Full Paper accepted to the ACM conference on Intelligent User Interfaces (IUI2013).
In press; this is the camera-ready author copy.

Helping Users with Information Disclosure Decisions: Potential for Adaptation

Bart P. Knijnenburg
Donald Bren School of
Information and Computer Sciences
University of California, Irvine, USA
bart.k@uci.edu

Alfred Kobsa
Donald Bren School of
Information and Computer Sciences
University of California, Irvine, USA
kobsa@uci.edu

ABSTRACT
Personalization relies on personal data about each individual user. Users are quite often reluctant though to disclose information about themselves and to be "tracked" by a system. We investigated whether different types of rationales (justifications) for disclosure that have been suggested in the privacy literature would increase users' willingness to divulge demographic and contextual information about themselves, and would raise their satisfaction with the system. We also looked at the effect of the order of requests, owing to findings from the literature. Our experiment with a mockup of a mobile app recommender shows that there is no single strategy that is optimal for everyone. Heuristics can be defined though that select for each user the most effective justification to raise disclosure or satisfaction, taking the user's gender, disclosure tendency, and the type of solicited personal information into account. We discuss the implications of these findings for research aimed at personalizing privacy strategies to each individual user.

Author Keywords
Recommender systems; information disclosure; user experience; satisfaction; user characteristics; adaptive interfaces; privacy; trust

ACM Classification Keywords
H.1.2 [**Models and Principles**]: User/Machine Systems, H.5.2 [**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, Measurement

INTRODUCTION
User modeling and personalization typically rely on implicitly gathered or explicitly requested personal information about the user [1]. Privacy research has however shown that quite a few people do not feel comfortable disclosing diverse personal information [2,3]. A typical remedy in this conflict is to give users control over what information they disclose [4], since this would allow them to trade off the potential personalization benefits of disclosure with the ensuing privacy risks [5,25]. However, users often have a hard time making this trade-off, since they lack knowledge about its positive and negative consequences [6,7]. This paper explores strategies to help users with their disclosure decisions in such a way that the amount of disclosure increases without decreased user satisfaction. We demonstrate that there is a potential for adaptation in this endeavor: the optimal strategy depends on the characteristics of the user and the optimization goal of the system (increasing the disclosure of two types of requested information, and/or raising users' satisfaction). Our work thus stands in the tradition of tailoring privacy to users' needs [8]. We discuss how the potential for adaptation that became manifest in our work can be translated into real-world systems that adapt their information request strategy to the user, in order to improve information disclosure while at the same time respecting users' privacy and satisfaction.

RELATED WORK

Strategies for Helping Users with Disclosure Decisions
Recent studies show that users can be assisted in their disclosure decisions by means of *justifications*. For instance, by informing users about the disclosure decisions of other users [9], or more specifically their friends [10], they become slightly more likely to conform to the perceived social norm in their own disclosure decisions. Others suggest justifications that inform users about the reason for requesting certain information [11] or the benefit of disclosing the information [8,12].

Another way to influence users' disclosure decisions is by changing the *order* of the disclosure requests. Acquisti et al. [13] showed that the average level of disclosure was higher when requests were made in decreasing order of intrusiveness. This was especially true for the more intrusive requests.

Potential for Adaptation
In [21, 27], we compared the effect of four different types of justifications against a system without such justifications. To our surprise, the system *without* justifications resulted in the most disclosure and highest satisfaction. Upon further analysis of these surprising results, we noticed that the

optimal justification depended on the characteristics of the user, type of information requested, and the goal of the system. In this paper we present this analysis, but we first outline existing work that corroborates our general findings.

User characteristics

Privacy concerns vary extensively across the population. One of the most cited results in privacy research is that people can be divided into three broad categories: privacy fundamentalists, pragmatists, and unconcerned [14]. Moreover, females and older people tend to have higher privacy concerns [15], especially when it comes to their location [16]. Due to these personal differences in privacy concerns, optimal strategies to assist disclosure decisions may not be universal, but may rather depend on the characteristics of each individual user.

Type of information requested

Users have different attitudes towards disclosing different types of information. Specifically, information that the user provides explicitly (demographic data) seems to raise less concern than information that the system extracts from user behavior on the quiet (context data). Although the latter is more convenient, it is also less controllable and may be incorrectly interpreted [17,18]. Due to these differences, the optimal strategy may be different for these two types of information.

Optimization goal of the system

Most studies on information disclosure have a behavioral goal: they try to increase the amount of users' disclosure. If these strategies are effective, one could argue that this is because they reduce privacy fears and/or increase trust in the company's privacy practices [12]. However, the increased disclosure may be unrelated to rational decision processes [13]. In fact, the same strategies that increase disclosure could inadvertently increase users' privacy fears and decrease their trust. While increased disclosure may eventually lead to higher personalization benefits, these unintended effects could at the same time lower user satisfaction. It therefore seems crucial to strive for both goals simultaneously.

ONLINE EXPERIMENT

Experimental Setup

Our experiment tests the hypothesis that the optimal strategy to help users with their disclosure decisions depends on the characteristics of the user, the type of information requested, and the optimization goal of the system. It considers a mockup of a mobile app recommender, which uses demographic data (e.g. age, hobbies, income) and context data (e.g. app usage, calendar, location) to provide users with recommendations for new applications for their phones [19]. The experiment only concerns 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, users were told that their data would in fact be disclosed to the developer, a company named "Appy"[1]. 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).

Participants

441 participants were recruited via Amazon Turk, a recruitment source that has become very popular for conducting user studies [20]. We only allowed participants from the United States and asked a number of reverse-coded and comprehension-testing questions to ascertain validity. An additional 52 participants were recruited via Craigslist, and we found no significant differences between the two groups. All 493 participants (223 males, 266 females, 4 did not disclose) were adult smartphone users, with a median age range of 25-30.

