decrease and even diminish when individuals are provided with credible advice about how to make the decision [7]. The intensity of framing effects also changes dramatically across different task domains. In 1996, Wang studied framing effects in risky choices across three different task domains, where participants had to choose between a sure outcome and a gamble of the same expected value. In addition to finding significantly different intensities of framing effects across domains, he also found that, within a given domain, framing effects tend to appear and disappear depending on the expected value of the gamble [23]. Since expected value manipulations change the objective benefits provided, this result suggests that when objective content is varied framing effects may change in intensity and may even disappear.

In the typical framing studies, participants are asked to choose between a sure option and a gambling option, while the framing of these options is manipulated between conditions. Under the positive frame, participants choose between winning an amount X for sure and a gamble in which they may win an amount greater than X but with some probability less than 1 (for example, a gamble to win $2X$ with probability $\frac{1}{2}$). Under the negative frame, participants are given some money in advance and asked to choose between giving back an amount X for sure and a gamble in which they may lose an amount greater than X but with some probability less than 1 (for example, a gamble to lose $2X$ with probability $\frac{1}{2}$). Participants tend to prefer the sure option to the probabilistic one when choices are framed as gains, and the opposite occurs when choices are framed as losses. But some researchers have found that framing effects can disappear if the payoffs from the choices are too small [9]. Therefore, it is important to investigate the interplay between objective payoffs and framing. In the case of privacy decisions, framing effects may disappear when the objective risk associated with disclosing personal information is perceived to be too low. In our studies, we test the interplay between objective risk and framing effects by varying both factors: 1) the objective risk participants face when disclosing personal information and 2) the manner in which the risks from disclosure are framed.

Previous researchers have also argued that non-normative factors, such as framing, tend to have an impact on decisions when consumers' preferences are ambiguous, but this effect diminishes (and even disappears) when preferences are more certain [19]. This is because, when preferences are ambiguous, individuals tend to

look for additional cues (such as how a problem is framed) to help them construct preferences [18, 20]. This suggests that framing may significantly impact decisions when individuals' preferences are ambiguous, and may fail to impact decisions when preferences are certain. In our work, we test this conjecture by measuring participants' level of ambiguity or uncertainty with their sharing decisions, and evaluating whether uncertainty can explain how framing effects impact disclosure decisions.

3. EXPERIMENTAL APPROACH

In two studies, we manipulate the manner in which privacy notices that inform participants about the risks from disclosure are framed. These studies were conducted on Amazon's Mechanical Turk (MTurk) platform. Previous researchers have shown that MTurk workers are more demographically diverse than the typical convenience samples of American college students, and established results have been replicated with this population, confirming its reliability [3, 4]. Amazon Mechanical Turk also allows researchers to approve or reject participants' payment based on their performance. Therefore, each participant has an approval rating, which is the percentage of his or her previously completed surveys or tasks that have been approved. We implemented a minimum requirement of a 95% approval rating during our recruitment process, and also used attention check questions in our surveys to ensure high quality data.

Our framing manipulation was embedded in the choice presented to our participants. In one condition, we asked participants if they would like to "allow us to share your information" and in another we asked if they would like to "prohibit us from sharing your information." Similar to framing manipulations used in previous literature, this manipulation allowed us to change the format of the notice while keeping the objective content of the notices constant. Query Theory research [8, 10] suggests that under the 'Allow' frame individuals will be more likely to give permission to share their information because 'Allow' frames typically make individual more likely to think about reasons to act as described in the question (in our context, allow the sharing of their information); whereas individuals will be less likely to give permission to share their information under the 'Prohibit' frame because it makes them more likely to think about reasons to prohibit the sharing of their information. Therefore, we hypothesize that participants in the 'Allow' condition will be more likely to accept the privacy policy than those in the 'Prohibit' condition.

Based on the literature highlighted in the previous section, we expect that, at low levels of objective privacy risk, individuals' sharing decisions will not vary with the framing of the notice. On the other hand, at moderate levels of disclosure risk, framing will have a significant effect on individuals' sharing decisions. We tested this hypothesis by varying the objective risk from disclosing personal information between conditions, and comparing the impact of framing at the different objective risk levels. In addition, we also tested the conjecture that framing effects depend on the level of ambiguity or uncertainty that participants have towards these disclosure decisions. Specifically, we investigated whether framing effects only tend to appear when participants are less certain about their preferences. All our studies were approved by the Institutional Review Board at Carnegie Mellon University. The IRB review covered the deception that we implemented in our studies. All participants were debriefed at the end of the study to

clarify that their data will not be shared with anyone other than the researchers conducting this study.

