

[short paper] Privacy in Mobile Personalized Systems: The Effect of Disclosure Justifications

Bart P. Knijnenburg^{1,2}
Bart.K@uci.edu

Alfred Kobsa¹
Kobsa@uci.edu

Gokay Saldamli²
Gokay.S@sisa.samsung.com

¹ Bren School of Information and Computer Sciences
University of California, Irvine, USA

² Samsung R&D Research Center
San Jose, USA

ABSTRACT

Mobile applications increasingly use personalization, to maximize the relevance of displayed content and to simplify the user interaction. However, disclosing the personal information that is required for personalization makes many users feel uncomfortable. It is believed that empowering users to make their own privacy decisions by giving them adequate notifications and disclosure controls would increase their trust and satisfaction. However, in our study targeted at a mobile app recommender system, we saw the opposite user behavior: providing disclosure justifications decreased users' disclosure, trust in the company, and satisfaction with the system. We present and discuss our unexpected findings and point at approaches to solve this problem.

Categories and Subject Descriptors

H.1.2 [Models and Principles]: User/Machine Systems, H5.2 [Information Interfaces and Presentation]: User Interfaces—evaluation/methodology, theory and methods, K.4.1 [Computers and Society]: Public Policy Issues—privacy

General Terms

Design, Experimentation, Human Factors, Theory.

Keywords

Privacy calculus, information disclosure, justification, decision-making, personalization, recommender systems, user experience.

1. INTRODUCTION

Mobile applications often use personalization (e.g. recommendations or filtering) to reduce the amount of displayed information, increase data relevancy, and simplify navigation [25, 26, 32]. In general, the information required for personalization consists of users' demographical data (e.g. age, hobbies, income) and contextual data (e.g. app usage, calendar, location), which are implicitly gathered or explicitly requested from the user [17]. Privacy research has however shown that many people are not comfortable with disclosing diverse personal information [5, 16]. A suggested remedy is to give users control over what information they disclose [19, 30]; this allows them to trade off between the potential personalization benefits and the ensuing privacy risks [22, 27]. However, users often have a hard time making this trade-off due to their lack of knowledge about its positive and negative consequences [1, 2, 7, 8].

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) 2012, July 11-13, 2012, Washington, DC, USA.

Recent studies show that users can be assisted in their disclosure decisions by means of informative *justifications* for disclosing (or not disclosing) certain information [3, 6, 18, 24, 29]. Justifications allow users to make more meaningful disclosure decisions, and are considered as a requirement in many privacy norms [12, 31]. In fact, justifications may also “nudge” [28] users to disclose useful information (e.g. information needed for effective personalization) that they normally do not disclose.

Several different types of justifications have been studied. One is to explain *the reason* why the information is requested. For instance, Consolvo et al. [9] observe location disclosure in a mobile activity tracker, and show that users are mostly interested in knowing by whom and for what reason their personal data is requested. A reasonable explanation may prove the rightness of the disclosure request. In fact, the legitimacy of a request has been shown to increase protection beliefs [21] and disclosure [4]. Moreover, the proposed U.S. Consumer Privacy Bill of Rights suggests that “companies should provide clear descriptions of [...] why they need the data, how they will use it” [31].

Another type of justification is to highlight the *benefits* of disclosure. The reasoning behind this is that people trade off the benefits of disclosure with its potential risks (‘privacy calculus’, [10, 20]). Explaining the benefits of disclosure should tip the scales in favor of the benefits. Kobsa and Teltzrow [18] used contextualized explanations of privacy practices and personalization benefits. They found that their participants “were significantly more willing to share personal data with the website, rated its privacy practices and the perceived benefit resulting from data disclosure significantly higher, and also made considerably more purchases” when they saw such justifications. Wang and Benbasat [29] show that explanations of the benefits of disclosure increase trust in the competence, benevolence and integrity of their system.

