

Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems

BART P. KNIJNENBURG, University of California, Irvine
ALFRED KOBSA, University of California, Irvine

Recommender systems increasingly use contextual and demographical data as a basis for recommendations. Users, however, often feel uncomfortable providing such information. In a privacy-minded design of recommenders, users are free to decide for themselves what data they want to disclose about themselves. But this decision is often complex and burdensome, because the consequences of disclosing personal information are uncertain or even unknown. Although a number of researchers have tried to analyze and facilitate such information disclosure decisions, their research results are fragmented, and they often do not hold up well across studies. This article describes a unified approach to privacy decision research that describes the cognitive processes involved in users' "privacy calculus" in terms of system-related perceptions and experiences that act as mediating factors to information disclosure. The approach is applied in an online experiment with 493 participants using a mock-up of a context-aware recommender system. Analyzing the results with a structural linear model, we demonstrate that personal privacy concerns and disclosure justification messages affect the perception of and experience with a system, which in turn drive information disclosure decisions. Overall, disclosure justification messages do not increase disclosure. Although they are perceived to be valuable, they decrease users' trust and satisfaction. Another result is that manipulating the order of the requests increases the disclosure of items requested early but decreases the disclosure of items requested later.

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, Measurement, Theory

Additional Key Words and Phrases: Privacy, information disclosure, decision-making, recommender systems, user experience

ACM Reference Format:

Knijnenburg, B. P., and Kobsa, A. 2013. Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems. *ACM Transactions on Interactive Intelligent Systems*. X, X, Article XX (Month 2013), X pages.
DOI:<http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

While traditional recommender systems are commonly trained through users' feedback on the recommended items, recommenders (and specifically mobile recommenders) increasingly also employ users' demographical and contextual data. Those data allow them to instantly generate recommendations that are relevant to the user and to the situation in which the recommender system is being used

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

Author's addresses: B. P. Knijnenburg and A. Kobsa, Donald Bren School of Information and Computer Sciences, University of California, Irvine. Email: {bart.k, kobsa}@uci.edu.

Permission to make digital or hardcopies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credits permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2010 ACM 1539-9087/2010/03-ARTXX \$15.00

DOI:<http://dx.doi.org/10.1145/0000000.0000000>

[Adomavicius and Tuzhilin 2011]. Privacy research however indicates that quite a few people feel uncomfortable disclosing demographical data [Ackerman et al. 1999], and that they dislike being ‘tracked’ for the purpose of gathering contextual data [Turow et al. 2009; Xu et al. 2009].

One approach to this problem is to create a privacy-preserving system architecture that can compute recommendations without explicitly knowing the users’ input data [Canny 2002b; 2002a; Polat and Du 2005a; 2005b]. However, this disregards the fact that users’ perception of the potential privacy threats may differ from the actual threats [John et al. 2011]. Another remedy is to give users explicit control over what information they disclose [Wenning and Schunter 2006; Kolter and Pernul 2009]. Information disclosure then becomes an explicit decision, in which users have to make a trade-off between the potential benefits of disclosure and the possibly ensuing privacy risks [Mabley 2000; Chellappa and Sin 2005; Taylor et al. 2009].

Decision-making is an inherently complex problem though, especially when the outcomes are uncertain or unknown [Kahneman and Tversky 1979; Kahneman et al. 1982; Gigerenzer and Goldstein 1996]. In the field of privacy, this complex decision process has been aptly dubbed “privacy calculus” [Culnan 1993; Laufer and Wolfe 1977]. When users have to decide whether or not to disclose personal information to a recommender system, they typically know little about the positive and negative consequences of disclosure [Acquisti and Grossklags 2005; Acquisti and Grossklags 2008].

Another problem is that users’ information disclosure decisions are highly dependent on the context [Lederer et al. 2003; Li et al. 2010; Nissenbaum 2010; John et al. 2011]. Researchers have looked at various techniques to assist or influence users in such decisions, such as reordering the disclosure requests to increase disclosure [Acquisti et al. 2011], providing justifications for disclosing (or not disclosing) certain information [Kobsa and Teltzrow 2005; Besmer et al. 2010; Patil et al. 2011; Acquisti et al. 2011], or displaying privacy seals or statements [Rifon et al. 2005; Hui et al. 2007; Egelman et al. 2009; Xu et al. 2009]. While these studies yielded interesting and occasionally even counterintuitive results, those results are mostly quite isolated. For instance, some research focuses on increasing disclosure behavior, but disregards users’ perception of the system and their satisfaction with the experience of using it (see section 2.1). Others study users’ general privacy concerns, but disregard their impact on disclosure behavior (see section 2.2). Research relevant to privacy-related decision-making is scattered across several disparate thrusts, including research on increasing information disclosure, research on user perception and satisfaction (also called ‘user experience’), and research on privacy concerns as personal traits.