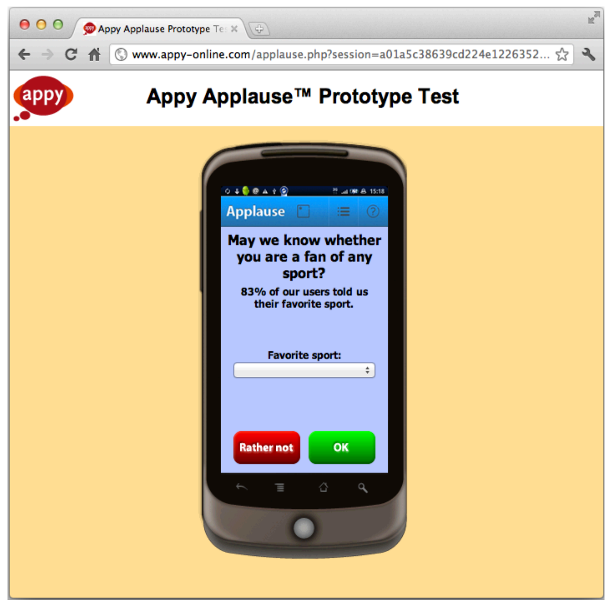

Figure 1: The website of "Appy", on which participants perform the disclosure part of the experiment.

Procedure

Participants were first given a short introduction to the mobile app recommender, and were specifically informed that they would be helping the Appy company to test the information disclosure part of the system. They were then randomly assigned to one of the experimental conditions (see below) and transferred to the Appy website. There, the system would make 31 disclosure requests, for 12 pieces of context data and 19 pieces of demographic data. In the context data requests, users were asked to indicate whether or not they would disclose the respective data (with 'yes' and 'no' as answers). In the case of demographics requests, they were asked to provide the actual information or decline

---

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

its disclosure[2]. We logged users' disclosure decisions in our database. After 31 decisions, participants were transferred back to the experimenters' website, where they answered questions about their subjective valuations.

Conditions: Type of Justification and Request Order

Four different justifications (see Table 1) are tested against the baseline of no justification, resulting in a total of five conditions. The 'useful for you' and 'useful for others' justifications explain the benefits of disclosure (cf. [12]) in two different ways. The 'number of others' justification appeals to the social norm (cf. [9,10]). The 'explanation' justification gives the reason for requesting the information (cf. [11]). The percentages that are shown in some of the justifications are randomly generated in the mockup, which makes the justifications more realistic. The reason presented in the explanation also varies enough per item to avoid excessive monotony.

Since our app recommender requests both demographic data and context data, we also manipulated the order in which these types of data are requested: demographic data first or context data first.

Justification	Message to user
None	[no justification message]
Useful for you	"The recommendations will be about [XX]% better for you when you tell us/allow us to use…"
Number of others	"[XX]% of our users told us/allowed us to use…"
Useful for others	"[XX]% of our users received better re–commendations when they told us/let us…
Explanation	"We can recommend apps that are [reason for request]"

Table 1: The justification messages evaluated in our study. Figure 1 shows how these messages are presented to the user.

User Characteristics: Gender and Disclosure Tendency

Two user characteristics are considered that may influence the relative impact of our strategies: gender and disclosure tendency[3]. Disclosure tendency is measured by splitting the users into two groups: one with a low disclosure tendency (up to 22 disclosed items), which comprises 33.3% of the participants, and one with high disclosure tendency (23-31 disclosed items). Several researchers found that people's actual disclosure behavior is only weakly related to their stated concerns [22,23,24], which is why we opt to split users based on their actual behaviors. The results of our analyses are robust under different splits: a 25-75% split provided very similar results, and also a three-way split failed to provide additional insights.

Gender and disclosure tendency each split the participants into two groups. To gauge the differences between these groups, we measured several other user characteristics in the post-experimental questionnaires, subjected them to a Confirmatory Factor Analysis, and regressed them on disclosure tendency and gender:

- Stated privacy concerns (3 items, adapted from [17]; e.g. "I am concerned about threats to my personal privacy today")
- Data collection concerns (5 items, adapted from [17]; e.g. "I'm concerned that online companies are collecting too much personal information about me")
- Control concerns (2 items, adapted from [18]; e.g. "Control of personal information lies at the heart of online privacy")
- Mobile internet usage (2 items, e.g. "I regularly use my phone to browse the Internet")
- Tech-savviness (3 items, e.g. "People ask me to fix their computer")

Participants with a low disclosure tendency have higher stated privacy concerns ($\beta = 0.379$, $p = .001$) and data collection concerns ($\beta = 0.696$, $p < .001$), are less phone-savvy ($\beta = -0.388$, $p < .001$) and less tech-savvy ($\beta = -0.222$, $p = .035$). Moreover, females are less tech-savvy ($\beta = -0.766$, $p < .001$). Finally, the percentage of people with a low disclosure tendency is not significantly different between males (32.7%) and females (33.1%).

Dependent Variables

We consider three different goals in our system: increasing the disclosure of context data, increasing the disclosure of demographic data, and increasing the users' subjective experience of the system. The first two goals are tested in a single analysis in which disclosure is regressed on the conditions and moderators; the type of data (demographic vs. context data) is included as an additional moderator.

Several subjective valuations of the disclosure process are measured with questionnaires, which are submitted to a Confirmatory Factor Analysis (CFA)[4] and regressed on the conditions and moderators:

- Perceived disclosure help (3 items, e.g. "The system helped me to make a tradeoff between privacy and usefulness")
- Perceived privacy threat (3 items, e.g. "The system has too much information about me")

[2] The fact that the (real) demographics data disclosure was actually *higher* than the (make believe) context disclosure suggests that participants answered both demographics and context requests as if they were real.