A final type of justification is to appeal to the social norm: people can eschew doing their individual privacy calculus by conforming to the behavior of the majority. In an experiment by Acquisti et al. [3], participants were about 27% more likely to disclose when they learned that many others had decided to disclose the same information. On the other hand, Besmer et al. [6] find that social cues have barely any effect on users' Facebook privacy settings: only the small subset of users who take the time to customize their settings may be influenced by strong negative social cues. Patil et al. [24] also find that social navigation cues provided merely a secondary effect on disclosure.

Two trends seem to emerge from these previous studies with different types of justifications:

1. Most studies suggest or demonstrate that justifications increase disclosure.

2. Most researchers assume (and some even demonstrate [21, 29]) that providing justifications would ultimately increase user satisfaction.

In this paper we present a study of a mobile app recommender system that suggests the opposite: providing justifications not only decreased the disclosure rate but also users’ subjective valuations. Below we briefly present the main findings of the study¹, and discuss the reason why justifications decreased satisfaction and trust. Finally, we propose ways to solve this problem.

2. USER STUDY

Our study considers a mockup of a mobile app recommender that uses demographics data (e.g. age, hobbies, income) and context data (e.g. app usage, calendar entries, location) to provide users with recommendations for new applications for their phones [11]. The experiment only concerned the part of the system that collects personal information (i.e. no recommendations are given). The mockup does not run on a phone, but in an Internet browser. In order to make the experiment more realistic, participants were told that their data would be disclosed to the developer, a company named “Appy”². We reinforced this belief by ostensibly transferring users to the “Appy” website (with its own URL and branding) for the disclosure part of the experiment (see Figure 1).

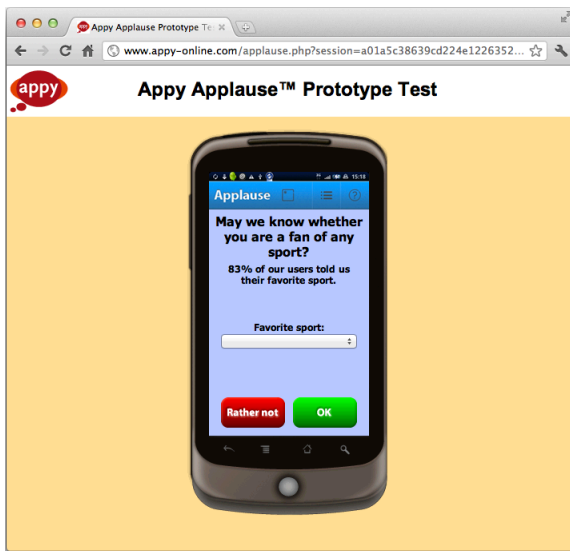


Figure 1. The website of “Appy”, on which participants do the disclosure part of the experiment.

2.1 Participants and procedure

491 participants were recruited via Amazon Turk and Craigslist. All participants (223 males) were adult smartphone users, with a median age range of 25-30. Participants were first given a short introduction to the system, and were specifically informed that they would be helping “Appy” to test the information disclosure part of the system. They were then randomly assigned to a condition (see below) and transferred to the “Appy” website.

¹ A comprehensive description of the study can be found in [14].

² This name was perceived as familiar and trustworthy in a pre-test that compared seven different company names and logos.

There, the system would make 31 disclosure requests, for 12 pieces of context data and 19 pieces of demographical data. In the context data requests, users were asked to indicate whether or not they would disclose the respective data (‘yes’ or ‘no’ answers). In the case of demographics requests, they were asked to provide the actual information or decline its disclosure. After 31 decisions, participants were transferred back to the experimenters’ website, where they would be asked for their subjective valuations.

2.2 Experimental manipulations

Participants were randomly assigned to one of the five justification conditions: no justification (control) or one of the four justification types shown in Table 1. The ‘useful for you’ and ‘useful for others’ justifications explain the benefits of disclosure (cf. [29]) in two different ways. The ‘number of others’ justification appeals to the social norm (cf. [3, 6, 24]). The ‘explanation’ justification states the reason for requesting the information (cf. [18]). The percentage figures that appear in some of the justifications are randomly chosen by the mockup.

