

10-2013

Is It the Typeset or the Type of Statistics? Disfluent Font and Self-Disclosure

Rebecca Balebako

Carnegie Mellon University, balebako@cmu.edu

Eyal Peer

Carnegie Mellon University, eper@andrew.cmu.edu

Laura Brandimarte

Carnegie Mellon University, lbrandim@andrew.cmu.edu

Lorrie Faith Cranor

Carnegie Mellon University, lorrie@cmu.edu

Alessandro Acquisti

Carnegie Mellon University, acquisti@andrew.cmu.edu

Follow this and additional works at: <http://repository.cmu.edu/heinzworks>

 Part of the [Databases and Information Systems Commons](#), and the [Public Policy Commons](#)

Published In

Proceedings of LASER 2013.

This Conference Proceeding is brought to you for free and open access by the Heinz College at Research Showcase @ CMU. It has been accepted for inclusion in Heinz College Research by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Is It the Typeset or the Type of Statistics? Disfluent Font and Self-Disclosure

Rebecca Balebako, Eyal Pe'er, Laura Brandimarte, Lorrie Faith Cranor, Alessandro Acquisti
Carnegie Mellon University

Abstract

- **Background.** The security and privacy communities have become increasingly interested in results from behavioral economics and psychology to help frame decisions so that users can make better privacy and security choices. One such result in the literature suggests that cognitive disfluency (presenting questions in a hard-to-read font) reduces self-disclosure.
- **Aim.** To examine the replicability and reliability of the effect of disfluency on self-disclosure, in order to test whether such approaches might be used to promote safer security and privacy behaviors.
- **Method.** We conducted a series of survey studies on human subjects with two conditions - disfluent and fluent font. The surveys were completed online (390 participants throughout the United States), on tablets (93 students) and with pen and paper (three studies with 89, 61, and 59 students). The pen and paper studies replicated the original study exactly. We ran an independent samples t-test to check for significant differences between the averages of desirable responses across the two conditions.
- **Results.** In all but one case, participants did not show lower self-disclosure rates under disfluent conditions using an independent samples t-test. We re-analyzed the original data and our data using the same statistical test (paired t-test) as used in the original paper, and only the data from the original published studies supported the hypothesis.
- **Conclusions.** We argue that the effect of disfluency on disclosure originally reported in the literature might result from the choice of statistical analysis, and that disfluency does not reliably or consistently affect self-disclosure. Thus, disfluency may not be relied on for interface designers trying to improve security or privacy decision making.

1 Introduction

Humans play a crucial role in security and privacy decisions. Technically secure systems can be vulnerable to attacks when users make suboptimal choices, such as choosing or re-using easy-to-guess passwords or falling prey to phishing or spear-phishing attacks [19, 7]. Similarly, privacy decisions are often left to the sole responsibility of the user, who is—at times quite unrealistically—expected to make optimal choices regarding what information to reveal or share based on her preferences and the trade-offs involved in information disclosure [1]. Realizing the complexity of such a task, security and privacy researchers have been trying to find ways to assist individuals' decision-making by designing interventions or interfaces informed by fundamental ideas from fields such as psychology and behavioral decision research.

One such idea is that of *cognitive fluency*. Disfluent conditions are those that induce “metacognitive difficulty” [4]. For instance, disfluency is induced during text processing tasks by using a hard-to-read font or a second language. Based on the literature on cognitive fluency, people perceive a task as more difficult when it is disfluent than when it is fluent, and process it analytically, as opposed to relying on quick judgments and heuristics [4, 17]. It has been argued that people make better decisions and learn better in disfluent conditions than in fluent ones [9, 16].

A potential implication of these findings could be that users could be assisted in making better security and privacy decisions simply by changing the font used to provide information relevant to that decision. The argument is that users would slow down to process the font, and in doing so, would slow down to truly contemplate their decision. This would lead to fewer errors in judgment when making a security decision (such as whether to open an attachment) or a privacy choice (such as whether to reveal sensitive information).

Describe the last time you were sexually aroused.

Describe the last time you were sexually aroused.

Figure 1: Examples of question in fluent font (top) and disfluent font (bottom)

2 Problem Being Solved

In a series of experiments, we investigated the role of disfluency in the decision to disclose sensitive information. While this is a privacy decision, it should be noted that the decision to reveal privacy-sensitive information could also lead to security issues. For example, when a user reveals information to challenge questions in authentication solutions, or when a businessperson reveals corporate information that could lead to a data breach.