[3] A third analyzed characteristic, age, turned out to be non-significant in the analyses.

[4] A CFA tests the internal consistency of a set of indicators, and combines them into a single, normally distributed scale.

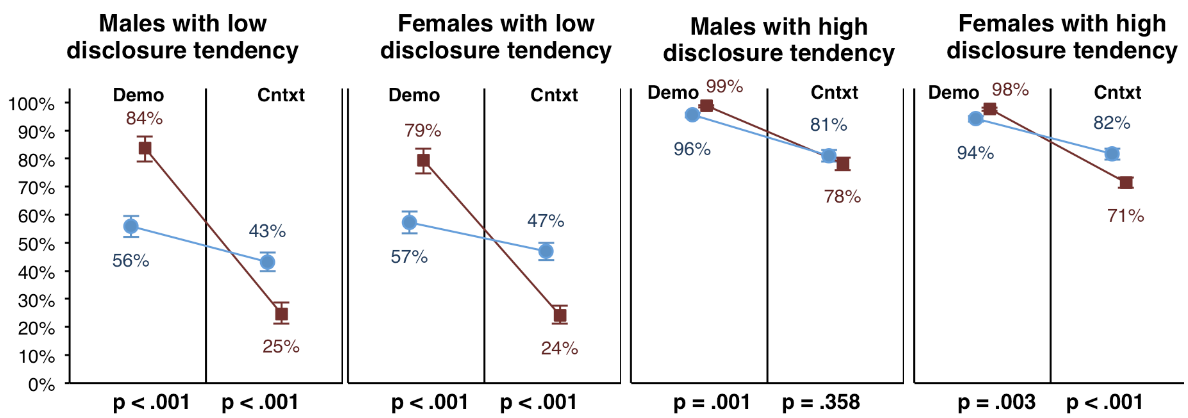

Figure 2: The effect of request order on disclosure per user group (disregarding justification type). Error bars are ±1 standard error.

- Trust in company privacy practices (3 items, adapted from [18]; e.g. "I believe this company is honest when it comes to using the information I provide")
- Satisfaction (6 items, from [22]; e.g. "Overall, I'm satisfied with the system")

RESULTS

Disclosure Behavior

The effect of our strategies[5] on subjects' disclosure decisions is analyzed using General Estimating Equations (GEE) with a log link function[6]. The dependent variable is the decision of participant *x* to disclose item *y* (yes/no). Because the 31 decisions of each participant are correlated, we impose a compound symmetric covariance structure.

The independent variables are the "strategies": the justification types (tested against no justification), the order of the disclosure requests (demographics first vs. context first), as well as their interaction. The interactions of the strategies with the user characteristics (gender and disclosure tendency) are also included. Finally, in order to make a distinction between the two goals of increasing context data disclosure and increasing demographics disclosure, the interactions of all aforementioned variables with the type of data (context data vs. demographics) are tested as well. The highest-order effect tested is thus: justification type × order × gender × disclosure tendency × type of data.

Disregarding justification type, the request order has a significant effect in almost all groups, and the main trend is that the disclosure of a certain type of data is higher if that data type is requested first (i.e. demographics disclosure is higher when demographics are requested first, and likewise for context data, $F(1,481) = 65.62$, $p < .001$). This effect is stronger for participants with a low disclosure tendency than for those with a high disclosure tendency (interaction: $F(1,481) = 43.68$, $p < .001$), but there is no difference between males and females. The mean disclosure rate for each request order and the p-value of the difference within each user group are displayed in Figure 2.

The request order differences can be explained in two ways: a growing concern with the accumulated information collected by the system, or boredom/fatigue resulting from answering 31 requests. If the latter were the case, the order effect should be most pronounced for demographics, because in our study disclosure is more laborious for demographics than for context data (selecting the answer from a dropdown menu vs. clicking a 'yes'/'no' button). However, the effect of order is stronger for context data than for demographics. The order effect is thus most likely *not* due to boredom or fatigue.

Figure 3 displays the combined effect of justification type and request order (i.e. the effect of the strategy) for the four user groups (males/females with low/high disclosure tendency). An ANOVA test of the five-way interaction of justification type × order × gender × disclosure tendency × type of data showed that the best strategy differs by gender, disclosure tendency and data type ($F(4,449) = 2.95$, $p = .020$). Aside from this five-way effect, there are significant lower-level interactions of strategy and disclosure tendency ($F(4,449) = 4.25$, $p = .002$) and strategy and data type ($F(4,449) = 2.87$, $p = .023$).

Within each user group, we compare the best strategy for each data type (marked with an arrow in Figure 3) against all other strategies. Strategies that perform significantly worse[7] than the best strategy are labeled with a *p*-value.

[5] We define "strategy" as the combination of a certain type of justification and a certain request order. We thus test 5×2 = 10 strategies.

[6] Our outcomes are correlated (each participant is represented by 31 data points) and non-normal (the data consists of 0's and 1's). GEE models are linear regression models that can robustly handle correlated, non-normal data.

[7] To reduce the risk of incorrectly claiming that a certain strategy performs equally well as the best strategy (Type I error), we take p < .10 as our cut-off value.