To make relevant and robust contributions, research on users’ reluctance to disclose personal data to context-based recommender systems should forge the divergent contributions into a unified approach. By incorporating system-related perceptions and experiences as mediators to information disclosure behavior, such an approach can provide insights into the cognitive processes involved in users’ privacy calculus, and explain how suggested system improvements as well as personal privacy concerns impact information disclosure decisions. This paper develops such an encompassing approach (section 2) and applies it to the analysis of an online user experiment with a mockup of a mobile app recommender system (section 3). Section 4 reflects on the results of this experiment and integrates them with qualitative findings from an interview study. Section 5 finally provides conclusions and suggestions for future research.

2. RELATED WORK

Existing approaches to privacy decision-making are scattered across a number of sub-fields, each of which studies only part of the problem. This section covers these approaches, and identifies potential synergies between them that can address their respective shortcomings. The purpose of the section is not to provide an exhaustive treatment of the entire body of privacy-related research (see [Solove 2006; Iachello and Hong 2007; Kobsa 2007; Smith et al. 2011; Bélanger and Crossler 2011] for more comprehensive surveys), but to provide a foundation for the subsequent discussion of our study and of future research in the field of privacy decision-making.

2.1 Research on Privacy as a Personal Trait

Arguably the first attempt to measure privacy as a personal trait was the Equifax survey by Westin and Harris & Associates [1981]. They use three core items across several surveys to classify people into three broad categories: privacy fundamentalists, pragmatists, and unconcerned [Harris et al. 1998; 2003].

Notwithstanding the simplicity of Westin's approach, most researchers agree that privacy is in fact a multi-dimensional concept [Laufer et al. 1973]. In this light, Culnan [1993] used the three items from the Equifax survey [Harris and Associates 1990; 1991] and two items from Smith et al. [1992] to construct two dimensions of concern for privacy: loss of control, and unauthorized secondary use of personal information. Smith et al. [1996] extended and refined this scale, resulting in the Concern For Information Privacy (CFIP) scale. The CFIP scale consists of fifteen items measuring four correlated factors: collection concerns, unauthorized access, fear of accidental errors, and secondary use. Smith et al. go at great lengths to validate the internal consistency and validity of the CFIP scale. However, Stewart and Segars [2002] demonstrate that CFIP can be more parsimoniously represented as a higher-order factor with the four sub-factors as indicators.

Malhotra et al. [2004] provide a different take on the CFIP scale, adapting it to an Internet environment. They produce two scales: a 6-item scale for General Information Privacy Concern (GIPC, partially adapted from Smith et al. [1992]), and an Internet Users Information Privacy Concern (IUIPC) scale with ten items measuring three factors: collection (adapted from Smith et al. [1992]), control (newly developed), and awareness (newly developed). Malhotra et al. claim that IUIPC is superior to CFIP because it has fewer factors, a better internal fit, a stronger relation to GIPC, and a slightly better fitting statistical model. Moreover, because it is based on social contract theory, it is also easily extensible to new types of information privacy. For instance, Buchanan et al. [2007] link IUIPC to more specific concerns and protection behaviors related to modern privacy-sensitive technologies (e.g. e-mail, e-banking), and Zhang et al. [2011] adapt IUIPC to Facebook privacy. In light of our goal of studying privacy in terms of information disclosure decisions, the control factor is the most interesting contribution of IUIPC, because people who desire to have control over their privacy may actually be relatively *more* willing to disclose information, as long as the decision is theirs to make [Nowak and Phelps 1995].

Malhotra et al. [2004] construct a structural model, linking their IUIPC scale to behavioral intentions via trusting beliefs and risk beliefs (taken from [Jarvenpaa and Tractinsky 1999]) as mediating concepts. Li et al. [2011] do the same for GIPC, and additionally show how emotions and cognitions influence this process. However, the aforementioned privacy scales do not explicitly consider information disclosure as a *decision* with inherent trade-offs of threats versus benefits. Likewise, these studies do not explicitly try to manipulate users' disclosure behavior.