Our intention was to replicate, and then build off of a seminal 2009 paper by Alter & Oppenheimer reporting that fluent conditions encouraged disclosure, while disfluent conditions discouraged it [3]. In particular, the manuscript found that disfluency reduces disclosure when it is measured using the Social Desirability Scale (SDS) [8]: participants admitted to fewer undesirable behaviors when the questions were harder to read. Our attempts to replicate this work showed that disfluency did not consistently impact self-disclosure. We were not able to find an explanatory variable or a method of testing that could show that disfluency reduces disclosure when compared to a fluent condition. Instead, we found that one unusual choice of statistical tests may explain the original effect reported by the authors. We conclude that disfluency may not be relied upon to influence users to disclose less information, and that this holds true for different measures of disclosure and different formats for the surveys.

3 Background and related work

A large body of research has explored how people make disclosure decisions. One of the rising themes from the literature is that people's personal preferences for self-disclosure are not always consistent, and that various factors affect people's decisions to disclose personal information [14]. People seem to rely on contextual cues, such as a survey's look and feel or implicit social norms, when disclosing intimate details about themselves [11]. A recent study has shown that when people perceive higher control over who can access and use their online personal information, they become more willing to disclose it even if that implies higher objective risks of privacy intrusions or security breaches, as compared to a condition where they perceive less control but are actually exposed to lower objective risks from their dis-

closure [6]. Also, people respond more honestly and with higher rates of disclosure to an online, versus a paper-and-pencil version, of the same questionnaire [21], and are more inclined to divulge information online than when they communicate face-to-face [10, 22].

Following this line of research, a recent manuscript [3] showed how cognitive disfluency (presenting questions in a hard-to-read font) suppressed people's propensity to disclose personal and sensitive information about themselves. When questions were printed in a disfluent font (50% gray, italicized 10-point Times New Roman font), participants exhibited lower rates of self-disclosure, compared to a condition where the questions were printed in a clear font (black, 12-point Times New Roman font). In a first study (Study 1a), 33 undergraduate participants showed higher percentages of socially desirable responses (indicating lower self-disclosure) when the 33 items in the Crowne-Marlow Social Desirability Scale (SDS) [8] were printed in disfluent vs. clear fonts. The second study (Study 1b) replicated these results using the 10-item version of the SDS. The third study (Study 2) showed that disfluency increased participants' tendency to generate thoughts associated with risks. The fourth study (Study 3) showed that the effect of disfluency on self-disclosure was mediated by negative emotions. The last study (Study 4) was a field experiment that showed that when an online disclosure web site (Group Hug) changed its design, making it easier to read, self-disclosure went up. Our efforts to replicate focused on Study 1a and Study 1b, which used the SDS measure.

4 Approach

The purpose of all our studies was to test the following hypothesis:

H: Disfluency reduces self-disclosure

We report several failed attempts to replicate one of the findings reported in [3] - namely, the results associated with Study 1a and Study 1b in the original paper. We focused on these two studies because they a) demonstrated the (alleged) effect of disfluency and disclosure and b) did so without any additional mediators (as Studies 2 and 3 did) and c) were replicable (Study 4 could not have been, for all practical purposes replicated reliably). Our original objective was actually to validate, using the disfluency manipulation, several self-disclosure

measures, which we planned to use for a different research project. We predicted that, as disfluency affected SDS scores, it should also affect other measures of self-disclosure.

5 Method

To test hypothesis **H**, we ran survey studies on human subjects. Our study was approved by Carnegie Mellon's Institutional Review Board. We designed our study so that the only differences between the two conditions was whether a disfluent or fluent font was used. Participants were assigned randomly to each condition. There were no significant differences between conditions in each study for age or gender. All information about the study, including consent forms, information about the researchers, and instructions, were identical in the different conditions for each study.

After the initial failure to replicate the results presented in one of the studies in [3], we attempted several variations of the study. The original authors of the manuscript were extremely prompt and helpful, and provided us with the original experimental material employed in their experiments. We ran our surveys online, in person using tablets, and in person using pen and paper. We describe the materials used and the subjects tested for each survey below.