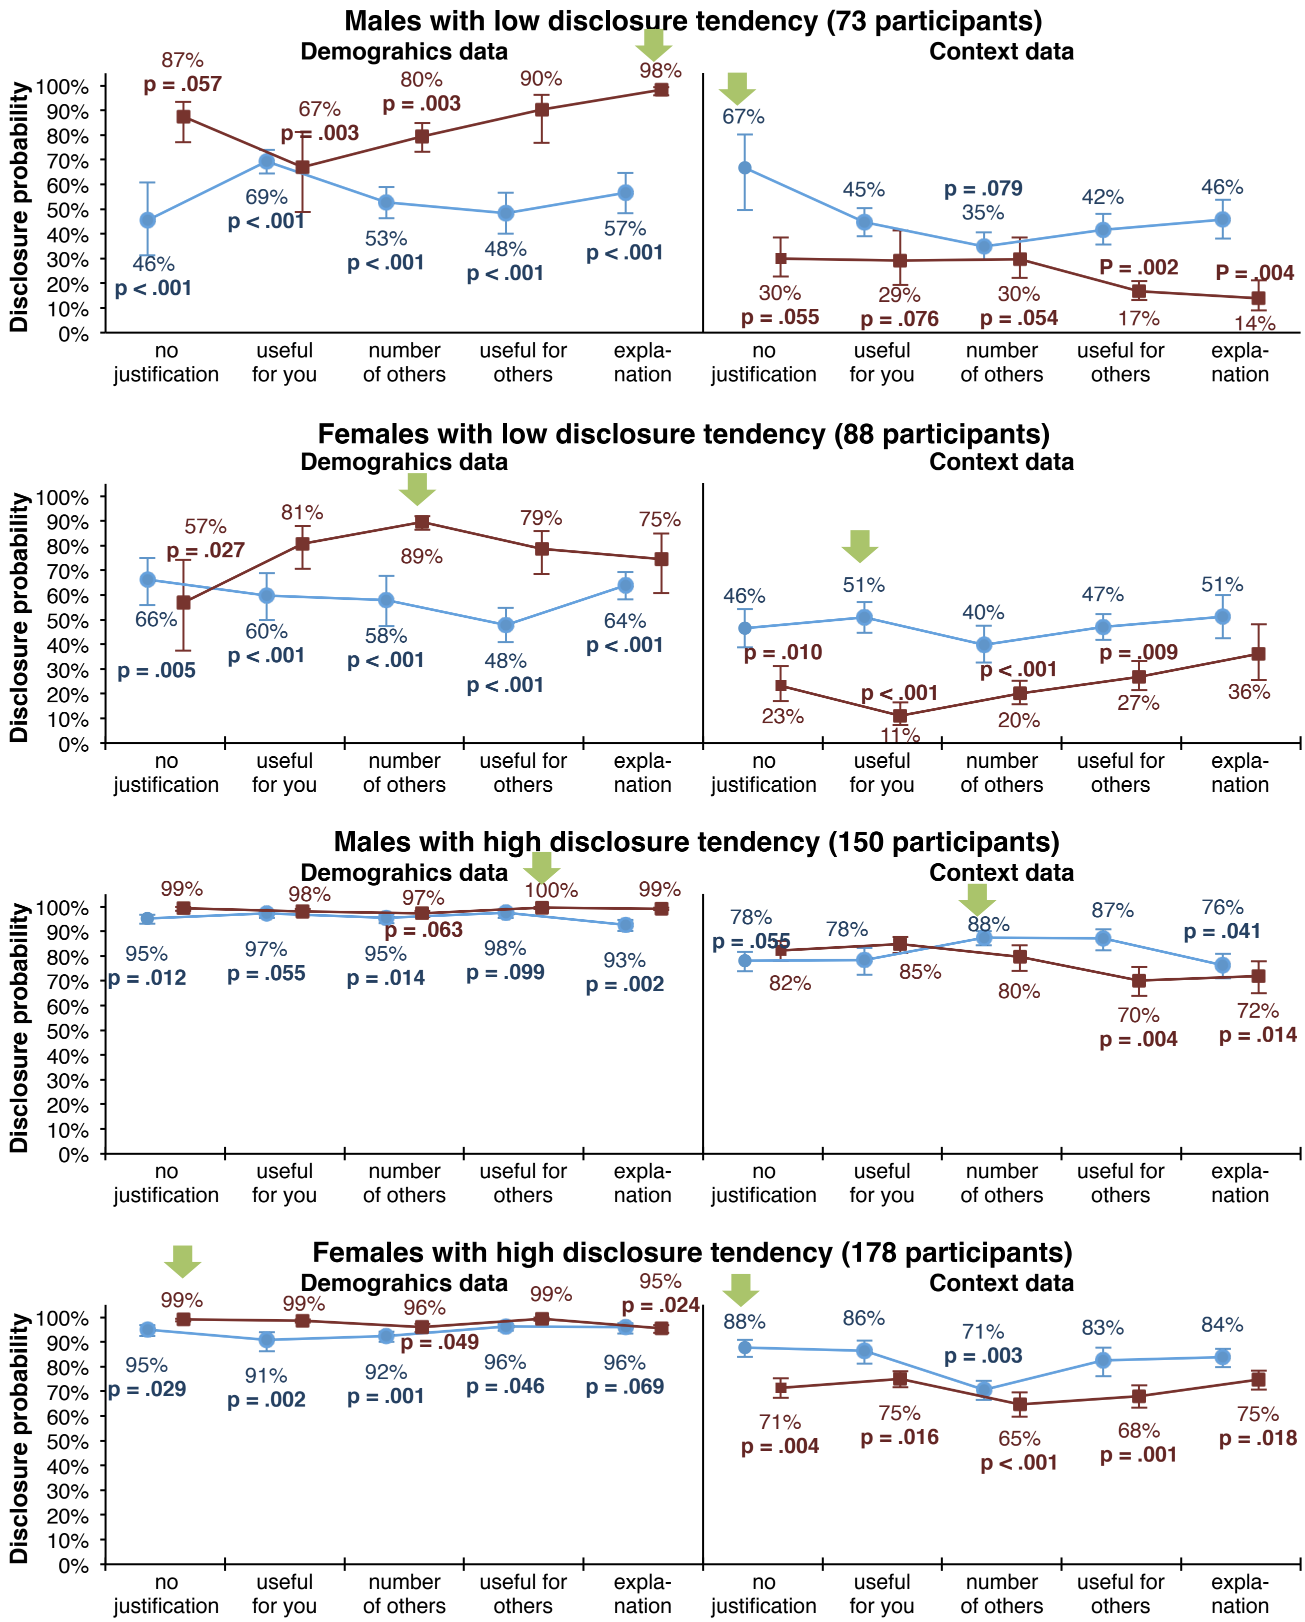

Figure 3: The effects of justification type and request order on disclosure (for each type of data, gender, and disclosure tendency). Error bars are ±1 standard error. The best strategy is labeled with an arrow; strategies with a p-value perform significantly worse.