Since our app recommender requests both demographical data and context data, the other manipulation is the order in which these data are requested: demographical data first or context data first.

Table 1. The justification messages used in our study

Type of justification	Message to user
Useful for you	“The recommendations will be about [XX]% better for you when you tell/allow us to use...”
Number of others	“[XX]% of our users told/allowed us to use...”
Useful for others	“[XX]% of our users received better recommendations when they told/let us...”
Explanation	“We can recommend apps that are [reason for request]”

2.3 Dependent variables

As outlined in the introduction, we were interested in the effect of justifications on the amount of disclosure of our participants and on their subjective valuations. Disclosure behavior was logged in our database as a yes/no decision of participant x for item y . Additionally, several subjective valuations of the disclosure process were measured with post-study questionnaires. The questionnaire items were submitted to a confirmatory factor analysis, resulting in the following factors:

- Perceived value of disclosure help (3 items, e.g. “The system helped me to make a tradeoff between privacy and usefulness”)
- Perceived privacy threats (3 items, e.g. “The system has too much information about me”)
- Trust in the company (4 items, e.g. “I believe this company is honest when it comes to using the information I provide”)
- Satisfaction with the system (6 items, from [15]; e.g. “Overall, I’m satisfied with the system”)

Information about the full questionnaires and the factor model can be found in [14].




Figure 2. The estimated effects of justification type (☐= ‘no justification’, ▨= ‘useful for you’, ▩= ‘number of others’, ▪= ‘useful for others’, ▫= ‘explanation’) and request order (different groups) on disclosure for each type of data (left vs. right). (¹ $p < .10$, * $p < .05$, ** $p < .01$, *** $p < .001$)

3. STUDY RESULTS

3.1 Disclosure behavior

The effect of the experimental manipulations on subjects’ disclosure decisions is analyzed using General Estimating Equations with a log link function. The dependent variable is the decision of participant x to disclose item y (yes/no). Since the 31 decisions of each participant are correlated, we impose a compound symmetric covariance structure.

The independent variables are the four justification types (tested against no justification), the order of the requests (demographics first vs. context first), as well as their interaction. To make a distinction between context and demographics disclosure, we also include the interactions of these variables with the type of data (context data vs. demographics). This results in a full factorial model: justification type \times order \times type of data³.

Figure 2 shows the estimated effects of justification type and request order on disclosure for each type of data. Disregarding justification type, the request order has a significant effect for both demographics data (79.2% for context first versus 93.1% for demographics first, $p < .001$) and context data (67.0% for context first versus 60.4% for demographics first, $p = .009$). The main trend is that the disclosure of a certain type of data is higher if that data type is requested first.

The interaction effect of justification type and request order is significant for both demographics data ($F(4, 9357) = 17.9$, $p < .001$) and context data ($F(4, 5906) = 5.32$, $p < .001$). We see that demographics disclosure is lower than the baseline (no justification) when requesting context first (87.1%) for the ‘useful for others’ justification (68.9%, $p = .001$) and somewhat lower for the ‘number of others’ justification (78.3%, $p = .064$). When requesting demographics first, the justifications are not significantly different from the baseline. Context data disclosure is lower than

the baseline when requesting context first (75.8%) for the ‘useful for you’ (65.5%, $p = .039$), ‘number of others’ (62.1%, $p = .006$), and ‘useful for others’ (63.3%, $p = .013$) justifications. When requesting demographics first (baseline: 65.9%), the disclosure is lower for the ‘number of others’ (51.6%, $p = .014$) and ‘useful for others’ (51.6%, $p = .013$) justifications. If anything, the justification messages thus *decrease* the level of disclosure, which, from the perspective of personalization, is a disappointing result.

3.2 Subjective valuations

To test the effect of the experimental manipulations on subjects’ subjective valuations, we regressed the factors from our factor analysis on the experimental manipulations and their interaction using a MIMIC model (a confirmatory factor analysis model with covariates) [23]. Figure 3 shows the estimated effects of justification type on the different factors. Request order and its interaction with justification type had no significant effect on these factors.