The first study was completed online using all four measures of self-disclosure, as described below. To explore whether the fact that the study was conducted online could explain the null results of Study 1, we conducted several follow-up studies (Study 2, 3, and 4) focusing specifically on the SDS measure used in [3]. The second study was done in-person using tablets. The third and fourth studies were done using pen-and-paper surveys. Specifically, the third study was done using material that we independently developed, while the fourth study used the original research materials used by Alter and Oppenheimer. The results of our studies are shown in Table 5.2. Finally, a follow-up study was also online and was a replication of one measure in the first study—namely, the SDS.

The survey questions are included in Appendix A, and samples of the pen and paper versions are provided in Appendix B. All surveys were anonymous; users were not asked to provide any identifying information.

Across the studies, participants were asked to complete several measures of self-disclosure. All of these measures were either shown in a clear, regular, font (Times new-roman, 12 points, black) or a disfluent font (Times new-roman, 10 points, grey). Examples of each are shown in Figure 1).

5.1 Manipulation Check

Based on previous findings [3], we predicted that the disfluent font would make the survey harder to read and thus reduce participants' rates of self-disclosure. In the first online study, participants in the disfluent condition indeed rated the survey as less easy to read than those in the clear font condition ($M_{\text{disfluent}}=5.38$, $SD=1.78$, vs. $M_{\text{clear}}=6.38$, $SD=.91$, $t(383)=6.96$, $p < .01$ where 1="not at all" and 7="very much" easy to read).

5.2 Measures of Self-Disclosure

We used four measures of self-disclosure in our first online study. We describe these measures below. In the pen-and-paper studies and tablet studies, we focused solely on SDS, in an effort to replicate the original study presented in [3].

5.2.1 Social Desirability Scale

The Social Desirability Scale, developed by Crowne and Marlowe in 1960 [8], includes questions about socially desirable behaviors, as well as undesirable behaviors. Many of the socially desirable behaviors reflect an improbable but idealized behavior, while the undesirable behaviors may be more realistic. The original version of the scale consisted of 33 yes or no questions. This scale has been used both as a measure of how respondents self-report their own behavior (along with their need for approval) and as a measure of their actual behavior [12]. Researchers have developed short forms of the SDS scale and tested their reliability and homogeneity [18, 20, 13]. We also found high reliability for this scale in our study (Cronbach's $\alpha = .85$ in Study 1). Participants' score on the SDS is calculated by summing up the number of undesirable responses (i.e., responding 'yes' to a socially undesirable behavior, such as "I like to gossip" or responding 'no' to a socially desirable behavior, such as "I'm always a good listener"). This sum is then divided by the number of items, resulting in a proportion of undesirable responses. This score could be considered a proxy of self-disclosure, as the more people are willing to admit to undesirable responses, the higher their self-disclosure is. Thus, a higher score on the SDS means higher self-disclosure. We believe SDS may be an imperfect measure of self-disclosure, as it may conflate actual conformance to social norms with disclosure. Therefore, we specifically used three additional measures to examine disfluency and self-disclosure.

5.2.2 Unethical Behaviors

We used 14 questions about unethical behaviors adapted from [11]. The questions had been tested for privacy sen-

Study	Measure	Disfluent		Control		Independent t-test		Paired t-test	
		Mean	SD	Mean	SD	t(df)	p	t(df)	p
A&O 1a	SDS 33-items ^a	40.92%	13.73	34.94%	12.72	1.26 (31)	0.22	2.26 (32)	0.03
A&O 1b	SDS 10-items ^a	42.11%	16.53	34.71%	18.75	1.26 (34)	0.22	2.6 (9)	0.03
1: Online	SDS 33-items ^b	45.97%	19.62	45.75%	17.68	-.12 (388)	0.91	.24 (32)	0.81
	Unethical behaviors ^c	29.82%	15.86	31.89%	16.47	1.27 (388)	0.21	1.88 (14)	0.08
	\$15 gift card ^d	3.23	5.73	3.37	5.65	.25 (388)	0.81	-	-
	\$10 gift card ^d	2.27	3.65	2.19	4	-.22 (388)	0.83	-	-
	Sensitive questions (1) ^e	1.22	0.45	1.30	0.54	2.09 (388)	0.04	1.6 (4)	0.19
	Sensitive questions (2) ^{e,f}	1.36	0.47	1.17	0.46	-2.16 (112)	0.03	-2.3 (4)	0.08
2: Tablets	SDS 33-items ^b	43.62%	13.60	41.35%	13.5	-.81 (91)	0.42	2.17 (32)	0.04
3: P&P1	SDS 33-items ^b	46.46%	16.70	48.28%	16.5	.49 (78)	0.63	-.89 (32)	0.38
4: P&P2	SDS 33-items ^b	43.81%	16.02	42.61%	13.3	.31 (57)	0.76	.5 (32)	0.62
	SDS 10-items ^b	25.16%	12.10	25.67%	11.7	-.17 (59)	0.87	.71 (9)	0.49

