The Effects of Introspection on Creating Privacy Policy

Stephanie Trudeau Sara Sinclair Sean W. Smith
Dept. of Computer Science
Dartmouth College
Hanover, NH USA
{trudeau,sinclair,sws}@cs.dartmouth.edu

ABSTRACT

Prior work in psychology shows that introspection inhibits intuition: asking human users to analyze judgements they make can cause them to be quantitatively worse at making those judgments. In this paper, we explore whether this seemingly contradictory phenomenon also occurs when humans craft privacy policies for a Facebook-like social network. Our study presents empirical evidence that suggests the act of introspecting upon one's personal security policy actually makes one worse at making policy decisions; if one aims to reduce privacy spills, the data indicate that educating users before letting them set their privacy policies may actually *increase* the exposure of private information.

Categories and Subject Descriptors

K.6.5 [**Management of Computing and Information Systems**]: Security and protection; H.1.2 [**Information Systems Applications**]: User/Machine Systems

General Terms

Human Factors, Experimentation

1. INTRODUCTION

Computer security classically depends on a *policy*: a specification of who should be allowed to do what, to whom, and when. A computing system typically provides to the user a set of "knobs" the user can adjust in order to define the policy the system enforces. The intention is that the user can use these inputs to ensure the system's de facto policy matches the user's mental model of what the policy should be.

Let P_k denote the policy enforced by the system with knob setting k, and let P_U denote the policy the user intends.[1]

[1]This is when the user already knows what she would like the policy to be; in unfamiliar domains or specialized systems, this may not be the case.

WPES'09, November 9, 2009, Chicago, Illinois, USA.

Whether this approach actually works depends on two implicit assumptions:

- There actually exists a knob setting k such that $P_k \equiv P_U$
- A user who wishes to ensure that the system follows some policy P_U can actually construct such a knob setting k.

Prior work (such as [1, 2, 8]) explored the first point, i.e., whether the space of knob-settings of particular systems was sufficiently nuanced to capture the well-designed policies users often have in their heads.

In this paper, we explore the second point. Prior work in psychology (e.g., [14]) shows that introspection inhibits intuition: asking human users to analyze judgements they make can cause them to be quantitatively worse at making those judgments. Does this phenomenon happen with security? Does leading users to think about how to craft security policies qualify as such introspection—and result in a P_k that does not match the user's initial intended P_U?

To examine this question, we developed a mock-up of a Facebook-like social networking site, and had users make a series of access control decisions about sharing various categories of personal data. Our results show that users who were first asked to think about issues regarding privacy and potential privacy policies subsequently made access control decisions that were significantly different from the control group. More surprisingly, the results show that introspection doesn't just change users' treatment of private information: it made it worse! Subjects who thought about privacy policies were subsequently far more willing to share private data.

This Paper.

Section 2 presents the psychology study that motivated this work. Section 3 presents the social networking application we developed to examine this question. Section 4 presents our experimental methodology. Section 5 presents our results; Section 6 analyzes them. Section 7 concludes and considers possible areas of future work.

2. INTROSPECTION INHIBITS INTUITION

Background.

As a jumping-off point for this work, we consider a 1991 psychology experiment by Wilson and Schooler [14]. Taste

Figure 1: The GUI for a sample "InnerCircle" profile. Two features not present in Facebook: *A* shows a note the user adds when he or she "friends" this party; *B* shows information InnerCircle calculates to convey the user's current connectedness to this party.

testing experts ranked 44 commercially-offered jams in order, from best to the worst. In the experiment, college students were asked to rank 5 of these jams (those ranking 1st, 11th, 24th, 32nd, and 44th, according to the experts) in the order they thought tasted the best. The students were divided into two groups. In the control group, students were asked just to rank the jams; in the experimental group the students were asked both to rank the jams and to provide a written explanation of why they ranked them in the order they did. The rankings produced by the control group were were likely to agree with the experts' ranking on both the best and the worst jams, and showed an overall correlation coefficient of 0.55. In contrast, the experimental group's rankings were much less similar to the experts' ranking, demonstrating a correlation of only 0.11.

Gladwell summarizes this experiment in his pop-science *Blink* [6]: "by making people think about jam, Wilson and Schooler turned them into jam idiots." Both the study authors and Gladwell offer various theories to explain the observed phenomenon; one such argument goes as follows. A non-expert is unused to making judgements based on identifiable nuanced qualities like texture, stickiness, and sweetness, whereas an expert has learned to quantify these features in evaluating a jam. The control group of non-experts produced judgements similar to the experts', but by priming the experimental group with an analysis of the *specific characteristics* that determine jam quality, the researchers somehow confused them or inhibited their capacity to make judgements that correspond with expert opinion.