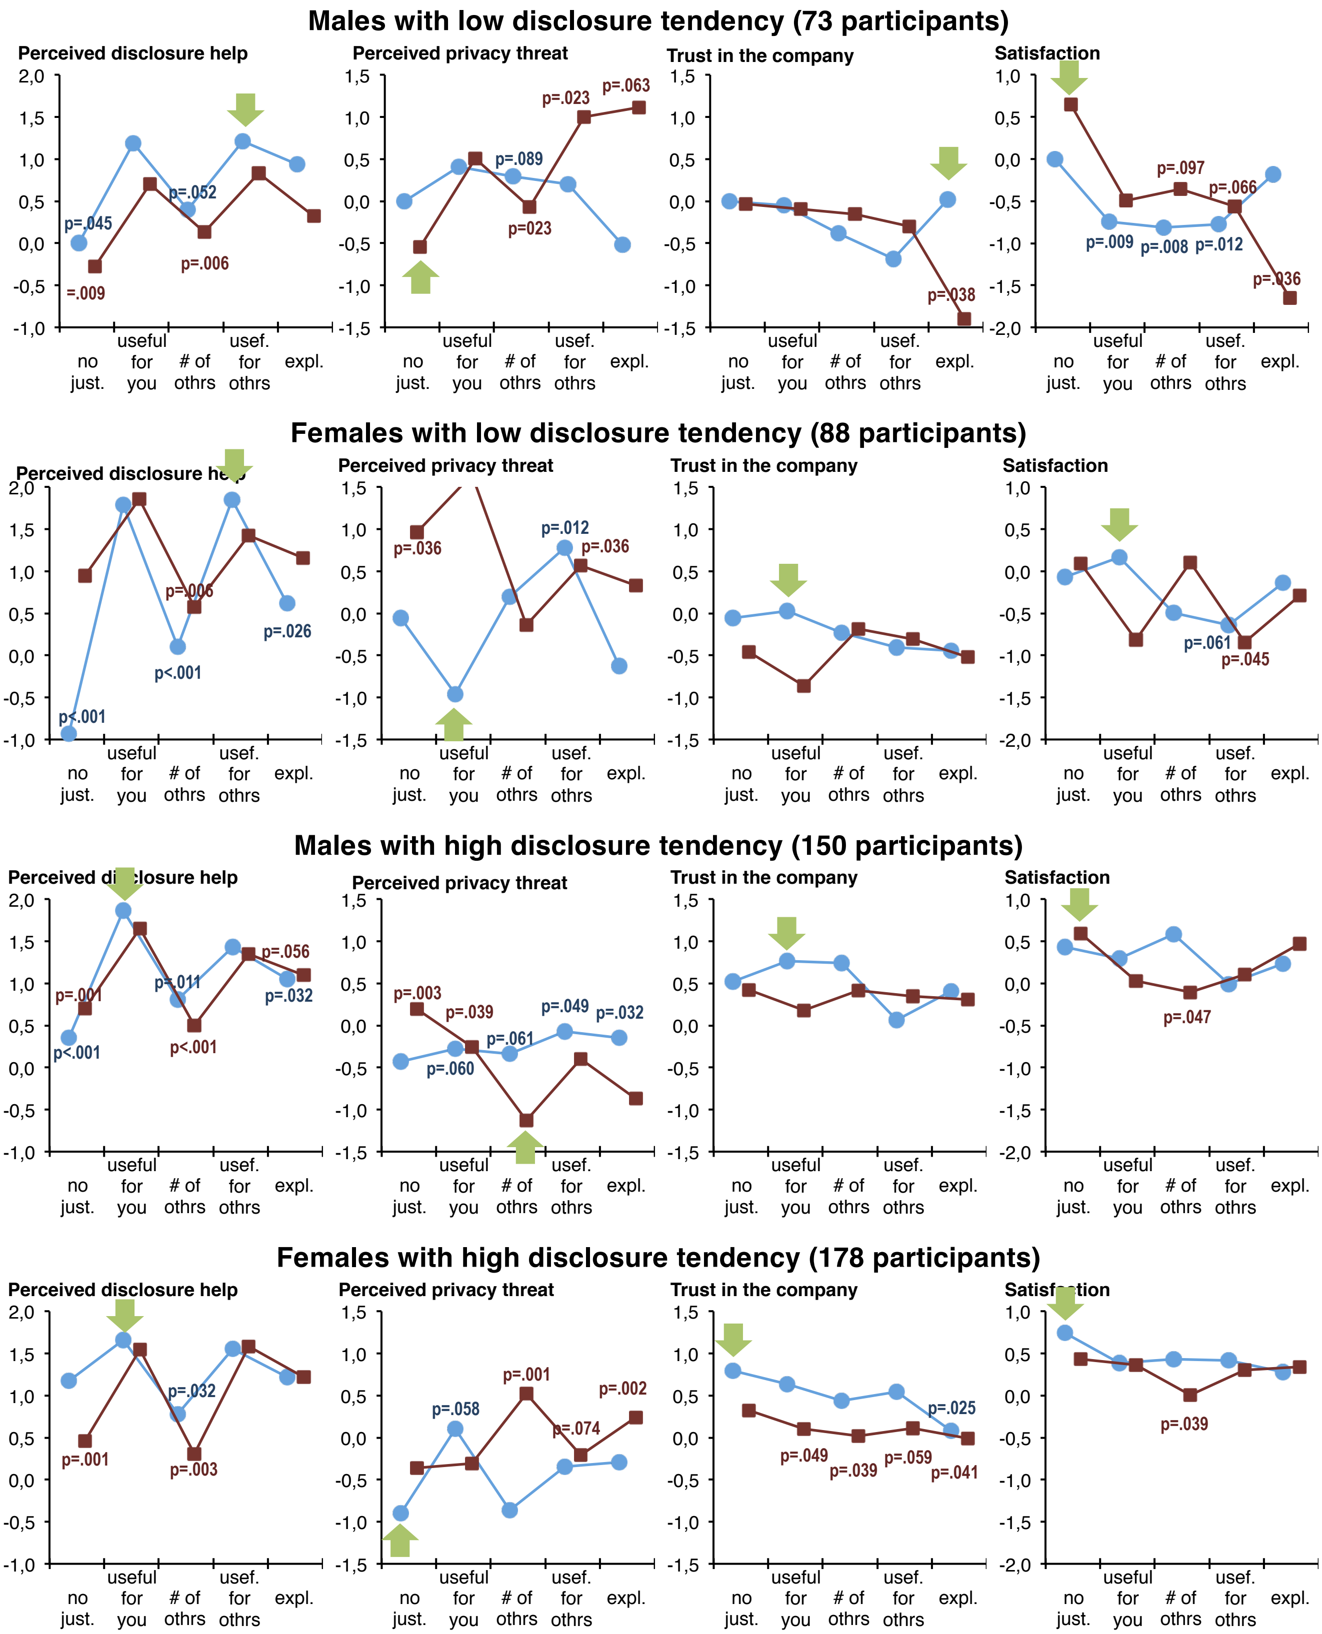

Figure 4: The estimated effects of justification type and request order on user satisfaction for each gender and disclosure tendency. Since the outcomes are scale-free factor scores, the y-axis is scaled in sample standard deviations (i.e. 95% of the participants fall within a 4.0-range), and the value for [male, low disclosure tendency, context first, no justification] is set to zero.

Subjective Valuations

We analyze the effect of the strategies on subjective valuations by submitting the questionnaire items to a Confirmatory Factor Analysis and by regressing resulting satisfaction factors on the strategies and user characteristics. The dependent variables in these analyses are: perceived disclosure help, perceived privacy threat, trust in company privacy practices, and overall satisfaction. The independent variables are the strategies (5 justification types × 2 orders), and the interactions of the strategies with gender and disclosure tendency. In effect, each dependent variable is regressed on justification type × order × gender × disclosure tendency.

Figure 4 displays the estimated effects of justification type and request order on the subjective valuations for each gender and disclosure tendency. Disregarding justification type, the request order has a significant effect on perceived disclosure help for males with a low disclosure tendency (β = -0.533, p = .023), where requesting context data first generally leads to a higher level of perceived disclosure help. The request order also has a significant effect on perceived privacy threat (β = 0.425, p = .024) and trust in the company (β = -0.340, p = .043) for females with a high disclosure tendency, where requesting context data first leads to less threat and more trust.

Figure 4 compares for each group the best strategy (marked with an arrow) against all other strategies. Strategies that perform significantly worse than the best strategy are labeled with a p-value.

HEURISTICS FOR SELECTING THE BEST STRATEGY

The results show that the best strategy depends on users' disclosure tendency and gender. It also depends on the goal of the system: some strategies increase disclosure of one type of data but not the other, and some increase disclosure but at the same time reduce users' satisfaction. We therefore suggest that the strategy should be adapted to the optimization goal of the system and the characteristics of the user. Table 4 outlines heuristics for selecting the best strategy for each type of user, given a certain system goal. Below we reflect on these suggested heuristics.