^a Results from Alter & Oppenheimer (2009) [3]

^b Percentage of responses agreeing (or disagreeing) with socially desirable (or undesirable) behaviors. Higher scores indicate lower self-disclosure.

^c Percentage of responses admitting to unethical behaviors. Higher scores indicate higher self-disclosure.

^d Mean difference (in dollars) participants were willing to pay for an anonymous vs. identified gift card of the same value (\$10 or \$15). A positive difference indicates that the participant was willing to pay more for the anonymous card, and, larger (positive) differences indicate less self-disclosure.

^e Two independent judges rated the depth of disclosure in participants' responses to the open-ended questions on a scale between 0 (non-responses) to 4 (highly revealing responses). The scores presented are averages of the five sensitive questions. Higher scores indicate higher self-disclosure.

^f Results of the follow-up study to Study 1.

Table 1: Comparing disfluent and control conditions for all studies and measures

sitivity. These questions included financial (“Have you had credit card debt above \$100”), illegal (“Have you tried LSD, ecstasy or similar drugs”), and sexual themes (“Have you masturbated at work, school, or in a public restroom”). Participants were given the options to respond with “yes,” “no,” or “Prefer not to answer.” We used affirmative answers (“yes”) as a measure of willingness to disclose and averaged the rate of self-disclosure across the 14 questions. Notice that these answers rely on self-reporting, and thus do not attempt to measure participants’ actual level of unethical behaviors, but only their willingness to disclose such behaviors. Assuming that a participants who either committed or did not commit a certain unethical behavior is equally likely to be randomly assigned to the control or disfluent condition, different scores between the conditions would suggest different levels of self-disclosure.

5.2.3 Sensitive Questions

The next dependent variable we considered consisted of five open-ended questions that pertain to sensitive issues [15]. For instance, this set included questions such as: “What is your most common sexual fantasy?” Participants were asked to respond using several sentences. Following Moon [15], two independent judges rated the depth of disclosure in participants’ responses to the open-ended questions on a scale from 0 (non-responses) to 4 (highly revealing responses) for amount and detail of disclosure. Inter-rater reliability ranged from .86 to .91 for all questions. Thus, the scores for all five questions were averaged into one score for disclosure depth on the sensitive questions in overall.

5.2.4 Gift Cards

Previous literature has used willingness to keep or exchange gift cards of different value and with different privacy protection features as a measure of informational privacy concerns [2]. Using a similar approach, we asked four pairs of questions (8 total) about willingness to buy gift cards that were either described as anonymous (not requiring any personal information to use) or identified (requiring email, name and address to use). We calculated, for each participant, the difference in their willingness to pay for an identified vs. anonymous gift card of a certain value (\$10 and \$15). A positive difference would indicate that the participant was willing to pay more for the anonymous card, and larger (positive) differences would suggest a higher preference toward the anonymous card (and, thus, less disclosure).

6 Analysis

As most research that compares the effectiveness of a treatment on a reliable scale between a treatment and control conditions does, we calculated the percentage of desirable responses per participant, and ran an independent samples t-test to check for significant differences between the averages of desirable responses across the two conditions (these are reported in 5.2). Such a test is typically used to determine whether the difference in the mean between two unrelated groups is significant (between-subjects design), under the assumption that the samples are drawn from normally distributed populations with equal variances (but adjustments for unequal variances are trivial with modern statistical software). The significance of the results is highly reliant on the sample size.

7 Results

In this section we describe the data and the results from the four studies we completed. For each study, we describe the participant selection criteria, the number of participants, and the gender balance of participants. We also describe the questionnaires and the results for each study.

7.1 Study 1: Online

Study 1 consisted in an online survey that included 390 participants from Amazon Mechanical Turk ¹ (64.6% were female; mean age was 35.34, SD=13). Berinsky found that studies using this platform often had samples more representative of the United States population than most in-person studies. Also, they were able to replicate several studies using this platform [5] Therefore, we are reassured that our results can be generalized to the US population and their online disclosure decisions. We included several attention check questions (e.g. “Have you ever had a fatal heart attack while watching television?”), which we used to eliminate participants who gave the wrong answer.