Idea.

Do these results from the psychological jam study also apply in the world of security and privacy policies? Does the quality of a user's performance in decision-making *decrease* when the user attempts to apply conscious analysis to their policy authoring?

In lamenting the challenges to designing a user interface for privacy settings, Cranor notes that "privacy policies are complex, user privacy preferences are often complex and nuanced, users tend to have little experience articulating their privacy preferences, and users are generally unfamiliar with much of the terminology used by privacy experts" [1]. Current wisdom dictates that introspective analysis *improves* the quality of security and privacy policies. Establishing the opposite would not only add a new challenge to the design of relevant interfaces, but indicate that many current methods for designing policy-authoring interfaces may be ineffective (or even counterproductive) in situations involving non-expert users. This result would thus be an appeal for a fundamental reformation of current security system design methodology.

3. OUR APPLICATION DOMAIN

Our first step in constructing an appropriate experiment was to pick a specific application domain in which to examine this problem.

File systems permissions are one domain where policy as implemented often fails to match policy as intended. In the 2001 "Memogate" scandal, a Republican staffer on the United States Senate Judiciary Committee Committee discovered that he had the ability to read confidential memos and many other documents belonging to the Democrats serving on the Committee, due to a misconfigured file server [5]. Good and Krekelberg showed in 2003 that most KaZaA users could not accurately determine which files they were sharing and which they were not. At times, the users would falsely assume that "they were not sharing any files when in fact they were sharing all files on their hard drive" [7]. Maxion and Reeder noted that "Microsoft publishes a list of 'best practices' for NTFS security that advises users not to use several of the provided features of the NTFS permissions model" because they "'could cause unexpected access problems or reduce security"' [10]. Reeder went on to do user studies having test subjects try to configure permissions in a distributed file system to achieve certain high-level policy goals [11].

Making trust decisions about online content is another area—either the rich space of phishing vulnerability in general (e.g. [3]) or the problem of trying to decide whether to trust email from an unknown sender (e.g. [9]). Other research (e.g. [1, 2, 12])) explores the challenge of expressing and making decisions about privacy policies of websites.

An interesting and scientifically meaningful user study involving file system permissions would require a non-trivial amount of user training. Given the timeframe for completing this experiment, the richness of prior work on evaluating website policies, and the burgeoning importance of social networking in modern culture, we decided to instead study privacy policies in the context of a social network. Specifically, we decided to test what information subjects would choose to share with members of *InnerCircle*, a fictitious social networking site modeled on Facebook.

We chose to built a mock-up of InnerCircle instead of studying Facebook itself for several reasons; first, presenting subjects with the profile of a fictitious Facebook user opened the possibility that subjects would not be able to get a good enough sense of the fictitious relationship with the user to make a meaningful trust judgement. We decided that offering information about the context in which the subject met the user, how the subject interacts with the user, etc. in the site's interface would more accurately simulate the real-world conditions under which subjects might operate. Second, we did not want subjects' existing Facebook habits or customs to unduly influence their behavior in the study. (For example, some Facebook users might restrict the number of Facebook friends they have during a busy semester to reduce associated distraction; such a decision includes factors beyond basic privacy concerns.) Finally, Facebook changes its user interface frequently, and we worried that performing the experiment with a mockup of an obsolete version would confuse subjects.

Figure 1 depicts one of the fictitious profiles used in the study. It contains components similar to a Facebook profile, including a Profile Picture, and the ideas of Networks, Photo Albums, and Groups. One important addition found in InnerCircle profiles is information on how the test subject is to have met the fictitious person in the profile. Area A in Figure 1 shows the question, "How do you know Michael?", with the corresponding answer: "Michael was your boss last year for your summer internship. You generally got along well and had a good professional relationship." Subjects are told that they entered this information when they added this person as an InnerCircle friend. We hypothesized that the privacy settings subjects generated for a given profile would be strongly influenced by this information. (E.g., we expected that a subject might choose to let a roommate see