Best Strategy to Achieve High Demographics Disclosure

To get high demographics disclosure, one should ask for demographics first. Users with high disclosure tendency do not require a justification. Users with low disclosure tendency require a justification, the best one being 'number of others' for females and 'explanation' for males.

Best Strategy to Achieve High Context Data Disclosure

To get high context data disclosure, one should ask for context data first. No justification is required, but males with high disclosure tendency disclose more with the 'number of others' or 'useful for others' justification.

Best Strategy to Achieve High Total Disclosure

Since it is best to ask demographics first to increase demographics disclosure, and context first to increase context disclosure, increasing total disclosure requires a compromise. The best way to attain this compromise is to first choose a preferred request order, and then to select a

User type	Context first	Demographics first
Males with low disclosure tendency	The 'useful for you' justification gives the highest demographics disclosure.	Providing no justification gives the highest context disclosure.
Females with low disclosure tendency	Providing no justification gives the highest demographics disclosure.	The 'explanation' justification keeps context disclosure on par.
Males with high disclosure tendency	The 'useful for others' justification keeps demographics disclosure almost on par.	The 'useful for you' justification keeps context disclosure on par.
Females with high disclosure tendency	Providing no justification gives a high demographics disclosure.	The 'useful for you' justification gives the highest context disclosure.

Table 2: Best strategies to achieve high overall disclosures.

User type	Best strategy
Males with low disclosure tendency	Demographics first with 'useful for you'.
Males with high disclosure tendency	The 'useful for you' justification in any order.
Females with low disclosure tendency	Context first with 'useful for you'.
Females with high disclosure tendency	Context first with no justification, but 'useful for you' is second best.

Table 3: Best strategies to achieve high user satisfaction.

justification message that minimizes the disclosure decrease in the data that is requested last. The best message for this purpose depends on user characteristics (see Table 2).