The four measures (SDS, unethical behaviors, sensitive questions, and willingness to buy gift cards) were presented to participants in random order, and so were the questions composing each measure. The surveys were anonymous, and participants were paid 80 cents. About half of the participants received all of these questions in a clear font, while the others received them in a disfluent font. For each of these measures, we computed an overall average score for each of the conditions and compared these using an independent samples t-test. As

¹<http://aws.amazon.com/mturk/>

can be seen in 5.2, only the open-ended sensitive questions showed a statistically significant difference in the expected direction, while for all other measures no significant differences were found between the disfluent and clear font conditions.

In a follow-up study, conducted six months later, we tried to replicate the effect of disfluency on these sensitive questions. In this replication effort, we recruited 114 additional participants from Amazon Mechanical Turk and used only the sensitive question measure (as opposed to all four measures in Study 1). We found statistically significant results in the *opposite* direction: disfluency increased disclosure, as can be seen in 5.2.

7.2 Study 2: Tablets

We hypothesized that the null results from Study 1 were due to the fact that the study was conducted online. Therefore, we decided to run a study in person using tablets. Furthermore, in an effort to replicate the original work, we used a similar population: undergraduate students on a college campus. We asked 93 students (58.1% females, average age 19.5, $SD=1.7$) to complete the 33-items SDS using 8-inch tablets. The tablets were configured in a way that a participant could not change the zoom of the screen (which was, theoretically, possible in the previous study, and may have potentially tampered with our manipulation of disfluency). In these studies, run at the university campus center, a research assistant asked students to take the survey immediately using the provided tablet. Participants received a candy bar for their time. We found, as shown as the second study in 5.2, no significant difference between the disfluent and control conditions.

7.3 Study 3: Pen and Paper 1

Study 2 was done electronically, not using pen and paper, as in the original study. We thus decided to run our third study using paper-and-pencil questionnaires (as in [3]). We asked 80 undergraduate students (58.8% female, average age 19.6, $SD = 1.7$) to complete the 33-items SDS questionnaire either in the clear font or disfluent font conditions. Once again, we ran the study at the university campus center. Students were recruited by a research assistant, took the paper survey immediately, and were given a candy bar for their participation. Again, we found no statistically significant difference between the conditions (see 5.2).

7.4 Study 4: Pen and Paper 2

We then suspected that perhaps we were not manipulating disfluency as in the original study [3]. We contacted

the authors, requesting the original research materials, which they promptly supplied. Their questionnaire did not use a table format as our did, and was more compact. Otherwise, the font, grey, and italics were the same. The disfluent versions of the 33-item materials for Study 3 and 4 are shown in Appendix A to enable comparison.

We ran our fourth study on two additional samples: one using the long, 33-item version ($N=59$, 69.5% females, average age 19.3, $SD=1.8$), and one using the short, 10-item version ($N=61$, 57.4% female, average age 19.3, $SD=2.7$). Unfortunately, as shown in 5.2, we again found no statistically significant difference between the disfluent and clear conditions in either of these samples.

7.5 Results Summary

We ran four different studies, one with 390 on-line participants, and three studies with between 59 and 93 participants. We compare these to the original studies which used 31 and 34 participants. Our smaller studies of undergraduate students more closely resembled the original study's participants, which were also undergraduates. Across all our studies we found disfluency to not have a statistically significant effect in eight of the ten instances. In only one case (the sensitive questions in Study 1), disfluency had a significant effect in the expected direction (participants in the disfluent condition disclosed less). However, in another case (the follow-up experiment to Study 1 using sensitive questions only), the effect was statistically significant, but in the opposite direction (participants in the disfluent condition disclosed more). Thus, in summary, it can be seen that the effect of disfluency on self-disclosure is not reliable or consistent. We find no evidence that disfluency impacts self-disclosure.

8 Re-Analysis and Discussion

Collectively, these failed attempts at replicating the effect of disfluency on self-disclosure puzzled us. We therefore decided to share our results with Alter and Oppenheimer and asked for their expert opinion on the matter. In response, they openly and promptly shared their data from the original tests, and commented that they had used a different statistical test to analyze their data.