	Name	Summary of connection to user	common friends	common tags	last time talked	distance (miles)
1	**Wade Spurlock**	lived on freshman floor; haven't talked since	16	18	infinite	0
2	**Chathum Nielsen**	randomly sat at your table in food court	0	0	infinite	0
3	**Arthur Patterson**	dated in 18 months in HS	23	88	9 months	126
4	**Danny Wilson**	uncle	9	0	9 months	587
5	**Jake Mehrens**	danced with once at Thu Night Salsa	1	3	1 month	656
6	**Andrew Van Winkle**	HS friend, family connections	35	5	1 month	3000
7	**Amanda Hartley**	same Greek house, but don't know well	36	0	6 months	0
8	**Phil Sanders**	met at a party; sketchy	26	0	infinite	0
9	**Beth Franz**	friend of older sister	1	3	infinite	2946
10	**Andrew Parrish**	met at party last term; funny	23	0	3 months	0
11	**Samantha Miller**	same camp in HS; used to hang out	14	18	5 months	2364
12	**Michael Holloway**	boss last summer	1	3	3 months	656
13	**Darcy Shapiro**	same top 5 favorite movies	0	0	infinite	0
14	**Megan Lundeby**	best friend since preschool	24	82	0 months	2719
15	**Maddie Petrin**	track teammate first 2 years of college	35	2	0 months	2997
16	**Colleen Kirsten**	both like Queen	0	0	infinite	361
17	**Peggy Clark**	camp director; you worked; family went	44	87	12 months	2688
18	**Cam Schnur**	met in a hostel in Prague during LSA	0	0	12 months	256
19	**Kate Farrington**	friend of roommates	13	2	infinite	0
20	**Sarah Watkins**	hung out at conference	0	0	2 months	126

Table 1: Our study asked users to make decisions about sharing information with these "Inner Circle" profiles.

their InnerCircle photos, but not give their boss that access privilege.)

Another question in this section that we added to InnerCircle was "When did you and [Person X] meet?" This information was included to give the subject a sense of how long they have been friends with the fictitious person, in case that has an impact on how how they choose to share information with that InnerCircle friend.

The InnerCircle GUI also has a box, labeled B in Figure 1, which provides information on how the subject is connected to the person in the profile. Subjects were told that InnerCircle automatically generates this data; it includes how much their friend circles overlap (that is, how many friends they have in common), how many photos both of these people are tagged in together, the last time the subject sent or received a message from this user, and how close geographically the subject is to this user (based on the Current Location listed on their profile).

After the connectedness information, there is an "Interesting Fact" box in each InnerCircle profile. This was included simply to help test subjects with the experimental suspension of disbelief; such small tidbits of information are an inherent part of real-life relationships.

4. METHODOLOGY

Having picked a sample application domain in which users make privacy policy decisions, we proceeded to carry out a user study modeled on the jam study of Wilson and Schooler (Section 2). Figure 2 shows our overall process: subjects make a series of access control decisions about sharing their InnerCircle information (thus giving us a way to measure P_U)—except the experimental group is first asked to think about a large set of potential issues that might be considered in specifying such policies. Will being explicitly confronted with these "knobs" change the resultant P_U?

Subjects.

The subjects who participated in this study consisted of 100 undergraduate students—56 women and 44 men. Students were drawn from various majors and class years. (While it is possible that the limited age range of participants may have influenced their responses, the ease of recruitment from this population makes such selection common for initial user studies.) The test subjects were divided evenly between the control and experimental groups, placing 50 in each group. (We also balanced gender mixes in each group.) Subjects were recruited by an email sent out to the entire Dartmouth campus, advertising a behavioral experiment related to Facebook that needed research participants. Subjects volunteered to participate in the half-hour study in exchange for $15. The only requirement that was demanded of each subject was that he or she currently has a Facebook account. This study was approved by the Committee for the Protection of Human Subjects (CPHS), the Institutional Review Board (IRB) at Dartmouth College.

Questionnaire.

The study was divided into two tasks. The first task required five to ten minutes to complete a questionnaire. The control group used this time to fill out a questionnaire that asked questions related to how they picked their major in college. The content of this questionnaire was unimportant, and was picked simply because each test subject was a

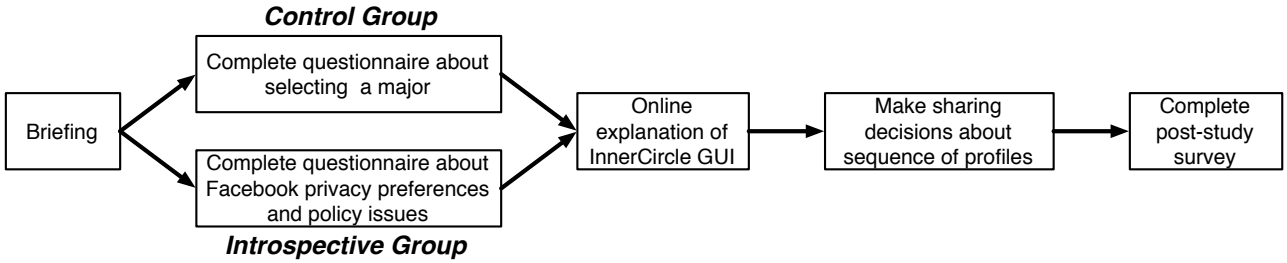

Figure 2: Our experimental process.

college student at Dartmouth College and would find these questions relevant to his or her life. The questionnaire was designed to be twenty-five questions long, consisting of a variety of yes-or-no questions, questions asking the user to pick a value on a scale from one to ten, and questions about when in their college career certain decisions were made. Several of the questions related to one's happiness with their major, or whether happiness ought to even be a factor in one's decision of what to major in during college.