Best Strategy to Achieve High User Satisfaction

To get high subjective valuations, the 'useful for you' justification generally works well, but see Table 3 for a more detailed appraisal.

Best Strategy Overall

When both the disclosure rates and satisfaction are to be optimized, one has to strike a compromise. This is possible with the following strategies: For users with low disclosure tendency, request demographics first, and use no justification for males and the 'explanation' justification for females. For users with high disclosure tendency, request demographics first with the 'useful for you' message for males, and request context first with no justification for females.

Best strategies for MALES with LOW disclosure tendency	
Goal	*Best strategy*
High demographics disclosure	Demographics first, 'explanation' justification.
High context data disclosure	Context first, no justification.
High overall disclosure	Context first, 'useful for you' justification.
High satisfaction	Context first, 'useful for others' justification or demographics first, 'useful for you' justification.
All of the above	Demographics first, no justification.
Best strategies for FEMALES with LOW disclosure tendency	
Goal	*Best strategy*
High demographics disclosure	Demographics first, 'number of others' justification.
High context data disclosure	Context first, 'useful for you' justification.
High overall disclosure	Demographics first, 'explanation' justification.
High satisfaction	Context first, 'useful for you' justification.
All of the above	Demographics first, 'explanation' justification.
Best strategies for MALES with HIGH disclosure tendency	
Goal	*Best strategy*
High demographics disclosure	Demographics first with any justification except 'number of others'.
High context data disclosure	Context first, 'number of others' or 'useful for others' justification.
High overall disclosure	Demographics first with no justification or the 'useful for you' justification, or context first with 'useful for others' justification.
High satisfaction	Demographics first, 'useful for others' or 'explanation' justification.
All of the above	Demographics first, 'useful for you' justification.
Best strategies for FEMALES with HIGH disclosure tendency	
Goal	*Best strategy*
High demographics disclosure	Demographics first with no justification, the 'useful for you' justification, or the 'useful for others' justification.
High context data disclosure	Context first with no justification.
High overall disclosure	Context first with no justification.
High satisfaction	Context first with no justification.
All of the above	Context first with no justification.

Table 4: Heuristics to find the best strategy, based on user characteristics (gender, disclosure tendency) and optimization goals.

Implementing the Strategies

To follow the heuristics, a system would have to discover the users' characteristics before or during the interaction. Gender can just be the first item to request. In fact, gender disclosure was the highest of all items in our study (94.9%), and hence we expect that asking for it first will not raise any concerns. To correctly determine the users' disclosure tendency, the system would have to first ask a number of potentially invasive questions, which is not desirable. Alternatively, one could ask about (or otherwise determine) the users' stated privacy concerns, mobile Internet usage and/or tech-savvyness, since these characteristics are related to users' disclosure tendency (see the section on user characteristics, and also [15,16]).

CONCLUSION AND FUTURE WORK

Information disclosure by users is an indispensible prerequisite for personalized systems, and strategies that increase disclosure may improve the accuracy of personalization. Our study shows that there is no single best justification strategy for all users to raise disclosure. Yet, heuristics can be developed to select an optimal strategy given the user characteristics, the type of information requested, and the optimization goal.

In terms of complexity, the heuristics presented in this paper fall between simple HCI design guidelines and more complex user models often used in adaptive systems. Although they can be realistically implemented, we do not have an intuitive rationale behind the specific heuristics we found. Future work should aspire a more fundamental understanding of the context-dependence of privacy behaviors (e.g. [26]).

In our current research, a user is only exposed to one of five types of justifications (including "none") for all information requests, and only to one of two request orders. A simple improvement would be to use different justifications for demographics and context requests. Although we have not tested this option, Figure 2 suggests that some improvements in disclosure can be attained in this way (for males with high disclosure tendencies, for instance, one could request demographics first with the 'useful for you' justification, and then request context data with the 'useful for others' justification). A fully adaptive system could go much further though: it could select the optimal justification for each specific requested item, and order the requests dynamically to ensure the highest possible disclosure (without reducing satisfaction, of course). Our results suggest that such a fully adaptive system is viable, but one should note that a lot of data is needed to optimize its performance. Although this is not feasible in an academic research setting, a commercial system will likely provide enough data to implement this adaptive approach.

Participants in our study tested the information disclosure part of a mobile app recommender system. It will still need to be verified whether different types of systems may result in different user behaviors, and whether moving from a mockup study to real-world systems bears any differences. Moreover, different types of justifications, such as privacy protection reassurances (which play more on the risk side of the risk-benefit trade-off) may work better than the justifications tested in our study (e.g. [8]). Despite these limitations, we feel confident that a careful consideration of user characteristics and optimization goals can significantly improve strategies for helping users with information disclosure decisions.

ACKNOWLEDGEMENTS

This work has been supported through NSF grants IIS-0308277, IIS-0808783 and CNS-0831526, and by the Samsung R&D Research Center.