2.2 Information Disclosure Research

Information disclosure research investigates the factors that may influence how much information users disclose. Information disclosure is seen as a *decision problem*, in which the decision-maker has to trade off several uncertain (and sometimes unknown) consequences. The fundamental finding of decision-making research, namely that humans typically do not follow rational economical principles in their decision processes, has been shown to also apply to disclosure decisions [Acquisti and Grossklags 2008].

One of these non-rational influences in decision-making is the ‘endowment effect’: people are usually less willing to give up something they already have than they are willing to pay for acquiring something they do not have [Thaler 1980; Kahneman et al. 1990]. Both Acquisti et al. [2009] and Tsai et al. [2010] show that people are indeed less willing to pay for gaining privacy than what they would demand to give it up. This may be the main reason why explicit monetary rewards seem to have varying effects on disclosure. Hui et al. [2007] find that participants are proportionally more willing to fill out a marketing survey with increasing monetary rewards ranging from \$0.60 to \$5.40. In a study on a location-based coupon service, Xu et al. [2009] find that a rebate of \$0.20 on the monthly phone bill increases disclosure only when the system pushes the coupons to the user. However, when studying information disclosure in an online fax service, Li et al. [2010] find an “undermining effect of rewards” (p. 21) when users do not perceive the requested information to be relevant to the purpose of the e-commerce transaction. It has no effect when the information is perceived as relevant to begin with.

A more subtle strategy to influence disclosure is to change the order of disclosure requests. Acquisti et al. [2011] demonstrated that people disclose less information when requests are made in increasing order of intrusiveness (compared to a random order). This effect is particularly pronounced for more intrusive questions: asking those upfront significantly increases their likelihood of being answered. Arguably, people become more wary of disclosing very personal information as the disclosed information accumulates; the most relevant information should thus be requested upfront. Acquisti et al. did not consider subjective evaluations of the decision process. It is thus unclear whether their manipulation resulted in people feeling ‘tricked’ into disclosing more information than they would have liked.

A somewhat more explicit strategy to improve disclosure is to provide justifications for disclosing the information. Such justifications include providing a reason for requesting the information [Consolvo et al. 2005], the benefits of disclosure [Kobsa and Teltzrow 2005; Wang and Benbasat 2007], and appealing to the social norm [Besmer et al. 2010; Patil et al. 2011; Acquisti et al. 2011]. The effect of such justifications seems to vary. In the study of Kobsa and Teltzrow [2005], users were about 8.3% more likely to disclose information when they knew the benefits of disclosing the information. In an experiment by Acquisti et al. [2011], they were about 27% more likely to do this when they learned that many others decided to disclose the same information. However, Besmer et al. [2010] 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. Similarly, Patil et al. [2011] rate social navigation cues as a secondary effect.

Another strategy is to provide a privacy indicator, statement or seal. Egelman et al. [2009] show that privacy indicators next to search results can entice users to pay a premium to vendors with higher privacy scores. In their study, participants paid

about \$0.15 extra for a pack of batteries and about \$0.40 for a sex toy (on top of a \$15.50 average base price). Users of Xu et al.'s [2009] location-based coupon service were more likely to disclose information when the site displayed either a TRUSTe seal or a legal statement, with the seal working best. In Hui et al.'s [2007] marketing survey, however, privacy statements only had a marginal effect, and a privacy seal did not significantly increase disclosure. Studying an online CD retailer, Metzger [2006] found that neither seal nor policy had an effect. Rifon et al. [2005] show that warnings and seals at an online retailer website influence users in certain situations only.

John et al. [2011] demonstrate that compared to an unofficial and unprofessional looking site, a professional looking site garners *higher* privacy concerns, because its design reminds users of privacy. While most likely being more risky to entrust one's information with, the unofficial-looking site downplays privacy concerns and thus increases disclosure. If even a professional-looking site can instill privacy concerns, it seems plausible that any reference to privacy may inadvertently prime users to become more concerned about it. This hypothesized phenomenon may explain the seemingly disappointing effects of justifications, seals and statements, as they inadvertently remind users of the concept of privacy. In this light, it seems important to consider users' *perceptions* of the privacy threat and of the value of the help offered by the system as important mediators of any effects on disclosure behavior.

Similarly, privacy desire as a personal trait (as discussed in section 2.1) is a surprisingly bad predictor of disclosure behavior [Spiekermann et al. 2001; Metzger 2006; van de Garde-Perik et al. 2008]. Presumably, people's information disclosure decisions are more strongly dependent on the context in which they are made [John et al. 2011]. Indeed, as suggested by Hui et al. [2006] and Xu et al. [2009], disclosure is a trade-off between experienced system-specific concerns and experienced system-specific benefits. *Experiential evaluations* of the system may thus be another important mediator of any effect on behavior.