4. STUDY 1

4.1 Methods

In this study, we manipulated risk and framing in a context that involved making real information disclosure decisions. Participants were recruited from Amazon's Mechanical Turk for a study about ethical behaviors. Following the design put forward by Adjerid et al. in their 2014 work, we first asked participants for their demographic information, and informed them that they would be asked several questions of a sensitive nature, such as "Have you ever had a one-night stand?" The goal of collecting demographic information before presenting the disclosure choice was to elicit a level of quasi-identifiability, so the subsequent disclosure decisions would not seem entirely risk-free [2]. Then we asked participants to make a disclosure choice for whether they would be willing to share their responses to the ethical behavior questions with a specific audience. We manipulated objective privacy risk by changing this audience: in one condition, the audience was research assistants for this study, whereas in the other condition the audience was a marketing company. We expected that participants would perceive sharing information with research assistants as not very risky, but sharing with a marketing company would be perceived as being somewhat risky. Framing was manipulated using the 'Allow' and 'Prohibit' frames. For instance, participants in the 'Marketing Company' and 'Allow' frame condition were shown the following sentence: "Allow my responses to be shared with a marketing company." As participants were aware that they were going to be asked a set of highly intrusive questions, the decision to share their responses with the specified audience involved evaluating actual risks, as opposed to hypothetical ones. Following this question, we asked them ten questions related to ethically questionable activities.¹ Participants were informed that if they were not comfortable answering any of these questions, they could skip them and proceed with the survey without any penalty. Note that while participants were told that their information would be shared with the specified audience, we did not actually share their information with anyone outside of the primary researchers associated with this study. Participants were debriefed about our real motivations at the end of the study.

4.2 Results

Three hundred and seventy-six individuals (Mean Age = 32.5; 58% Male) from Amazon's Mechanical Turk participated in Study 1. Participants in the 'Allow' condition were 21% more likely to share their responses compared to those in the 'Prohibit' condition (86% vs. 71%, $\chi^2(1) = 12.22$, $p < 0.001$). Furthermore, participants in the 'Research Assistants' condition were 55% more likely to share their responses when compared to those in the 'Marketing Company' condition (96% vs. 62%, $\chi^2(1) = 64.74$, $p < 0.001$). Therefore, we observe main effects of framing and risk.

Looking at the two risk conditions individually, we find that there is no significant effect of framing among participants in the

'Research Assistants' condition (98% vs. 93%, $\chi^2(1) = 2.27$, $p = 0.132$) but the effect is significant in the 'Marketing Company' condition (74% vs. 49%, $\chi^2(1) = 11.77$, $p = 0.001$). In other words, at the level of risk presented in the 'Research Assistants' condition, almost everyone is willing to share his or her responses irrespective of framing. Therefore, we find that when the objective privacy risk is low, our framing manipulation does not significantly impact disclosure rates.

Next, we test the following econometric model:

$$Share_i = \beta_0 + \beta_{ResearchAssistants} ResearchAssistants + \beta_{Allow} Allow + \beta_{ResearchAssistants*Allow} ResearchAssistants*Allow + \epsilon_i$$

where 'Share' represents whether participants gave permission to share their data or not (binary choice), 'ResearchAssistants' is a dummy variable equal to 1 for the 'Research Assistants' condition, 'Allow' is a dummy variable equal to 1 for the 'Allow' condition, 'ResearchAssistants*Allow' is the interaction term, and ' ϵ ' is the random error term. Our dependent variable is a binary choice so we estimate the model as a probit. Since interaction terms are easier to interpret with linear regression coefficients, we report the OLS regression coefficients here. Probit model results are consistent with the OLS results. Both probit and OLS complete results are reported in Appendix B.