The experimental group had a different questionnaire to fill out during this first task. This questionnaire asked questions about the subject's privacy settings on his actual Facebook account. Most of these questions asked the subject to assume that Facebook had an interface that allowed him to easily configure his security settings in any way that he wants, and then asked how the various settings would be configured. Some additional questions also asked the subject if she had anything on her Facebook account that she would not want various people (an employer, a professor, a parent, etc.) to see. This questionnaire was twenty-seven questions long, and was designed in tandem with the control group's questionnaire to take approximately the same time to complete. (Appendix A contains the introspection questionnaire; further information on the experiment design can be found in the first author's thesis [13].)

Privacy Decisions.

Following the questionnaire, each test subject was then shown twenty InnerCircle profiles in the same pre-generated random order. (Table 1 summarizes these; see [13] for the full images.) Subjects were asked the same seven questions (Figure 3) about each profile. These questions were selected to directly address the seven key areas of Facebook to which users are allowed to configure access by other Facebook users.[2] On Facebook, these settings are made more broadly. For example, in the Privacy Settings area of Facebook, one can set their Basic Info to be viewable by "Everyone," "My Networks and Friends," "Friends of Friends," etc. We designed InnerCircle to be more explicit in these decisions, however, by requiring the user to make these choices with respect to every profile. This task of the experiment took users anywhere from fifteen to twenty-five minutes to complete.

Post-Test Feedback.

[2]The mechanisms for managing Facebook policies have changed at the time of this writing since the experiment was designed.

	Mean	Standard Deviation	Min	Max
Basic Info	88.61	10.07	5.00	95.00
Personal Info	76.60	17.80	25.00	95.00
Email	77.42	18.48	3.00	95.00
Phone	48.92	27.52	0.00	95.00
Address	48.00	29.39	0.00	95.00
Photos	57.45	22.64	0.00	95.00
Videos	54.44	25.45	0.00	95.00

Table 2: This table displays the mean percentage of all 100 users (averaged over all twenty profiles) who granted access to each of the seven types of information in the experiment. The standard deviation of these averages is also shown, as well as the minimum and maximum percentage of yes-answers for each category that were given.

Finally, upon completing the InnerCircle profile evaluations, the user was asked to fill out a post-test feedback survey.

5. RESULTS

Our experimental hypothesis was that introspection inhibits intuition. The act of introspecting upon one's own Facebook security policies (P_k) would cause the experimental group to lose track of their own intuitive security policy (P_U) that they would have liked their InnerCircle privacy policies to express. We expected this confusion to be manifested in the data: perhaps the control group's InnerCircle decisions would be more clustered, but the experimental group's answers would have a higher standard of deviation, varying greatly not only from the control group's answers but from each other's answers as well.

What actually happened was not what we expected—it was more interesting than that.

Openness.

To start with, we can look at the number of yes-answers (which indicate a subject would choose to share a particular bit of information with the fictitious InnerCircle user associated with a particular profile) that were given during the second task of the experiment. The mean number of yes-answers given by the control group was 86.6, while the mean for the introspective group was 98.56. An analysis of