2.3 User Experience Research

Strategies aimed at influencing users' disclosure behavior may have unforeseen effects on their perceptions and experiences, and these effects could arguably cancel out or even negate the intended effects on disclosure (cf. [Archer and Berg 1978; Fitzsimons and Lehmann 2004]). Likewise, users' perceptions and experiences may mediate the effect of users' personal privacy preferences on disclosure. However, only few researchers on information disclosure consider such perceptual and experiential aspects of the systems they evaluate (Hui et al. [2006] and Xu et al. [2009] are notable exceptions). Explicitly measuring those mediating concepts may strengthen the link between personal privacy preferences and disclosure behavior.

User experience research typically takes perceptual and experiential aspects into account [Hassenzahl 2005]. In the field of recommender systems, Knijnenburg et al. [2012] and Pu et al. [2011] developed and validated frameworks for user experience research. Their frameworks show considerable overlap, both describing how perceptions and experiences influence user behavior. Knijnenburg et al. additionally describe how these constructs mediate the effect of objective system aspects (e.g. the strategy to influence disclosure), and also consider personal and situational characteristics, which include personal privacy concerns.

Fig. 1 shows how in the Knijnenburg et al. [2012] framework, user behavior (or interaction; INT) is related to users' evaluation of the interaction with the system (or experience; EXP). The effect of any objective system aspects (OSA) on the experience

and interaction is (at least partially) mediated by users' perceptions of the system aspects (or subjective system aspects; SSA). Not only the system, but also personal and situational characteristics (PC and SC) can influence the perceptions, experience, and interaction. The Knijnenburg et al. framework consists of the higher-level concepts displayed in Fig. 1, as well as a number of operationalized factors for each of these concepts. However, since the framework has mainly been used to study the overall user experience of recommender systems, our current study will not use these lower-level factors (with the exception of system satisfaction). Rather, we will use the higher-level concepts, to integrate aspects considered in research on information disclosure and on privacy as a personal trait. These aspects related to privacy and information disclosure are listed as bullet points in Fig. 1. The specific nature of their integration will be described in the next section.

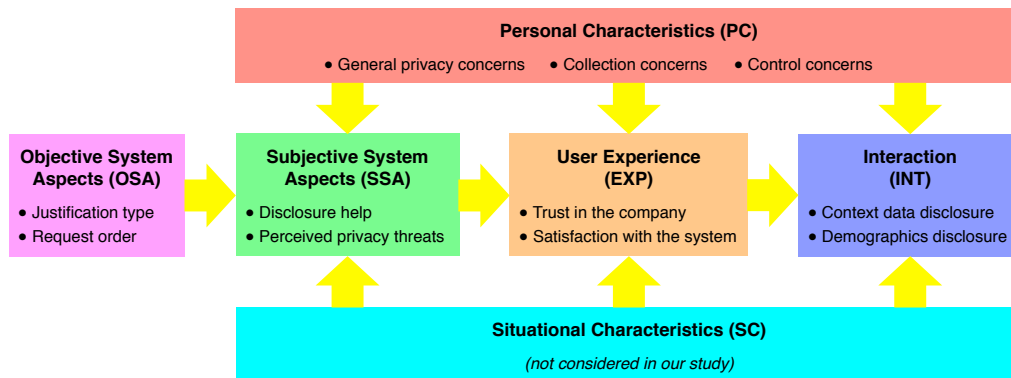


Fig. 1. The framework for user-centric evaluation of recommender systems (Knijnenburg et al. [2012]), populated with the concepts used in this current study (bullet points).

3. INTEGRATING EXISTING APPROACHES TO PRIVACY

As mentioned before, existing work on privacy decision-making faces two main handicaps. The first problem is that research typically either looks at privacy as a personal trait or at factors that influence disclosure behavior, rarely comparing the relative importance of their respective impacts. The second problem is that the influence of both personal traits and system characteristics on information disclosure varies extensively from system to system. The main contribution of the current paper is a unified approach to studying privacy decisions, which integrates privacy-related concepts into the Knijnenburg et al. [2012] framework. This allows us to remedy both problems.