Estimation of the model without the interaction term confirms the main effects of framing and objective risk ($\beta_{Allow} = 0.145$, $p < 0.001$; $\beta_{ResearchAssistants} = 0.340$, $p < 0.001$). The model with the interaction term also shows a significant and positive coefficient on 'Allow' ($\beta_{Allow} = 0.242$, $p < 0.001$), indicating that framing has a significant effect on sharing decisions in the 'Marketing Company' condition. A significant and positive coefficient on 'Research Assistants' ($\beta_{ResearchAssistants} = 0.440$, $p < 0.001$) indicates that the level of objective privacy risk affects decisions to share in the 'Prohibit' condition. The interaction term in this model is also statistically significant ($\beta_{ResearchAssistants*Allow} = -0.197$, $p = 0.009$), confirming that the intensity of the framing effect is significantly smaller in the 'Research Assistants' condition than in the 'Marketing Company' condition. Figure 1 shows these effects.

This interaction effect provides evidence for the argument that the intensity of the framing effect depends on the level of risk that participants are faced with. When the risk of information sharing is low (such as sharing survey responses with research assistants) then framing effects may disappear, but they appear when the risk of information sharing is relatively higher (such as sharing survey responses with a marketing company). In addition, these results also suggest that an increase in objective privacy risk under the 'Allow' frame leads to a smaller adjustment of sharing behavior, as compared to an equivalent increase in risk under the 'Prohibit' frame. This implies that, when privacy policies are framed in a positive way (as most current day privacy policies are), individuals may be less likely to adjust their sharing behavior to account for an increase in objective privacy risk, compared to cases when privacy policies are framed in a negative way. This result is important as companies frequently make changes to their privacy policies and

¹ We used questions rated as highly intrusive in Acquisti et al.'s 2012 work [1]. See Appendix A.

these changes often involve increasing the privacy risk for consumers.

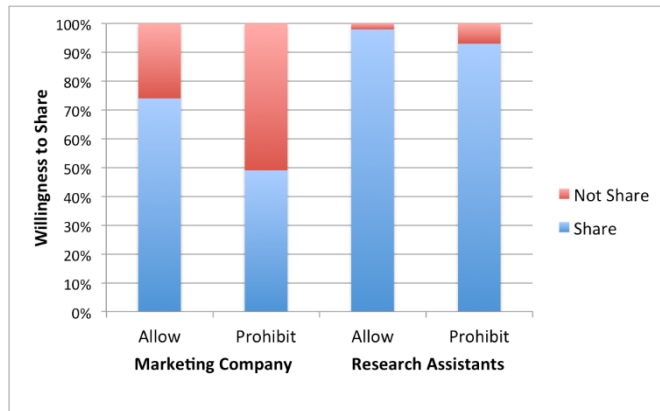


Figure 1. Percentages of participants willing to share their information in Study 1.

4.3 Discussion

The results from Study 1 provide evidence that framing nudges in privacy notices can influence disclosure decisions, but these effects are not universal across different levels of objective privacy risk. Positive frames such as ‘Allow’ make participants more likely to share their personal information when compared to negative frames such as ‘Prohibit’. But there may be credible levels of objective privacy risk at which individuals overcome the effect of framing, because the risk associated with sharing is perceived to be too low (such as sharing survey responses with research assistants). When the objective privacy risk associated with sharing information is moderate (such as sharing with a marketing company), framing effects are more likely to occur. This result is consistent with previous research that showed framing effects tend to disappear when objective payoffs from a gamble are small [9, 12, 13].

5. PRELIMINARY ANALYSIS

Following Study 1, we conducted a set of additional studies to measure the level of perceived risk associated with sharing personal information with different entities. These studies gave us insight into the relative perceptions of risk associated with the two conditions tested in Study 1 as well as other conditions which we later tested in Study 2.