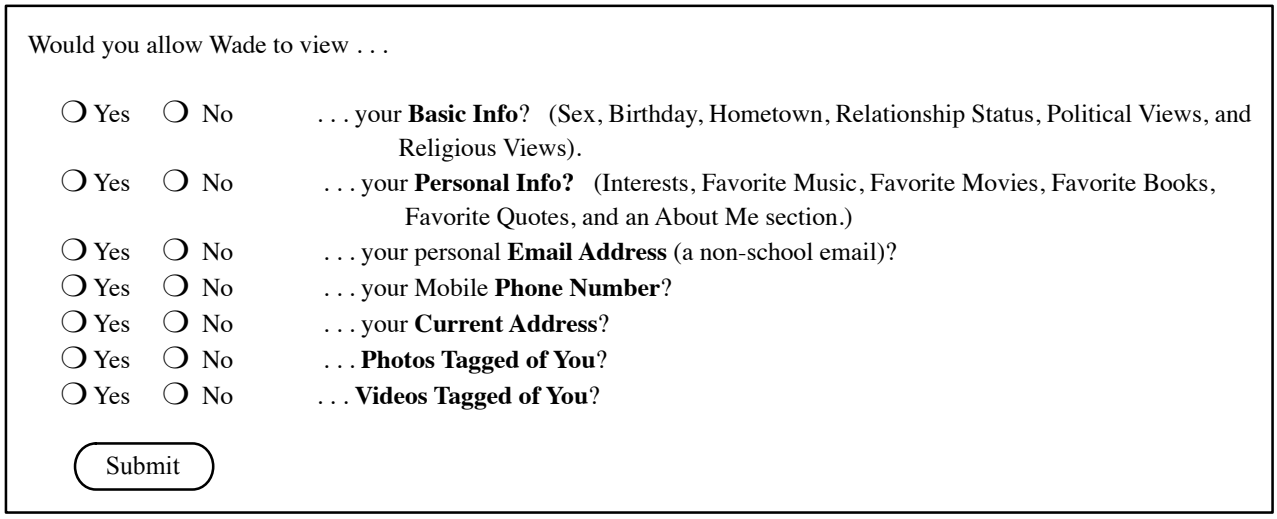

Figure 3: For each profile, users were asked these questions about sharing.

	% of Average (Control)	% of Average (Introspective)	Standard Deviation (Control)	Standard Deviation (Introspective)
Basic Info	88.37	88.85	10.23	10.01
Personal Info	73.88	79.38	20.06	14.86
Email	78.33	76.53	19.11	18.00
Phone	41.86	55.82	27.22	26.29
Address	43.96	52.13	29.96	28.53
Photos	51.00	63.88	25.90	16.75
Videos	47.55	61.33	28.40	20.15

Table 3: The first two columns of this table display the mean percentage of users (averaged over all twenty profiles) who granted access to each of the seven types of information in the experiment. The next two columns show the standard deviation corresponding to these averages.

variance (ANOVA) between the control and experimental groups yields an F statistic of 6.57; this means that the difference in average total yes-answer counts between the two groups is indeed statistically significant.

From this we are able to conclude that *the introspective group was significantly more willing to share their information than the control group.*

Types of Information.

The next data set we examined was that which describes how the two groups of subjects shared each type of information. As Figure 3 showed, the study had seven types of information one could either share or refuse to share. As Table 2 shows, the test subjects were most protective of their phone numbers and mailing addresses, followed by their photos and videos. Table 3 displays the mean number of yes-answers given in each of the seven question-type categories, this time splitting up the means by control group and introspective group. In the categories of Basic Info and Email, the difference between the averages of the control group and the introspective group are clearly negligible. However, the discrepancies are greater for the remaining five categories; ANOVA (Table 4) reveals a significant difference in behavior between the two groups.

The introspective group was not only more willing to share their information than the control group, but *the additional information they chose to share was in fact their most sensitive information.*

Profiles.

Figure 4 charts the difference between the introspective and control groups' willingness to share three sensitive pieces of data (phone number, photos, and videos). This reveals that in some instances the introspective users seemed to make clearly "wrong" decisions: e.g., choosing to strongly trust Profile 5 with sensitive data—even though this profile belonged to "Jake," whom the subjects had met only once. Maddie (Profile 15) and Phil (Profile 8) were given approximately the same degree of access to this sensitive data, yet she was "one of your former teammates on the track team during your first two years of college," and he was "someone you met at a party who seemed sketchy."