First, using the Knijnenburg et al. framework, we can integrate the effects of privacy as a personal trait (PC) and of design characteristics (OSA) on information disclosure decisions (INT) in a single model. Specifically, we amend the Knijnenburg et al. framework with three personal characteristics (PC): ‘general privacy concerns’, ‘collection concerns’ and ‘control concerns’ (cf. [Smith et al. 1996; Malhotra et al. 2004]). As for strategies to influence disclosure (objective system aspects, OSA), we consider two previously investigated strategies: justification messages (cf. [Kobsa and Teltzrow 2005; Besmer et al. 2010; Patil et al. 2011; Acquisti et al. 2011]) and request order (cf. [Acquisti et al. 2011]). Integrating these PCs and OSAs allows us to evaluate the relative contribution of each, thereby solving the first problem.

Second, the Knijnenburg et al. framework allows us to describe the aforementioned effects as mediated by system-specific perceptions (SSA) and experiences (EXP). Specifically, we consider the subjective system aspects (SSA) ‘perceived privacy threats’ (cf. [Xu et al. 2009; Xu et al. 2011]) and ‘perceived value of disclosure help’ (cf. [Wang and Benbasat 2007]) as well as the experience (EXP) variables ‘trust in the company’ (cf. [Jarvenpaa and Tractinsky 1999; Metzger 2004]) and ‘self-anticipated satisfaction with the system’ (cf. [Xu et al. 2009; Xu et al. 2011; Hui et al. 2006]). These mediating concepts may increase the robustness of the link between information disclosure behavior and its presumed antecedents, and in the absence of an effect, they may explain why the strategy or personal trait did not influence disclosure as expected. Any inconsistencies with existing work can thus be explained in terms of these mediating variables. Our work thus takes an important step towards solving the second problem; authors of future work can adapt our integrated approach to further increase the comparability of disparate privacy research efforts.

This paper is not the first to integrate several privacy aspects into a single structural model. Hui et al. [2006], Li et al. [2010; Li and Santhanam 2008], Xu et al. [2009; 2011], and Keith et al. [2011] provide similar models which inspired our work. However, despite the more comprehensive nature of their approach, their work fails to truly investigate privacy as a decision making process in adequate detail, because their outcome measure is a more generic form of behavioral intention (i.e. measured with generic questionnaire items such as “How likely would you provide your personal information (including your location) to use the M-Coupon service?”). Such intentions arguably do not directly relate to observable privacy behaviors (cf. [Spiekermann et al. 2001] who show that privacy preferences and actual behavior tend to be weakly related at best). Our approach, in contrast, considers users’ detailed privacy decisions (a yes/no decision for multiple disclosures), which is more compatible with existing information disclosure research (cf. [Acquisti et al. 2011]). In effect, our paper is arguably the first to answer the call by Smith et al. [2011] to integrate research on antecedents, privacy concerns and privacy calculus, with a focus on actual observable outcomes.

4. ONLINE EXPERIMENT

We validate our integrated approach using a mobile app recommender system, inspired by existing systems that have been developed both for research and commercial purposes (e.g., [Böhmer et al. 2010; Girardello and Michahelles 2010; Davidsson and Moritz 2011], *chomp.com*). The system recommends apps for Android phones based on users’ context (e.g. location, app usage, credit card purchases) and demographics (e.g. age, hobbies, religion, household income).

Although context has recently attracted the most attention in mobile recommender systems research [Ricci 2011; Adomavicius and Tuzhilin 2011], several researchers have explored the combination of context and demographics in mobile recommenders [Lee and Lee 2007; Lee and Park 2007; Oh and Moon 2012; Zheng et al. 2012]. In general, the users’ context provides a wealth of automatically accrued data that can be used to provide relevant recommendations tailored to the specific usage situation. Demographics, on the other hand, can be used to overcome the “new user problem” [Lee and Park 2007], and is typically easier to interpret than context data. Indeed, research shows that it is the combination of demographics and context that leads to the best recommendation quality: Lee and Lee [2007] show that there is added value of context over demographics, while Oh and Moon [2012] show the value

of demographics over context. Zheng et al. [2012] show context and demographics each add value to a standard collaborative filtering system.

4.1 System

The system we created for our experiment has the working title ‘Applause’. As the current experiment only considers the information disclosure aspect of the system, it uses a web-based mockup of the Applause system that collects personal information but does not make any recommendations. To increase the realism of the experiment, users were told that their data would be disclosed to the developer, a company named Appy¹. 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 (Fig. 2).