5.1 Survey A

Participants were recruited from Mechanical Turk for a study in which they would be asked for their opinions about a hypothetical scenario. Participants were asked to imagine that they were answering a survey about ethical behaviors on Mechanical Turk, which involved answering sensitive questions such as “Have you

ever had a one-night stand?” They were then told that the researchers of that study first asked them whether they would be willing to share their information with a specific audience. We manipulated this audience between conditions, testing four different audiences: 1) research assistants, 2) marketing companies, 3) publicly on the Internet, and 4) on Mechanical Turk forums (such as Turker Nation, MTurk Grind, MTurk Forum, etc.) along with their Mechanical Turk ID. The last condition was included in an attempt to increase the personal relevance of the decision for Mechanical Turk participants, as individuals on Mechanical Turk often use such forums to discuss pragmatic considerations about MTurk tasks such as pay rates or requesters’ reputations [5]. Arguably, they may care a significant amount about their reputation on these forums. As Mechanical Turk does not permit the collection of any personally identifiable information, we attempted to achieve quasi-identifiability by claiming that the responses would be shared along with participants’ Mechanical Turk ID. Next, participants were asked four questions, each on a 1–7 scale from ‘Not at all’ to ‘Very much’: 1) how risky they thought it would be to share their survey responses with this audience, 2) how likely they would be to share their survey responses with this audience, 3) how comfortable they would be sharing their survey responses with this audience, and 4) how concerned they would be about sharing their survey responses with this audience.

5.1.1 Results

One thousand, two hundred seventy-nine participants (Mean Age = 32.3; 55% Male) from Amazon Mechanical Turk completed this survey. As shown by the mean values reported in Table 1, participants did not think sharing their survey responses with research assistants was very risky. Sharing survey responses with marketing companies, publicly online, and on Mechanical Turk forums with Mechanical Turk IDs were perceived to be increasingly riskier scenarios, in that order.² These results confirm the assumption made in Study 1, that the perceived risk of sharing survey responses with research assistant is very low risk while that of sharing survey responses with marketing companies is moderate.

Table 1. Mean values of variables measured in Survey A.

Dependent Variable	Research Assistants	Marketing Companies	Publicly Online	On MTurk with Turk ID
Risky	2.46	3.75	4.32	4.44
Likely	5.67	3.98	3.12	2.88
Comfortable	5.37	4.11	3.32	3.22
Concerned	2.60	3.84	4.56	4.53

² The differences between ‘Publicly Online’ and ‘On MTurk with Turk ID’ for all four variables are not statistically significant. All other differences are statistically significant at the 0.001 level.

5.2 Survey B

Participants were recruited for a survey about information sharing preferences. Each participant was shown five scenarios, two of which involved sharing information with a hypothetical news website and three of which involved sharing survey responses to ethical behavior questions with a specific audience. Therefore, this study was conducted with a within-subjects design as opposed to the between-subjects design of Survey A. In the two scenarios involving the sharing of information with a hypothetical news website, participants were asked to imagine that they want to read an article on a news website but are faced with the decision to accept or reject the news website's privacy policy before they can read the article. We tested participants' opinions about two different amounts of personal information purportedly being collected by the news website. (These two scenarios are used in studies that are not reported in this paper.) The scenarios involving sharing survey responses with specific audiences were presented in the same way as in Survey A (by asking participants to imagine they are answering a survey about ethical behaviors). The three audiences tested in this survey were: 1) research assistants, 2) other research organizations, and 3) marketing companies. We asked participants how risky they thought it would be to share the information and how likely they would be to share it, on a 1–7 scale ranging from 'Not at all' to 'Very much.' The same two questions were used for the hypothetical news website scenarios as well, and the order of these two questions was randomized across participants.

5.2.1 Results

One hundred twenty participants (Mean Age = 34.3; 67% Male) from Amazon Mechanical Turk completed this study. As shown by the mean values reported in Table 2, sharing survey responses with research assistants, other research organizations, and marketing companies are perceived to be increasingly riskier scenarios, in that order.³

Table 2. Mean values of variables measured in Survey B.