REFERENCES

1. Kobsa, A. Tailoring Privacy to Uses' Needs. In: Bauer, M., Gmytrasiewicz, P.J., and Vassileva, J. (eds.) *UM 2001. LNAI, vol. 2109*, pp. 303--313. Springer, Heidelberg (2001). DOI: 10.1007/3-540-44566-8_52
2. Kobsa, A. Privacy-Enhanced Web Personalization. In: Brusilovsky, P., Kobsa, A., and Nejdl, W. (eds.) *The Adaptive Web: Methods and Strategies of Web Personalization*. pp. 628--670. Springer, Heidelberg (2007). DOI: 10.1007/978-3-540-72079-9_21
3. Assad, M., Carmichael, D., Kay, J., Kummerfeld, B. Giving Users Control Over Location Privacy. In: *Workshop on Ubicomp Privacy*. (2007).
4. Teltzrow, M., Kobsa, A. Impacts of User Privacy Preferences on Personalized Systems: a Comparative Study. In: Karat, C.-M., Blom, J., and Karat, J. (eds.) *Designing Personalized User Experiences for eCommerce*. pp. 315--332. Kluwer, Dordrecht (2004). DOI: 10.1007/1-4020-2148-8_17
5. Culnan, M.J. and Armstrong, P.K. Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation. *Organization Science*. 10(1), 104--115 (1999). DOI: 10.1287/orsc.10.1.104
6. Chellappa, R.K., Sin, R.G. Personalization Versus Privacy: An Empirical Examination of the Online Consumer's Dilemma. *Information Technology and Management*. 6, 181--202 (2005). DOI: 10.1007/s10799-005-5879-y
7. Brodie, C., Karat, C.-M., Karat, J. Creating an E-Commerce Environment Where Consumers are Willing to Share Personal Information. In: Karat, C.-M., Blom, J.O., and Karat, J. (eds.) *Designing Personalized User Experience for eCommerce*. pp. 185--206. Kluwer, Dordrecht (2004). DOI: 10.1007/1-4020-2148-8_11
8. Kobsa, A., Teltzrow, M. Contextualized Communication of Privacy Practices and Personalization Benefits: Impacts on Users' Data Sharing Behavior. In: Martin, D. and Serjantov, A. (eds.) *PET 2005*. 329--343.

Springer, Heidelberg (2005). DOI: 10.1007/11423409_21

9. Besmer, A., Watson, J., Lipford, H.R. The Impact of Social Navigation on Privacy Policy Configuration. In: *6th Symposium on Usable Privacy and Security*. ACM Press, New York (2010). DOI: 10.1145/1837110.1837120
10. Patil, S., Page, X., Kobsa, A. With a Little Help from My Friends. In: *ACM 2011 Conference on Computer Supported Cooperative Work*. pp. 391--394. ACM Press, New York (2011). DOI: 10.1145/1958824.1958885
11. Consolvo, S., Smith, I., Matthews, T., LaMarca, A., Tabert, J. and Powledge, P. Location disclosure to social relations: why, when, & what people want to share. In: *ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 81–90. ACM Press, New York (2005). DOI: 10.1145/1054972.1054985
12. Wang, W., Benbasat, I. Recommendation Agents for Electronic Commerce: Effects of Explanation Facilities on Trusting Beliefs. *Journal of Management Information Systems*, 23, 217--246 (2007). DOI: 10.2753/MIS0742-1222230410
13. Acquisti, A., John, L.K., Loewenstein, G. The Impact of Relative Standards on the Propensity to Disclose. *Journal of Marketing Research*. (2011). DOI: 10.1509/jmr.09.0215
14. Privacy and American Business: Consumer Privacy Attitudes: A Major Shift Since 2000 and Why. *Privacy and American Business Newsletter*, 10, (2003).
15. Cho, H., Rivera-Sánchez, M., Lim, S.S. A Multinational Study on Online Privacy: Global Concerns and Local Responses. *New Media & Society*, 11, 395--416 (2009). DOI: 10.1177/1461444808101618
16. Li, N., Chen, G. Sharing Location in Online Social Networks. *IEEE Network*, 24, 20--25 (2010). DOI: 10.1109/MNET.2010.5578914
17. Smith, H.J., Milberg, S.J., Burke, S.J. Information Privacy: Measuring Individuals' Concerns About Organizational Practices. *MIS Quarterly*, 20, 167--196 (1996).
18. Malhotra, N.K., Kim, S.S., Agarwal, J. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research*, 15, 336--355 (2004). DOI: 10.1287/isre.1040.0032
19. Davidsson, C., Moritz, S. Utilizing Implicit Feedback and Context to Recommend Mobile Applications from First Use. *IUI 2011 Workshop on Context-awareness in Retrieval and Recommendation*. pp. 19--22. ACM, New York (2011). DOI: 10.1145/1961634.1961639
20. Kittur, A., Chi, E.H., Suh, B. Crowdsourcing user studies with Mechanical Turk. *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*. pp. 453--456. ACM press, New York (2008). DOI: 10.1145/1357054.1357127
21. Knijnenburg, B.P., Kobsa, A., Saldamli, G. Privacy in Mobile Personalized Systems: The Effect of Disclosure Justifications. In: *SOUPS Workshop on Usable Privacy & Security in Mobile Devices*. (2012). URL: http://cups.cs.cmu.ecu/soups/2012/u-prism/soups12-mobile-final1.pdf
22. Knijnenburg, B.P., Willemsen, M.C., Gantner, Z., Soncu, H., Newell, C.: Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction*. 22, (2012). DOI: 10.1007/s11257-011-9118-4
23. Spiekermann, S., Grossklags, J.: E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus actual Behavior. In: *3rd ACM Conference on Electronic Commerce*. pp. 38--47. ACM Press, New York (2001). DOI: 10.1145/501158.501163
24. Tsai, J.Y., Egelman, S., Cranor, L.F., Acquisti, A.: The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research*. 21, (2010). DOI: 10.1287/isre.1090.0260
25. Laufer, R.S., Wolfe, M.: Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory. *Journal of Social Issues*. 33(3), 22--42 (1977). DOI: 10.1111/j.1540-4560.1977.tb01880.x
26. Knijnenburg, B.P., Kobsa, A., Jin, H.: Dimensionality of Information Disclosure Behavior. *Submitted to the International Journal of Human-Computer Studies, special issue on Privacy Methodologies in HCI*. http://bit.ly/privdim
27. Knijnenburg, B.P., Kobsa, A.: Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems. *Accepted to the ACM Transactions on Interactive Intelligent Systems, special issue on Human Decision Making and Recommender Systems*. Available as ISR tech report: http://www.isr.uci.edu/tech_reports/UCI-ISR-12-1.pdf