The justification messages have a significant effect on the perceived value of the disclosure help ($\chi^2(4) = 48.936$, $p < .001$). Specifically, the ‘useful for you’ ($\beta = 0.896$, $p < .001$), ‘useful for others’ ($\beta = 0.767$, $p < .001$), and ‘explanation’ ($\beta = 0.496$, $p = .003$) justifications are perceived to be more helpful than the baseline (no justification).

Overall, the justification messages have no significant effect on the perceived privacy threat ($\chi^2(4) = 6.410$, $p = .171$). If anything, participants perceive more threat when they received the ‘useful for others’ justification ($\beta = 0.351$, $p = .044$).

Overall, the justification messages have a marginally significant effect on participants’ trust in the company ($\chi^2(4) = 8.079$, $p < .089$). Specifically, participants who received the ‘useful for others’ ($\beta = -0.409$, $p = .006$), and somewhat the ‘explanation’ ($\beta = -0.261$, $p = .071$) justifications have less trust in the company.

The justification message have a significant effect on participants’ satisfaction with the system ($\chi^2(4) = 18.978$, $p < .001$). Specifically, they are less satisfied when they received the ‘useful for you’ ($\beta = -0.400$, $p = .010$), ‘number of others’ ($\beta = -0.440$, $p = .005$), ‘useful for others’ ($\beta = -0.661$, $p < .001$), and somewhat the ‘explanation’ ($\beta = -0.290$, $p = .064$) justifications.

³ We also ran a model that included the percentage shown in the justification messages. Although this percentage has a significant effect in the “number of others” condition, its effect is small compared to the main effects of justification type. We therefore left it out of the current analysis.




Figure 3. The estimated effects of justification type (‘no justification’ (fixed at zero), \diagup = ‘useful for you’, \diagdown = ‘number of others’, \boxplus = ‘useful for others’, \blacksquare = ‘explanation’) on subjective valuations. Scales are in sample standard deviations. ($^1p < .10$, $^*p < .05$, $^{**}p < .01$, $^{***}p < .001$)

4. REFLECTION

Disclosure justifications did not have the expected effects in our study, in contrast to what prior research had suggested. They did not increase users’ disclosure, trust in the company or satisfaction with the system, but rather decreased them. Less disclosure may negatively affect the personalization quality, which could further reduce satisfaction.

Why did this happen? Previous work has already demonstrated that people’s disclosure behavior is a trade-off between the potential benefits and risks of disclosure [7, 8]. As our justifications provided random percentages, some disclosures were randomly portrayed as positive (“The recommendations will be 94% better for you if you tell us...”) but others were portrayed as rather negative (“8% of our users told us...”). In terms of privacy calculus [10, 20], then, participants found the justifications helpful, but arguably only because they helped them decide that oftentimes the promised benefits were too low. As a result, they not only withheld the information, but they also became dissatisfied with the recommender for not providing better benefits. Indeed, a series of in-depth interviews about the system revealed that users typically treat the justification messages as warning signs: 11 out of 17 interviewees would take a low percentage in a justification message as a reason not to disclose (whereas only 5 participants mentioned a high percentage as a reason to disclose). Arguably then, since participants in our study acted more upon negative justifications, this also reduced their subjective valuations.

However, even at high percentages the justifications did not increase disclosure rates. It may thus be that these justifications carry an implicit warning, despite highlighting benefits, social norms, or the legitimacy of the request. By justifying the disclosure, we inadvertently signaled that the act of disclosure is not trivial and may involve risks: the justifications bring the disclosure decision to the foreground and demand users’ attention. Similarly, John et al. [13] demonstrate that professional looking sites may garner *higher* privacy concerns than their more ‘shady’ counterparts, because the professionalism may be so conspicuous that it primes users with the concept of privacy. Given the presence of such subtle priming effects, it may come as no surprise that justifications can also be interpreted negatively, despite their positive explicit message.