A closer examination of the Alter and Oppenheimer's original paper suggested that they had used a did not use the independent means t-test, we we did (described above). Namely, the number for the degrees of freedom reported in their studies does not correspond to an independent t-test analysis (32 in Study 1a and 9 in Study 1b, in which N s were 33 and 36 respectively). As confirmed through personal communications with the original authors, the tests used in the original paper were t-tests for

paired (and not independent) samples. Paired t-tests are typically used for within-subjects designs. For example, they can be used to show before and after results for a subject taking a treatment, or two different treatments on the same subject. In the case of Alter and Oppenheimer's studies, however, the researchers calculated a mean score for each item of the SDS in each of the conditions, resulting in a dataset containing 33 (or 10, in the short version) paired mean scores. Each of these mean item scores in each condition were treated as a single unit of observation, and subjected to a paired-samples t-test to examine the differences in the items' means between the conditions.

When treating individual questions as units of observation, one does not take into account the number of subjects that answered each question to determine significance. More importantly, the independence assumption on which the test is based is harder to justify, since the various questions are answered by the same participant and are, therefore, likely to be correlated. Additionally, the paired samples approach might provide an underestimation of the variance in the population, as within-subjects variance tends to be smaller than between-subjects variance.

After receiving from the authors of the original paper the original raw data for Studies 1a and 1b, we re-analyzed it using independent samples t-tests instead of the paired t-test that was reported in the original paper. We discovered that the results were not statically significant (see 5.2). Finally, we re-analyzed our own data, in all of our studies and for all of the measures that used multiple questions (i.e., excluding the gift-cards question), using a paired t-test approach. In our studies, a paired t-test did not yield significant results, except for one case (Study 2; see 5.2). A meta-analysis on the results of all independent t-tests showed the average test value to be of $Z = 1.02$, and not statistically significant ($p = .31$). A similar meta-analysis on the results of all of the paired t-tested showed a similar, non-significant, result ($Z = .93$, $p = .36$). To summarize, these re-analyses showed that the null hypothesis which states that disfluency does not affect self-disclosure could not be rejected. In other words, no consistent or reliable effect could be found for disfluency on self-disclosure, using either the independent samples t-test or the paired t-test.

9 Conclusions

Two conclusions may be drawn from the results of our studies. First, and most importantly, we did not find evidence that disfluency consistently impacts self-disclosure, or at least not in a manner that our studies, based on [3], could precisely estimate. The inconsistency of this effect was observed on a number of measures (in-

cluding SDS, unethical behaviors and sensitive personal questions), in a variety of administration forms (online surveys, in person, and pen-and-paper) and in different samples (students and non-students). Thus, if an interface or system designer wished to use fluent or disfluent fonts in an attempt to encourage higher or lower disclosure by users, she will find that disfluency is not a reliable solution. Security and privacy designers should be aware that disfluency might slow-down their users, but would not cause them to reveal less.

Second, the original reported effect of disfluency on self-disclosure might have been due to the use of a very specific statistical test, whose appropriateness in this case could be questionable. Specifically, the original analysis was conducted as an 'item-level' analysis instead of the more widely used 'subject-level' analysis. We are not fully convinced that a paired samples t-test is appropriate for this type of analysis. Our failures in replicating the disfluency results could therefore serve to raise a discussion in the community as to what is the more valid statistical approach to be used in similar instances.

10 Acknowledgments

We thank Adam Alter and Daniel Oppenheimer for sharing their data and analyses with us, and for their prompt responses and cooperation with our queries. We also thank Howard Seltman for his expert statistical advice. This work was funded by the National Science Foundation under Grants CNS-1012763 (Nudging Users Towards Privacy) and a Google Focused Research Award on Privacy Nudges.

References

- [1] ACQUISTI, A. Privacy in Electronic Commerce and the Economics of Immediate Gratification. In *Proceedings of the 5th ACM conference on Electronic commerce* (2004), ACM, pp. 21–29.
- [2] ACQUISTI, A., JOHN, L., AND LOEWENSTEIN, G. What is privacy worth. In *Workshop on Information Systems and Economics (WISE)* (2009).
- [3] ALTER, A. L., AND OPPENHEIMER, D. M. Suppressing secrecy through metacognitive ease cognitive fluency encourages self-disclosure. *Psychological science* 20, 11 (2009), 1414–1420.
- [4] ALTER, A. L., OPPENHEIMER, D. M., EPLEY, N., AND EYRE, R. N. Overcoming intuition: Metacognitive difficulty activates analytic reasoning. *Journal of Experimental Psychology-General* 136, 4 (2007), 569–576.
- [5] BERINSKY, A. J., HUBER, G. A., AND LENZ, G. S. Using mechanical turk as a subject recruitment tool for experimental research. *Submitted for review* (2011).
- [6] BRANDIMARTE, L., ACQUISTI, A., AND LOEWENSTEIN, G. Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science* (2012).