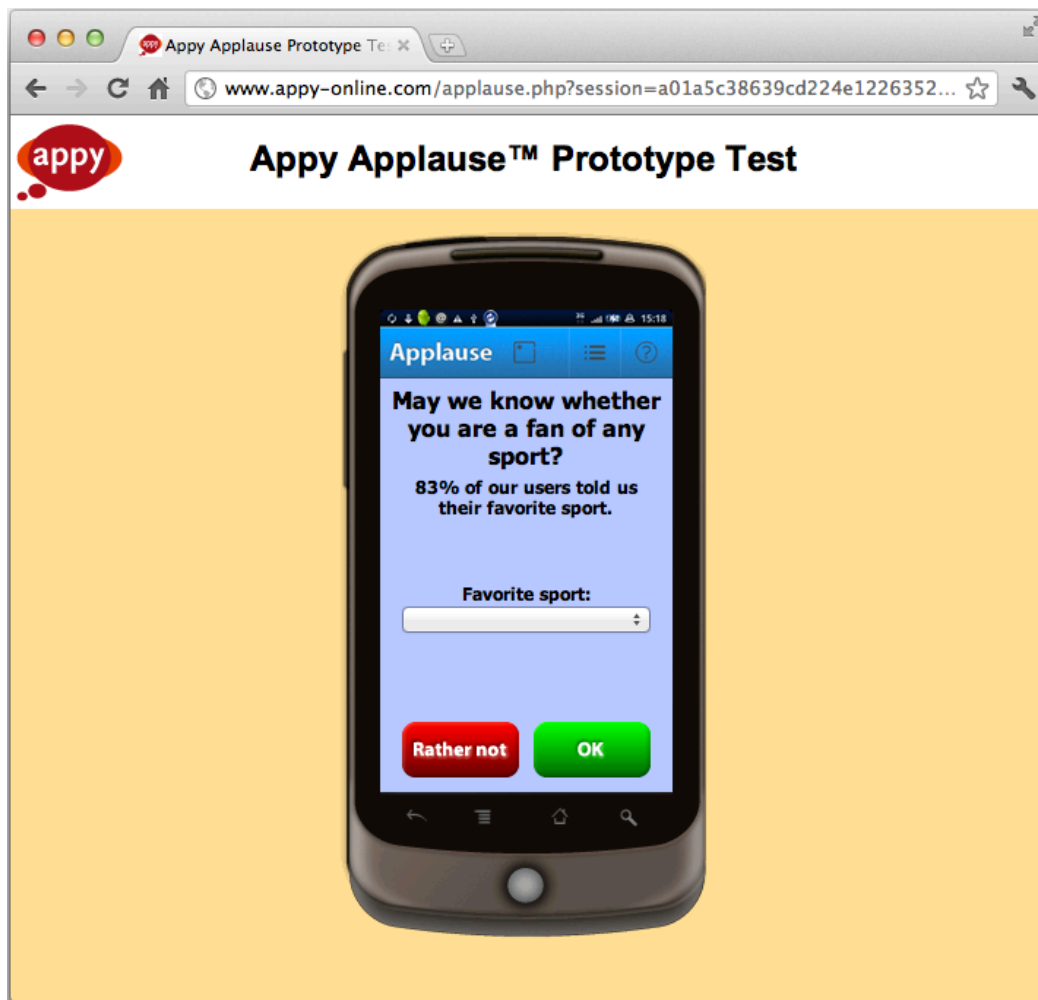


Fig. 2. The website of Appy with the Applause system mockup. On this website the participants do the disclosure part of the experiment

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

4.2 Setup

Participants were recruited between June 2011 and February 2012 in three rounds. We first enrolled 200 participants via Amazon Turk, a recruitment source that became very popular for conducting user studies [Kittur et al. 2008]. To improve the quality of our results, we only allowed participants from the United States, and asked a number of comprehension-testing questions. Moreover, we included several reverse-coded items in our questionnaires, and excluded participants who gave apparently inconsistent answers. In the second round we recruited an additional 52 participants via Craigslist.com to test for any anomalous differences between these two recruitment sources. No significant differences were found. Combined, these 252 participants formed our exploratory sample, on which different measurement models and structural models were tested to find an optimally fitting model. Finally, we recruited another 239 participants via Amazon Turk as a confirmatory sample. The optimal model found in the exploration phase was tested on this set, and any inconsistent effects were removed from the model. This “split-half” method increases the robustness of the model by removing unstable effects. However, we estimate the final model using the data from all 491 participants, because the full data set provides the most accurate estimates of the true effect sizes. The set of participants had an adequate distribution of gender (223 males, 266 females, 2 did not disclose) and age (ranging from 18 to older than 60, with 25-30 year-olds as the median age group).