Dependent Variable	Research Assistants	Other Research Organizations	Marketing Companies
Risky	2.53	3.24	4.03
Likely	5.43	4.70	3.56

6. STUDY 2

6.1 Methods

Based on the results of the surveys we designed the second study. In Study 2, we test framing effects at three different risk levels to get a better sense of how framing effects vary with risk. The design of this study is similar to that of Study 1, but instead of a 2-by-2 design, we used a 2-by-3 design. Two objective privacy risk levels tested in this study are similar to the ones used Study 1: sharing with research assistants and sharing with marketing companies. The third risk level is sharing on Mechanical Turk forums (such as

Turker Nation, MTurk Grind, MTurk Forum, etc.), along with their Mechanical Turk ID. The framing manipulation is implemented in the same way as in Study 1, by altering whether participants are asked to 'Allow' the sharing of their information or 'Prohibit' the sharing of their information. Just as in Study 1, participants were first informed that they were going to be asked sensitive questions. Then, they were asked for their sharing preferences (varying the scenarios in their framing and risk between conditions) and were subsequently asked the same ethical questions as used in the previous study. We include an additional question in this study that measures participants' level of uncertainty with the sharing decision, using the uncertainty subscale from the Decision Conflict Scale [16]. This question is included to test whether participants' level of uncertainty with the sharing decision varies in the same way as framing effects vary when objective risk is manipulated (i.e., more uncertainty correlates with larger framing effects). Such a result would provide support for the conjecture that framing effects only occur when individuals' preferences are ambiguous, consistent with previous research [19]. Participants were debriefed at the end of the study.

6.2 Results

Nine hundred ninety-five individuals (Mean Age = 37.2; 47% Male) from Amazon Mechanical Turk completed Study 2. Participants in the 'Allow' condition were 43% more likely to share their responses when compared to those in the 'Prohibit' condition (73% vs. 51%, $\chi^2(1) = 51.95, p < 0.001$). Furthermore, participants in the 'Research Assistants' condition were 64% more likely to share their responses when compared to those in the 'Marketing Company' condition (92% vs. 56%, $\chi^2(1) = 109.26, p < 0.001$), and participants in the 'Marketing Company' condition were 37% more likely to share their responses when compared to those in the 'Mechanical Turk Forums' condition (56% vs. 41%, $\chi^2(1) = 14.64, p < 0.001$). Therefore, we observe the main effects of framing and risk.

Looking at the three objective privacy risk conditions individually, we find a significant effect of framing at all three risk levels ('Research Assistants' condition: 96% vs. 87%, $\chi^2(1) = 7.84, p = 0.005$; 'Marketing Company' condition: 67% vs. 44%, $\chi^2(1) = 17.42, p < 0.001$; 'Mechanical Turk Forums' condition: 58% vs. 22%, $\chi^2(1) = 45.11, p < 0.001$). The relative size of the framing effect increases as risk increases (Cohen's d for 'Research Assistants' condition = 0.31; Cohen's d for 'Marketing Company' condition = 0.47; Cohen's d for 'Mechanical Turk Forums' condition = 0.78). Therefore, we find that when objective risk from disclosure is low, our framing manipulation has a smaller impact on sharing decisions, but as the objective risk increases, the effect of our framing manipulation also increases. It is important to note that even the highest level of risk tested in this study ('Mechanical Turk Forums') is perceived to be 'moderate', as shown by the mean value of perceived risk in Survey A (mean perceived risk for 'Mechanical Turk Forums' = 4.44 on a 1–7 scale). So, the increasing trend in the size of framing effects is observed when risk increases from low to moderate levels. We do not know how framing effects may vary when the perceived risk increases beyond moderate levels to high levels.

³ All differences are statistically significant at the 0.001 level.

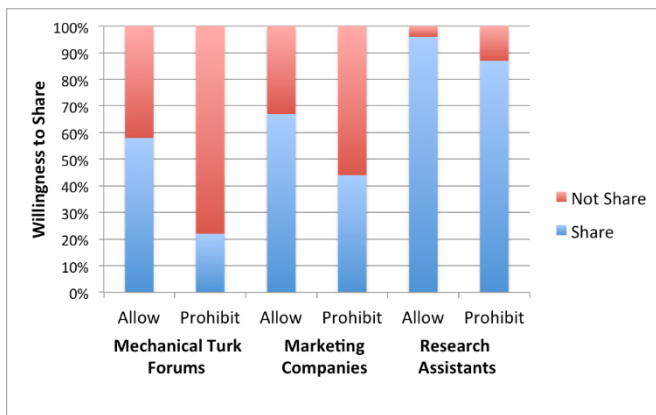


Figure 2. Percentages of participants willing to share their information in Study 2.