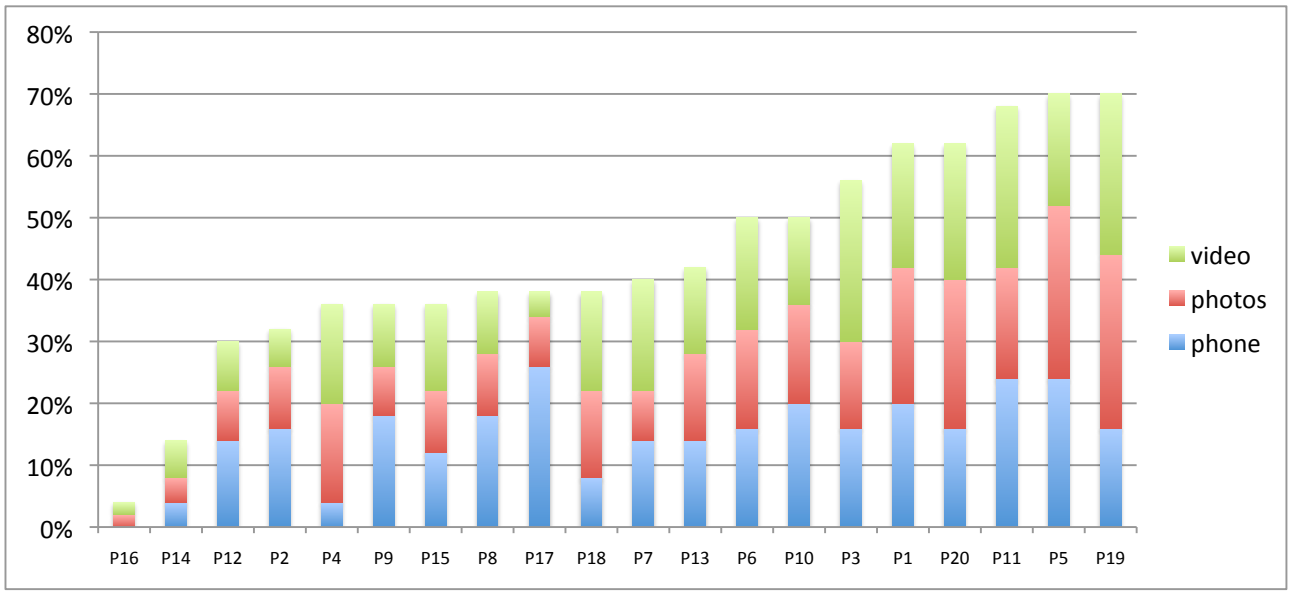

Figure 4: The introspective group was more willing to share their most sensitive data. This chart shows the profiles ordered by this "sensitivity index": the sum of the percentage increase (over the control group) in willingness to share the three most sensitive data types.

	Difference between averages	*Difference between standard deviations*	*F-Test*
Basic Info	0.48	0.22	0.06
Personal Info	5.50	5.20	2.34
Email	1.80	1.11	0.23
Phone	13.96	0.93	6.58
Address	8.17	1.43	1.85
Photos	12.88	9.15	8.51

Table 4: The first column shows the difference between the averages for the control and introspective group that were displayed in Table 3. The second column shows the differences between the standard deviations presented in Table 3. The third column shows the results of an ANOVA F-Test on these averages—shaded rows indicate significant differences.

The final type of analysis we performed on the data was to again compare the control group with the experimental group in terms of the average number of yes-answers given. This time, however, we looked at these averages by profile, rather than by question type. Table 5 shows the mean number of yes-answers given for each of the twenty profiles that were shown to test subjects during the study. Again there were a number of discrepancies between the control and introspective group for averages with respect to particular profiles. ANOVA (Table 6) confirmed that the subject groups differed significantly in their treatment of about half of the twenty profiles.

6. ANALYSIS

As the previous section presents, our results show that by requiring subjects in the experimental group to fill out a questionnaire that asks them them to introspect about social networking privacy policies, those subjects became far more willing to share their most sensitive types of information than their control group counterparts. Since the two groups were chosen from the same population, we can assume that they would by default exhibit the same P_U. However, being exposed to the policy knobs changed P_U—and in a direction that most privacy advocates would consider to be worse. Why?

Although it is not sufficient as an explanation, one phenomenon is clear: the introspective subjects appeared to only have two "buckets" into which they grouped the profiles. Many members of this group articulated this in the post-test feedback, saying that once they had friended someone then they felt weird denying that "friend" access to any of their personal information. (*None* of the control group made this observation.) In contrast, the control users seemed to ex-

		% of Average (Control)	% of Average (Introspective)	Standard Deviation (Control)	Standard Deviation (Introspective)
1	*Wade*	68.80	83.38	28.22	19.64
2	*Chathum*	50.29	61.51	32.17	33.66
3	*Arthur*	81.05	92.71	23.58	16.27
4	*Danny*	62.10	71.73	22.30	23.48
5	*Jake*	58.29	72.01	30.39	28.71
6	*Andrew V.*	86.00	95.43	18.14	10.18
7	*Amanda*	82.29	88.63	21.32	15.70
8	*Phil*	47.71	55.43	29.80	34.23
9	*Beth*	56.57	62.10	28.12	33.45
10	*Andrew P.*	68.86	77.71	25.27	24.69
11	*Samantha*	81.71	94.86	23.98	11.46
12	*Michael*	47.43	53.06	23.35	26.57
13	*Darcy*	40.29	47.14	31.58	37.21
14	*Megan*	94.86	98.00	12.83	7.08
15	*Maddie*	81.43	87.43	20.05	18.62
16	*Colleen*	29.74	24.86	29.42	30.52
17	*Peggy*	54.29	66.29	25.98	25.29
18	*Cam*	50.00	56.27	31.25	32.79
19	*Kate*	61.43	75.71	28.90	25.04
20	*Sarah*	65.71	78.00	28.86	23.66

Table 5: The first two columns of this table display the mean number of users (averaged over the seven different types of questions) who granted access to each of the twenty profiles in the experiment. The next two columns show the standard deviation corresponding to these averages.

hibit a much wider gradation in their profile grouping.