Does this mean that providing justifications is a bad idea? There seem to be pros and cons to it. Arguably, the system is originally “hiding” part of its costs of disclosure. By giving users a justification for disclosure, their satisfaction may go up or down to match what John et al. call the “true cost” of disclosure [13]. On the other hand, no message is purely objective, and hence even the “neutral” justifications that we provided to our participants may have nudged down their assessment because of the implicit warning they carried [28]. As the true cost of disclosure is hard to measure objectively, the utility of disclosure justifications remains a subjective issue.

5. CONCLUSION AND FUTURE WORK

In this paper we reported a study in which providing users with justifications to support their privacy-related decisions decreased their disclosure of personal data, and even their trust in the company and their satisfaction with the system. Although participants perceive our justifications as helpful, they ultimately seem not to benefit from them (at least not in the short run).

How could this problem be solved? We see two options. There may be a “golden justification” out there—one that we have yet to discover—that actually does increase users’ disclosure and/or subjective valuations. Future work could explore the space of possible explanation types more exhaustively and systematically. Note however that both screen size and attention span are limited in the mobile domain, so it will be a challenge to come up with a pithy message that actually works. One suggestion is to broaden the notion of “justification” to include visual representations.

Another solution is to tailor the justification to the user, in line with the above argument that the utility of justification is a subjective issue. Exploring our data in more depth, we indeed found evidence that certain justifications did work for some participants, but not for others. For instance, we found that gender and people’s general tendency to disclose information are determinants for the effectiveness of the justifications. Specifically, our preliminary results suggest that no justification will work for males with a low disclosure tendency, but that for females with a low disclosure tendency the ‘explanation’ justification increases subjective valuations. For males and females with a high disclosure tendency, the ‘useful for you’ justification works best. We plan to further explore this idea of tailored justifications in the future.