Estimation of the model as in Study 1 without the interaction term (coding the risk variable as ‘0’ for ‘Research Assistants’, ‘1’ for ‘Marketing Companies’, and ‘2’ for ‘Mechanical Turk Forums’) confirms the main effects of framing and objective risk ($\beta_{\text{Allow}} = 0.226, p < 0.001$; $\beta_{\text{Risk}} = -0.255, p < 0.001$). The model with the interaction term also shows a significant and positive coefficient on ‘Allow’ ($\beta_{\text{Allow}} = 0.087, p = 0.039$), indicating that framing has a significant effect on sharing decisions in the ‘Research Assistants’ condition. A significant and negative coefficient on ‘Risk’ ($\beta_{\text{Risk}} = -0.325, p < 0.001$) indicates that the level of risk affects decisions to share in the ‘Prohibit’ condition (increasing risk decreases willingness to share). The interaction term in this model is also statistically significant ($\beta_{\text{Risk} * \text{Allow}} = 0.137, p < 0.001$), confirming that the intensity of the framing effect significantly changes as risk is varied. The positive coefficient on the interaction term confirms that the size of the framing effect increases as the level of risk increases. Figure 2 shows these effects graphically.

Next, we analyze how participants’ uncertainty varies with objective risk, to see if uncertainty can explain the increase in the size of framing effects as risk increases. The three items in the uncertainty subscale of the Decisional Conflict Scale [16] show high reliability (Cronbach’s alpha = 0.927) so we follow the instructions provided by O’Connor [17] to code the 1–5 scale (from ‘strongly agree’ to ‘strongly disagree’) as 0–4, then average the three items and multiply the averaged value by 25. The final score ranges from 0 (participant feels extremely certain about the best choice for them) to 100 (participant feels extremely uncertain about the best choice for them). The uncertainty values do not vary significantly with risk ($\beta_{\text{Risk}} = 0.318, p = 0.715$). The mean uncertainty values across the three risk conditions do not show an increasing trend (mean uncertainty value for ‘Research Assistants’ condition = 21.67; mean uncertainty value for ‘Marketing Companies’ condition = 24.47; mean uncertainty value for ‘Mechanical Turk Forums’ condition = 22.33). This suggests that the level of uncertainty, as measured by the uncertainty subscale of the Decisional Conflict Scale, cannot explain the increase in size of framing effects as risk increases.

6.3 Discussion

The results from Study 2 confirm an increasing trend in the size of framing effects as objective privacy risk increases. We find a small but significant effect of framing on sharing decisions when perceived risk is low (such as sharing with research assistants). The

size of the framing effect increases as risk is increased to moderate levels (such as sharing on Mechanical Turk Forums along with Mechanical Turk ID).

These results are particularly important from a policy perspective as they suggest that regulators should work towards more nuanced requirements in terms of how privacy notices ought to be framed. Individuals’ propensity to framing effects varies considerably with objective risk. A single set of blanket requirements for all websites (irrespective of the amount of privacy risk consumers face from sharing information) may not be sufficient to protect consumers’ privacy. For instance, websites that merely collect users’ IP addresses versus those that collect web browsing and purchase behaviors should not be subject to the same regulations. Our results suggest that the latter category of websites may be more capable of nudging consumers’ sharing decisions by using framing nudges, and therefore should be held to higher standards by policy makers. While it is challenging to present information in a truly ‘neutral’ frame, further work should investigate solutions to mitigate the privacy risk faced by consumers, especially in situations with moderate objective risk.

It is important to discuss the limitations of this study. First, the highest level of objective privacy risk tested in this study (sharing on Mechanical Turk Forums along with Mechanical Turk ID) is only perceived to be moderately risky by our participants. So, the observed increasing trend in framing effects can only be claimed to occur when risk increases from low to moderate levels. While we do not know for sure how framing effects would vary when risk is increased from moderate to high levels, we suspect that framing effects may decrease at very high levels of risk. Second, this study also attempted to test whether participants’ level of uncertainty can explain the increase in size of framing effects as risk is increased. We could not find evidence for this, as uncertainty did not vary significantly across the different risk levels. Further research is required to better understand why framing effects increase when objective risk is increased from low to moderate levels.