- [7] CRANOR, L. F. A framework for reasoning about the human in the loop. In *Proceedings of the 1st Conference on Usability, Psychology, and Security* (2008), vol. 1, USENIX Association, pp. 1–1.
- [8] CROWNE, D. P., AND MARLOWE, D. A new scale of social desirability independent of psychopathology. *Journal of consulting psychology* 24, 4 (1960), 349.
- [9] DIEMAND-YAUMAN, C., OPPENHEIMER, D. M., AND VAUGHAN, E. B. Fortune favors the (): Effects of disfluency on educational outcomes. *Cognition* 118, 1 (2011), 111–115.
- [10] HARPER, V., AND HARPER, E. J. Understanding student self-disclosure typology through blogging. *The Qualitative Report* 11, 2 (2006), 251–261.
- [11] JOHN, L. K., ACQUISTI, A., AND LOEWENSTEIN, G. Strangers on a plane: Context-dependent willingness to divulge sensitive information. *Journal of Consumer Research* 37, 5 (2011), pp. 858–873.
- [12] JOHNSON, T. P., AND FENDRICH, M. A validation of the crowne-marlowe social desirability scale. In *57th Annual Meeting of the American Association for Public Opinion Research* (2002).
- [13] LEITE, W. L., AND BERETVAS, S. N. Validation of scores on the marlowe-crowne social desirability scale and the balanced inventory of desirable responding. *Educational and Psychological Measurement* 65, 1 (2005), 140–154.
- [14] MOON, Y. Intimate exchanges: Using computers to elicit self-disclosure from consumers. *Journal of Consumer Research* 26, 4 (2000), 323–339.
- [15] MOON, Y. Intimate exchanges: Using computers to elicit self-disclosure from consumers. *Journal of Consumer Research* 26, 4 (2000), pp. 323–339.
- [16] MOREWEDGE, C. K., AND KAHNEMAN, D. Associative processes in intuitive judgment. *Trends in Cognitive Sciences* 14, 10 (2010), 435 – 440.
- [17] OPPENHEIMER, D. M. The secret life of fluency. *Trends in Cognitive Sciences* 12, 6 (2008), 237 – 241.
- [18] REYNOLDS, W. M. Development of reliable and valid short forms of the marlowe-crowne social desirability scale. *Journal of clinical psychology* 38, 1 (1982), 119–125.
- [19] SMITH, S. Humans in the loop: human-computer interaction and security. *Security Privacy, IEEE* 1, 3 (2003), 75–79.
- [20] STRAHAN, R., AND GERBASI, K. C. Short, homogeneous versions of the marlowe-crowne social desirability scale. *Journal of clinical psychology* (1972).
- [21] TOURANGEAU, R. Survey research and societal change. *Annu. Rev. Psychol.* 55 (2004), 775–801.
- [22] WHITTY, M. T., AND JOINSON, A. N. Truth, lies and trust on the internet. (2009). *Truth, lies and trust on the Internet*. vi, 176 pp. New York, NY, US: Routledge/Taylor & Francis Group (2009).

A Pen and Paper Screenshots

The following screenshots show the disfluent (Times new-roman, 10 points, grey) version of the two pen and paper studies: Study 3 (top) and Study 4 (bottom). These screenshots demonstrate that the disfluent versions were light and difficult to read.

Following are 33 statements to which some people agree and some do not.

Please mark whether you agree or do not agree with each of the following statements.