Participants were first given a short introduction to the mobile app recommender, including two examples of how the system might use their data to provide context-aware and personalized recommendations. They were then informed that they would be helping Appy to test the information disclosure part of the system². After randomly assigning them to one of 5×2 conditions (see below), participants were ostensibly “transferred” to the Appy website, where they would make 31 information disclosure decisions on 12 pieces of context data and 19 pieces of demographical data. Context requests asked users to indicate whether they would disclose the respective data, and could be answered with a simple ‘yes’ or ‘no’. For demographics requests, users were asked to provide the actual information, or to decline disclosure. All decisions were logged to our database. After 31 decisions, participants were transferred back to the experimenters’ website, where they were asked about their personal privacy concerns and their subjective and experiential evaluation of the system.

4.3 Manipulations

The experiment introduces two strategies to influence information disclosure as between-subjects manipulations: the type of justification message (5 conditions) and the order in which disclosure requests are made (2 conditions). Although these strategies had been tested before (in different forms and different contexts), our unified approach may allow to measure the effect of these strategies more robustly, and to explain *why* their effects occur in terms of perceptions and experiences.

Four different types of justification messages are tested against the baseline system with no justification messages, bringing the total to five conditions (see Table I). The ‘useful for you’ and ‘useful for others’ justifications explain the benefits of disclosure (cf. [Wang and Benbasat 2007; Kobsa and Teltzrow 2005]) in two different ways. The ‘number of others’ justification appeals to the social norm (cf. [Besmer et

² To prevent any disappointment that might influence the study results, participants were explicitly told before and after the test that they would not be receiving any recommendations.

al. 2010; Patil et al. 2011; Acquisti et al. 2011]). The ‘explanation’ justification, which was added to the experiment after a preliminary interview study, gives the reason for requesting the information (cf. [Consolvo et al. 2005]).

Each user would see only one of the five conditions (either no justifications, or one of the four justification types). The percentages in the messages ‘useful for you’, ‘number of others’ and ‘useful for others’ were randomly chosen from 5% to 95% (in Knijnenburg and Kobsa [2013] we show that this percentage has barely any effect, and hence we disregard it in the current analysis).

Finally, since Acquisti et al. [2011] demonstrate that the request order may influence users’ disclosure decisions, we manipulate the order in which disclosure requests are made: demographical data first or context data first³ (see Table I).

Table I. Experimental manipulations: strategies to influence information disclosure

Manipulation	Conditions	Description
Justification type	None	[Baseline condition with no justifications]
	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 recommendations when they told us/let us...”
	Explanation	“We can recommend apps that are [reason for request]”
Request order	Demographical data first	The system first requested the 19 pieces of demographical data, then the 12 pieces of context data.
	Context data first	The system first requested the 12 pieces of context data, then the 19 pieces of demographical data.

4.4 Measures

The main dependent variable in our experiment is participants’ information disclosure decision. Table II shows the requested items and the percentage of participants disclosing this information. Participants seemed to view context data as generally more sensitive than demographical data.

We subjected the items to a Confirmatory Factor Analysis (CFA) with dichotomous indicators and a weighted least squares estimator, estimating two factors⁴: one for context data and one for demographical data. Items with a very high level of disclosure and items that showed very high residual correlations with some of the other items were not included in the analysis. The final factor model has 7 items for context data disclosure and 7 items for demographical data disclosure. Factor loadings of the included items are shown in Table II, as well as Cronbach’s alpha and

³ The main request order manipulation was ‘demographical data first’ versus ‘context data first’, and their respective items were therefore grouped together. Beside that, we grouped similar items within the demographical data requests as well. For example, we always requested ‘household income’, ‘household savings’ and ‘household debt’ in consecutive order. Because of this, the demographical data requests consist of four separate subgroups of requests (labeled ‘leisure’, ‘personal’, ‘family/income’, ‘cultural background’). We furthermore manipulated the order of these four subgroups of items, but this manipulation had no effect in our model (i.e. it did not have any significant arrows in the model of Fig. 3). We therefore disregard this manipulation in the further discussion.