One potential explanation for the results of this study is that after considering the many instances in which she might wish to deny or allow access, the subject is unable to discern her reasons for defining a policy in the first place. As a result, she might lose touch with any rules she had held, and resort to a default of allowing access unless she is given a "good" reason to deny it.

Similarly, another explanation might be that reflecting on the reasons he might grant or deny access to his personal information leads the subject to believe that he must have a strong reason for denying a person access. Without concrete reasons to back them up, perhaps the user is unable to feel comfortable trusting his intuitions, and thereby resorts to the default of granting access to everyone.

Another idea that we have considered is the introspection questionnaire primed the subjects in some way to think "allow." Figure 3 shows the GUI that subjects saw; perhaps the "allow" in every question pertaining to individual profiles subliminally influenced the experimental subjects? Perhaps if the questions had read "Would you deny [Person X] access to.." (instead of "would you allow [Person X] to view...") we may have seen different answers.

Another possible contributing factor could be the ratio of male subjects to female subjects. It has been reported (e.g., the talk accompanying [4]) that female users tend to associate computer privacy leaks with potential threats to their physical person, suggesting a hypothesis that men are more likely to share sensitive data than women are. But (as we noted), the control and experimental groups had balanced gender ratios, so we don't think this is a significant factor in our results.

For whatever reason, it would seem that introspection somehow causes the subjects to lose the inclination to deny, and thus causes them to default to granting access.

7. CONCLUSION

Effective privacy and security, in practice, require that the policy a computer system actually enforces matches the policy the users intend. Our user study provides empirical evidence that the act of introspecting on issues affective privacy policy in social networks changes a user's intuitive policy to be far more open. If our goal is to prevent privacy spills, the straightforward approach of informing users of the issues and then letting them specify their policy actually makes things worse.

Future Work.

As noted above, one interesting follow-on might be to repeat the study, but with different ways of asking the data challenges—to see if the question wording itself interacts with the introspection priming. Another might be to carry out a experiment in which users generate a k such that a P_k can be explicitly evaluated—rather than just answer questions which let us gauge P_U. Alternatively, it might be in-

		Difference between averages	Difference between standard deviations	F-Test
1	Wade	14.58	8.58	8.81
2	Chathum	11.22	1.49	2.88
3	Arthur	11.66	7.31	8.12
4	Danny	9.63	1.18	4.29
5	Jake	13.72	1.68	5.33
6	Andrew V.	9.43	7.96	10.28
7	Amanda	6.34	5.62	2.83
8	Phil	7.72	4.43	1.44
9	Beth	5.53	5.33	0.79
10	Andrew P.	8.85	0.58	3.14
11	Samantha	13.15	12.52	12.23
12	Michael	5.63	3.22	1.26
13	Darcy	6.85	5.63	0.99
14	Megan	3.14	5.75	2.30
15	Maddie	6.00	1.43	2.40
16	Colleen	4.88	1.10	0.66
17	Peggy	12.00	0.69	5.48
18	Cam	6.27	1.54	0.95
19	Kate	14.28	3.86	6.98
20	Sarah	12.29	5.20	5.42

Table 6: The first column shows the difference between the averages for the control and introspective group that were displayed in Table 5. The second column shows the differences between the standard deviations presented in Table 5. The third column shows the results of an ANOVA F-Test on these averages—shaded rows indicate significant differences.

teresting to select a problem domain which would clearly admit the *Subject Matter Experts (SME)*, to correspond to the "taste experts" in Wilson and Schooler.

Other future work could also include repeating the study with a broader population (i.e., not just undergraduate students), with different types of policy (i.e., file system permissions), or other configurations that would attest to the generalizability of the results.

On a different note, we were amused to observe how many in the control group (who weren't sharing much) indicated they wanted to go back to Facebook and double-check their policies—but many in the introspective group (who were sharing everything) thought their Facebook policies were fine; a few even observed that they were *more conservative* on InnerCircle. This anecdotal evidence from the post-test feedback survey suggests that *self-reported* views of privacy and security practices may differ from reality and may be influenced greatly by priming. It would be interesting to explore these potential differences further—and to see whether any convention wisdom derived from self-reported gets upended.

Finally, a vital area of future work is to apply these results to real-world systems. If introspection reduces users' ability to craft effective policies, we as a community clearly need to identify alternate mechanisms, interfaces, or methods to facilitate this activity.

Acknowledgments

The authors would like to thank the reviewers for their helpful feedback and suggestions, as well as Deanna Caputo and Denise Anthony for their help in experiment design and analysis.

This research was supported in part by the Institute for Security, Technology, and Society (ISTS), the Institute for Information Infrastructure Protection (I3P), and the U.S. Department of Homeland Security, under Grant Award Number 2006-CS-001-000001.