	<i>Agree</i>	<i>Disagree</i>
<i>1. Before voting, I thoroughly investigate the qualifications of all the candidates</i>		
<i>2. I never hesitate to go out of my way to help someone in trouble</i>		
<i>3. It is sometimes hard for me to go on with my work if I am not encouraged</i>		
<i>4. I have never intensely disliked anyone</i>		
<i>5. On occasion I have had doubts about my ability to succeed in life</i>		
<i>6. I sometimes feel resentful when I don't get my way</i>		
<i>7. I am always careful about my manner of dress</i>		
<i>8. My table manners at home are as good as when I eat out in a restaurant</i>		
<i>9. If I could get into a movie without paying and be sure I was not seen I would probably do it</i>		
<i>10. Of a few occasions, I have given up doing something because I thought too little of my ability</i>		
<i>11. I like to gossip at times</i>		
<i>12. There have been times when I felt like rebelling against people in authority even though I knew they were right</i>		
<i>13. No matter who I'm talking to, I'm always a good listener</i>		
<i>14. I can remember "playing sick" to get out of something</i>		
<i>15. There have been occasions when I took advantage of someone</i>		
<i>16. I'm always willing to admit it when I make a mistake</i>		
<i>17. I always try to practice what I preach</i>		
<i>18. I don't find it particularly difficult to get along with loud mouthed, obnoxious people</i>		
<i>19. I sometimes try to get even rather than forgive and forget</i>		
<i>20. When I don't know something I don't at all mind admitting it</i>		
<i>21. I am always courteous, even to people who are disagreeable</i>		
<i>22. At times I have really insisted on having things my own way</i>		
<i>23. There have been occasions when i felt like smashing things</i>		
<i>24. I would never think of letting someone else be punished for my wrongdoings</i>		

	<i>Agree</i>	<i>Disagree</i>
25. <i>I never resent being asked to return a favor</i>		
26. <i>I have never been irked when people expressed ideas very different from my own</i>		
27. <i>I never make a long trip without checking the safety of my car</i>		
28. <i>There have been times when I was quite jealous of the good fortunes of others</i>		
29. <i>I have almost never felt the urge to tell someone off</i>		
30. <i>I am sometimes irritated by people who ask favors of me</i>		
31. <i>I have never felt that I was punished without cause</i>		
32. <i>I sometimes think when people have a misfortune they only got what they deserved</i>		
33. <i>I have never deliberately said something that hurt someone's feelings</i>		

Please state your gender: Male Female

What is your age? _____

Are you a(n) undergraduate or graduate student?

Thank you very much!

Listed below are a number of statements concerning personal attitudes and traits. Read each item and decide whether the statement is True or False as it pertains to you personally.

1. Before voting I thoroughly investigate the qualifications of all the candidates.
2. I never hesitate to go out of my way to help someone in trouble.
3. It is sometimes hard for me to go on with my work, if I am not encouraged.
4. I have never intensely disliked anyone.
5. On occasion I have had doubts about my ability to succeed in life.
6. I sometimes feel resentful when I don't get my way.
7. I am always careful about my manner of dress.
8. My table manners at home are as good as when I eat out in a restaurant.
9. If I could get into a movie without paying and be sure I was not seen, I would probably do it.
10. On a few occasions, I have given up doing something because I thought too little of my ability.
11. I like to gossip at times.
12. There have been times when I felt like rebelling against people in authority even though I knew they were right.
13. No matter who I'm talking to, I'm always a good listener.
14. I can remember "playing sick" to get out of something.
15. There have been occasions when I took advantage of someone.
16. I'm always willing to admit it when I make a mistake.
17. I always try to practice what I preach.
18. I don't find it particularly difficult to get along with loud-mouthed, obnoxious people.
19. I sometimes try to get even rather than forgive and forget.
20. When I don't know something I don't at all mind admitting it.
21. I am always courteous, even to people who are disagreeable.
22. At times I have really insisted on having things my own way.
23. There have been occasions when I felt like smashing things.
24. I would never think of letting someone else be punished for my wrongdoings.
25. I never resent being asked to return a favor.
26. I have never been irked when people expressed ideas very different from my own.
27. I never make a long trip without checking the safety of my car.
28. There have been times when I was quite jealous of the good fortune of others.
29. I have almost never felt the urge to tell someone off.
30. I am sometimes irritated by people who ask favors of me.
31. I have never felt that I was punished without cause.
32. I sometimes think when people have a misfortune they only got what they deserved.
33. I have never deliberately said something that hurt someone's feelings.

Please fill out the following items:

Age: _____ Sex: _____ Ethnicity/Race: _____

How frustrated do you feel at the moment (rated from 1 to 10, where 1 is not at all frustrated, and 10 is very frustrated)? _____