6. REFERENCES

- [1] Acquisti, A. and Grossklags, J. 2005. Privacy and Rationality in Individual Decision Making. *IEEE Security & Privacy*. 3, 1 (2005), 26–33.
- [2] Acquisti, A. and Grossklags, J. 2008. What Can Behavioral Economics Teach Us About Privacy? *Digital Privacy: Theory, Technologies, and Practices*. A. Acquisti, S. De Capitani di Vimercati, S. Gritzalis, and C. Lambrinoudakis, eds. Taylor & Francis. 363–377.
- [3] Acquisti, A., John, L.K. and Loewenstein, G. 2011. The Impact of Relative Standards on the Propensity to Disclose. *Journal of Marketing Research*. (Nov. 2011), 1–15.
- [4] Annacker, D., Spiekermann, S. and Strobel, M. 2001. E-privacy: Evaluating a new search cost in online environments. *e-Everything: e-Commerce, e-Government, e-Household, e-Democracy: 14th Bled Electronic Commerce Conference* (Bled, Slovenia, 2001).
- [5] Anton, A.I., Earp, J.B. and Young, J.D. 2010. How Internet Users' Privacy Concerns Have Evolved since 2002. *Security Privacy, IEEE*. 8, 1 (Feb. 2010), 21–27.
- [6] Besmer, A., Watson, J. and Lipford, H.R. 2010. The impact of social navigation on privacy policy configuration. *Proceedings of the Sixth Symposium on Usable Privacy and Security - SOUPS '10* (Redmond, Washington, 2010).
- [7] Brodie, C., Karat, C.M. and Karat, J. 2004. Creating an E-commerce environment where consumers are willing to share personal information. *Designing personalized user experiences in eCommerce*. (2004), 185–206.
- [8] Chellappa, R.K. and Sin, R.G. 2005. Personalization versus privacy: An empirical examination of the online consumer's dilemma. *Information Technology and Management*. 6, 2 (2005), 181–202.
- [9] Consolvo, S., Smith, I., Matthews, T., LaMarca, A., Tabert, J. and Powledge, P. 2005. Location disclosure to social relations: why, when, & what people want to share. (Portland, Oregon, USA, 2005), 81–90.
- [10] Culnan, M.J. and Armstrong, P.K. 1999. Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation. *Organization Science*. 10, 1 (1999), 104–115.
- [11] Davidsson, C. and Moritz, S. 2011. Utilizing implicit feedback and context to recommend mobile applications from first use. *Proceedings of the 2011 Workshop on Context-awareness in Retrieval and Recommendation - CaRR '11* (Palo Alto, California, 2011), 19–22.
- [12] EU 2011. *Regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)*.
- [13] John, L.K., Acquisti, A. and Loewenstein, G. 2011. Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. *The Journal of Consumer Research*. 37, 5 (Feb. 2011), 858–873.
- [14] Knijnenburg, B.P. and Kobsa, A. 2012. Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems. *submitted to the ACM Transactions on Intelligent Interactive Systems, Special Issue on Human Decision Making and Recommender Systems*. http://www.isr.uci.edu/tech_reports/UCI-ISR-12-1.pdf. (2012).
- [15] Knijnenburg, B.P., Willemsen, M.C., Gantner, Z., Soncu, H. and Newell, C. 2012. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction*. 22, 4-5 (2012).
- [16] Kobsa, A. 2007. Privacy-enhanced web personalization. *The adaptive web* (2007), 628–670.
- [17] Kobsa, A., Koenemann, J. and Pohl, W. 2001. Personalized Hypermedia Presentation Techniques for Improving Customer Relationships. *The Knowledge Engineering Review*. 16, 2 (2001), 111–155.
- [18] Kobsa, A. and Teltzrow, M. 2005. Contextualized communication of privacy practices and personalization benefits: Impacts on users' data sharing and purchase behavior. *Privacy Enhancing Technologies* (2005), 329–343.
- [19] Kolter, J. and Pernul, G. 2009. Generating User-Understandable Privacy Preferences. *2009 International Conference on Availability, Reliability and Security* (Fukuoka, Japan, 2009), 299–306.
- [20] Laufer, R.S. and Wolfe, M. 1977. Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory. *Journal of Social Issues*. 33, 3 (1977), 22–42.
- [21] Li, H., Sarathy, R. and Xu, H. 2010. Understanding situational online information disclosure as a privacy calculus. *Journal of Computer Information Systems*. 51, 1 (2010), 62–71.
- [22] Mabley, K. 2000. *Privacy vs. Personalization: Part III*. Cyber Dialogue, Inc.
- [23] Muthén, B.O. 2011. Latent Variable Modeling in Heterogeneous Populations. *Psychometrika*. 54, 4 (Sep. 2011), 557–585.
- [24] Patil, S., Page, X. and Kobsa, A. 2011. With a little help from my friends: can social navigation inform interpersonal privacy preferences? *Proceedings of the ACM 2011 conference on Computer supported cooperative work*. (2011), 391–394.
- [25] Ricci, F. 2011. Mobile Recommender Systems. *Information Technology & Tourism*. 12, 3 (2011), 205–231.
- [26] Schwinger, W., Grün, C., Pröll, B. and Retschitzegger, W. 2009. Context-Awareness in Mobile Tourist Guides. *Handbook of Research on Mobile Multimedia, Second Edition*. IGI Global. 534–552.
- [27] Taylor, D., Davis, D. and Jilappalli, R. 2009. Privacy concern and online personalization: The moderating effects of information control and compensation. *Electronic Commerce Research*. 9, 3 (Sep. 2009), 203–223.
- [28] Thaler, R.H. and Sunstein, C. 2008. *Nudge: improving decisions about health, wealth, and happiness*. Yale University Press.
- [29] Wang, W. and Benbasat, I. 2007. Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs. *Journal of Management Information Systems*. 23, 4 (2007), 217–246.
- [30] Wenning, R. and Schunter, M. 2006. *The Platform for Privacy Preferences 1.1 (P3P1.1) Specification*. W3C Working Group Note.
- [31] White House 2012. *Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Economy*. White House.
- [32] Zhang, D. and Lai, J. 2011. Can Convenience and Effectiveness Converge in Mobile Web? A Critique of the State-of-the-Art Adaptation Techniques for Web Navigation on Mobile Handheld Devices. *International Journal of Human-Computer Interaction*. 27, 12 (2011), 1133–1160.