7. CONCLUSION

In two experiments, we studied the influence of a framing nudge that may be used in privacy notices to influence individuals’ willingness to disclose personal information. We found that the manner in which privacy notices are framed has a significant impact on individuals’ disclosure decisions. Consistent with our hypotheses, notices that use the ‘Prohibit’ frame reduce the likelihood that individuals will share their information as compared to notices that use the ‘Allow’ frame. However, and importantly, we also found that the intensity of this framing effect is small when objective privacy risk from disclosure is small, and it increases as the objective privacy risk increases to moderate levels.

These findings have implications for the design of privacy policies that can empower consumers to face the tradeoffs between privacy risks and the benefits associated with data sharing. Specifically, our results strengthen the notion that simply providing consumers with notice and choice may not be sufficient mechanisms to serve the goal of consumer privacy protection. The manner in which notice and choice are framed is also important as companies may use framing nudges to impact individuals’ sharing decision. Our results also assist in guiding the attention of organizations and policy makers towards cases where individuals might be most susceptible to framing nudges, specifically when the objective privacy risks are moderate. Owing to the substantially different intensities of

framing effects observed at different objective privacy risk levels, it appears that a blanket policy for all websites, irrespective of the amount of privacy risk consumers face from sharing information with the site, may not be sufficient to protect consumers' privacy. Companies and policy makers may consider more nuanced sets of rules concerning how privacy policies should be framed, keeping in mind the level of privacy risks put forward by different websites.

8. ACKNOWLEDGMENTS

This research has been in part supported by the National Science Foundation under grants 1012763, 1514192, and 1327992. We would like to thank all reviewers for their feedback and suggestions.

9. REFERENCES

- [1] Acquisti, A., John, L. K., & Loewenstein, G. (2012). "The impact of relative standards on the propensity to disclose," *Journal of Marketing Research*, 49(2), 160-74.
- [2] Adjerid I., Acquisti, A., & Loewenstein, G. (2014). "Framing and the Malleability of Privacy Choices," *In Workshop on the Economics of Information Security (WEIS) 2014*.
- [3] Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). "Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk," *In Political Analysis*, 20(3), 351-68.
- [4] 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.
- [5] Chandler, J., Mueller, P., & Paolacci, G. (2014). "Nonnaïveté among Amazon Mechanical Turk workers: Consequences and solutions for behavioral researchers," *Behavior Research Methods*, 46(1), 112-30.
- [6] Davis, M. A., & Bobko, P. (1986). "Contextual effects on escalation processes in public sector decision making," *Organizational Behavior and Human Decision Processes*, 37, 121-38.
- [7] Druckman, J. N. (2001). "Using credible advice to overcome framing effects," *Journal of Law, Economics, and Organization*, 17(1), 62-82.
- [8] Hardisty, D. J., Johnson, E. J., & Weber, E. U. (2010). "A dirty word or a dirty world? Attribute framing, political affiliation, and query theory," *Psychological Science*, 21(1), 86-92.
- [9] Hsee, C. K., & Weber, E. U. (1997). "A fundamental prediction error: Self-others discrepancies in risk preference," *Journal of Experimental Psychology: General*, 126(1), 45.
- [10] Johnson, E. J., Häubl, G., & Keinan, A. (2007). "Aspects of endowment: a query theory of value construction," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(3), 461.
- [11] Kühberger, A. (1998). "The influence of framing on risky decisions: A meta-analysis," *Organizational Behavior and Human Decision Processes*, 75(1), 23-55.
- [12] Kühberger, A., Schulte-Mecklenbeck, M., & Perner, J. (1999). "The effects of framing, reflection, probability, and payoff on risk preference in choice tasks," *Organizational Behavior and Human Decision Processes*, 78(3), 204-31.
- [13] Kühberger, A., Schulte-Mecklenbeck, M., & Perner, J. (2002). "Framing decisions: Hypothetical and real," *Organizational Behavior and Human Decision Processes*, 89(2), 1162-75.
- [14] Levin, I. P., & Gaeth, G. J. (1988). "Framing of attribute information before and after consuming the product," *Journal of Consumer Research*, 15, 374-78.
- [15] Levin, I. P., Schneider, S. L., & Gaeth, G. J. (1998). "All frames are not created equal: A typology and critical analysis of framing effects," *Organizational Behavior and Human Decision Processes*, 76(2), 149-88.
- [16] O'Connor, A. M. (1995). "Validation of a decisional conflict scale," *Medical Decision Making*, 15(1), 25-30.
- [17] O'Connor, A. M. (updated 2010). *User Manual – Decisional Conflict Scale*, https://decisionaid.ohri.ca/docs/develop/User_manuals/UM_Decisional_Conflict.pdf
- [18] Payne, J. W., Sagara, N., Shu, S. B., Appelt, K. C., & Johnson, E. J. (2013). "Life expectancy as a constructed belief: Evidence of a live-to or die-by framing effect," *Journal of Risk and Uncertainty*, 46(1), 27-50.
- [19] Schoorman, F. D., Mayer, R. C., Douglas, C. A., & Hetrick, C. T. (1994). "Escalation of commitment and the framing effect: An empirical investigation," *Journal of Applied Social Psychology*, 24(6), 509-28.
- [20] Slovic, P. (1995). "The construction of preference," *American Psychologist*, 50, 364-71.
- [21] Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving Decisions About Health, Wealth and Happiness*, New Haven and London: Yale University Press.
- [22] Tversky, A., & Kahneman, D. (1981). "The framing of decisions and the psychology of choice," *Science*, 211(4481), 453-58.
- [23] Wang, X. T. (1996). "Framing effects: Dynamics and task domains," *Organizational Behavior and Human Decision Processes*, 68(2), 145-57.