8. REFERENCES

[1] L. F. Cranor. *Web Privacy with P3P*. O'Reilly, 2002.

[2] L. F. Cranor. Designing a Privacy Preference Specification Interface. In *Workshop on Human-Computer Interaction and Security Systems*, 2003.

[3] R. Dhamija, J. D. Tygar, and M. Hearst. Why Phishing Works. In *SIGCHI Conference on Human Factors in Computing Systems*, pages 581–590, 2006.

[4] P. Dourish. Security as a Practical Problem: Some Preliminary Observations of Everyday Mental Models. In *Workshop on Human-Computer Interaction and Security Systems*, 2003.

[5] E. F. Gehringer. The 'Memogate' Affair: A Case Study on Privacy in Computer Networks. In

Proceedings of the 2005 American Society for Engineering Education Annual Conference and Exposition, 2005.

[6] M. Gladwell. *Blink : The Power of Thinking Without Thinking*. Back Bay, 2007.

[7] N. S. Good and A. Krekelberg. Usability and Privacy: a Study of Kazaa P2P File-sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 137–144, 2003.

[8] C. Lagoze and V. Weissman. Towards a Policy Language for Humans and Computers. In *Proceedings of the Eighth European Conference on Digital Libraries*, pages 513–525, 2004.

[9] C. P. Masone. *Attribute-Based, Usefully Secure Email*. PhD thesis, Dartmouth College, Department of Computer Science, 2008. Technical Report TR2008-633.

[10] R. A. Maxion and R. W. Reeder. Improving User-Interface Dependability through Mitigation of Human Error. *International Journal of Human-Computer Studies*, 63:20–50, 2005.

[11] R. W. Reeder. *Expandable Grids: A User Interface Visualization*. PhD thesis, Carnegie Mellon University, School of Computer Science, 2008.

[12] W. Roger. Privacy isn't Public Knowledge. *USA Today*, June 2000.

[13] S. Trudeau. *The Effects of Introspection on Computer Security Policies*. Senior honors thesis, Dartmouth College, Department of Computer Science, 2009. Technical Report TR2009-652.

[14] T. D. Wilson and J. W. Schooler. Thinking Too Much: Introspection Can Reduce the Quality of Preferences and Decisions. *Journal of Personality and Social Psychology*, 60:181–192, 1991.

APPENDIX

A. FACEBOOK QUESTIONNAIRE

The following questions are related to your Facebook account. Please answer all questions honestly, and to the best of your knowledge.

1. Have you ever looked at your "Privacy Settings" on Facebook? (Yes, No)

2. Have you ever changed any of the settings in your "Privacy Settings"? (Yes, No)

3. On a scale of 1 (very closed) to 10 (extremely open), how open do you consider your Facebook privacy settings?

For the following questions, assume that Facebook has an interface that allows you to easily configure your security settings in any way you want.

4. How would you set the "Search Visibility" for your Facebook account? (My Networks, My Networks and Friends of Friends, My Networks and Friends, Friends of Friends, Only Friends, Custom—specify)

5. Who would you allow to view your Profile? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

6. Who would you allow to view your Basic Info? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

7. Who would you allow to view your Personal Info? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

8. Who would you allow to view your photo albums? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

9. Who would you allow to view your Friends list? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

10. Who would you allow to view the Wall Posts on your profile page? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

11. Who would you allow to view your Education Info? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

12. Who would you allow to view your Work Info? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

13. Who would you allow to view your IM Screen Name? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

14. Who would you allow to view your Mobile Phone Number? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

15. Who would you allow to view your Current Address? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

16. Who would you allow to view your Websites? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

17. Who would you allow to view your Current Residence? (Everyone, My Networks and Friends, People at Dartmouth

and Friends, Friends of Friends, Only Friends, Custom—specify)

18. Who would you allow to view your Primary Email Address? (Everyone, My Networks and Friends, People at Dartmouth and Friends, Friends of Friends, Only Friends, Custom—specify)

19. Do you limit any of your Facebook friends to only being able to view your Limited Profile? (Yes, No)

20. Is there anything on your Facebook account that you would not want an employer to see? (Yes, No)

21. Is there anything on your Facebook account that you would not want a parent to see? (Yes, No)

22. Is there anything on your Facebook account that you would not want a professor to see? (Yes, No)

23. Are you Facebook friends with any professors? (Yes, No)

24. Are you Facebook friends with any of your past or present bosses or employers? (Yes, No)

25. Does one of your parents have a Facebook account? (Yes, No)

26. Are you Facebook friends with one of your parents? (Yes, No)

27. Are you Facebook friends with someone you have never met before in real life? (Yes, No)