APPENDICES

A. Appendix A

Questions used in Study 2 (from Acquisti et al., 2012):

1. Have you ever had sex with the current husband, wife, or partner of a friend?
2. Have you ever masturbated at work or in a public restroom?
3. Have you ever had a fantasy of doing something terrible (e.g., torturing) to someone?
4. Have you ever fantasized about having violent non-consensual sex with someone?
5. Have you ever, while an adult, had sexual desires for a minor?

6. Have you ever neglected to tell a partner about a sexually transmitted disease from which you were suffering?
7. Have you ever had sex with someone who was too drunk to know what they were doing?
8. Have you ever stolen anything that did not belong to you?
9. Have you ever tried to gain access to someone else's (e.g., a partner, friend, or colleague's) email account?
10. Have you ever looked at pornographic material?

B. Appendix B

OLS regression coefficients for sharing choice in Study 1

	(1)	(2)	(3)
	Share	Share	Share
ResearchAssistants	0.340*** (0.038)	0.439*** (0.058)	0.441*** (0.058)
Allow	0.145*** (0.038)	0.242*** (0.069)	0.246*** (0.068)
ResearchAssistants *Allow		– 0.197*** (0.075)	–0.194** (0.075)
Male			–0.048 (0.038)
Age			0.004** (0.002)
Constant	0.544*** (0.041)	0.495*** (0.052)	0.381*** (0.082)

*p<0.10; **p<0.05; ***p<0.01
Standard errors in brackets

Probit results for sharing choice in Study 1

	(1)
	Share
	Average Marginal Effects
ResearchAssistants	0.339*** (0.037)
Allow	0.143*** (0.037)
ResearchAssistants *Allow	–0.197***
[Cross-partial derivative]	(0.075)

*p<0.10; **p<0.05; ***p<0.01
Standard errors in brackets

OLS regression coefficients for sharing choice in Study 2

	(1)	(2)	(3)
	Share	Share	Share
Risk	–0.255*** (0.016)	–0.325*** (0.023)	–0.325*** (0.023)
Allow	0.226*** (0.027)	0.087** (0.042)	0.092** (0.042)
Risk *Allow		0.137*** (0.033)	0.133*** (0.033)
Male			0.052* (0.027)
Age			–0.002* (0.001)
Constant	0.765*** (0.025)	0.835*** (0.030)	0.845*** (0.040)

*p<0.10; **p<0.05; ***p<0.01
Standard errors in brackets

Probit results for sharing choice in Study 2

	(1)
	Share
	Average Marginal Effects
Risk	–0.240*** (0.012)
Allow	0.217*** (0.025)
Risk*Allow	0.098***
[Cross-partial derivative]	(0.042)

*p<0.10; **p<0.05; ***p<0.01
Standard errors in brackets