⁴ We conducted a series of Exploratory Factor Analyses in order to select the right number of factors for this data. We compared the Bayesian Information Criterion (BIC) for a one-factor model (10316), a two-factor model (9174), a three-factor model (9213), a four-factor model (9351), and a five-factor model (9426). The two-factor model has the lowest BIC, which means that it has the best parsimony.

average variance extracted (AVE) for each factor. Values for both Cronbach's alpha and AVE are good, indicating convergent validity, and the square root of the AVE is higher than the factor correlation, indicating discriminant validity of the two factors. An introduction to CFA as applied in this paper can be found in the electronic appendix.

Table II. Items used to measure participants' disclosure behavior

Type of data	Items	Level of disclosure	Factor loading
Context Alpha: 0.79 AVE: 0.652 Factor correlation: 0.432	Recommendation browsing	87.0%	
	Location	84.8%	0.767
	Phone model	84.6%	0.659
	App usage	82.2%	0.749
	App usage time	73.2%	
	App usage location	67.1%	
	Accelerometer data	65.3%	
	Calendar data	62.9%	0.835
	Microphone	50.9%	
	Web browsing	48.3%	0.874
	E-mail messages	36.7%	0.940
	Credit card purchases	20.1%	0.796
	Demographics Alpha: 0.86 AVE: 0.784 Factor correlation: 0.432	Gender	94.9%
Amount of reading		93.5%	
Age		93.3%	
Education		92.7%	
News interests		92.7%	
Amount of TV watching		92.3%	
Population density of area		90.7%	
Workout routine		90.1%	
Children		89.3%	
Race		89.1%	
Relationship status		88.6%	0.911
Phone data plan		87.6%	0.905
Housing situation		87.4%	
Favorite sports (fan)		86.8%	0.718
Political preferences		86.4%	0.802
Field of work		83.6%	0.915
Household income		74.2%	0.964
Household savings		66.3%	0.957
Household debt	64.5%		

After completing the disclosure part of the experiment, participants were asked about their privacy concerns and their subjective evaluation of the system. Participants indicated on a 5-point scale their level of agreement with the items presented in Table III. We subjected the items to a factor analysis with ordered categorical indicators and a weighted least squares estimator, estimating 7 factors. Items with low factor loadings, high cross-loadings, or high residual correlations were removed from the analysis. Factor loadings of the included items are shown in Table III, as well as Cronbach's alpha and average variance extracted (AVE) for each factor. Values for AVE are good for all factors, indicating convergent validity. Values of Cronbach's alpha range from acceptable to excellent, with the exception of control concerns. This factor borrows some of its stability from correlation with other factors. The square root of the AVE is higher than the factor correlation for all factors except general privacy concerns and collection concerns, indicating that these factors could be collapsed. Since Malhotra et al. [2004] proposed these factors as distinct constructs, we keep them separate though.

Table III. Items used to measure participants' privacy concerns and subjective evaluations of the system

Considered aspects	Items	Factor loading
General privacy concerns (PC) Alpha: 0.76 AVE: 0.774 Based on [Malhotra et al. 2004; Smith et al. 1996]	All things considered, the Internet causes serious privacy problems	0.785
	Compared to others, I am more sensitive about the way online companies handle my personal information	0.708
	To me, it is the most important thing to keep my privacy intact from online companies	
	I believe other people are too concerned with online privacy issues	0.824
Collection concerns (PC) Alpha: 0.86 AVE: 0.815 Based on [Malhotra et al. 2004; Smith et al. 1996]	I am concerned about threats to my personal privacy today	0.860
	It usually bothers me when online companies ask me for personal information	
	When online companies ask me for personal information, I sometimes think twice before providing it	0.829
	It bothers me to give personal information to so many online companies	-0.749
	Online companies may collect any information about me because I have nothing to hide (new)	0.855
Control concerns (PC) Alpha: 0.58 AVE: 0.526 Based on [Malhotra et al. 2004]	I'm concerned that online companies are collecting too much personal information about me	-0.774
	I'm not bothered by data collection, because my personal information is publicly available anyway (new)	
	Online privacy is really a matter of consumers' right to exercise control and autonomy over decisions about how their information is collected, used, and shared	0.735
	Control of personal information lies at the heart of online privacy	0.715
	I do not want to think about who controls my personal information (new)	
Perceived value of disclosure help (SSA) Alpha: 0.75 AVE: 0.581 Inspired by [Wang and Benbasat 2007]	I believe that online privacy is invaded when control is lost or unwillingly reduced as a result of a marketing transaction	
	I do not feel the need to control my personal information (new)	
	The system ⁵ helped me to decide what information I should disclose	0.788
	The system explained how useful providing each piece of information was	0.633
Perceived privacy threats (SSA) Alpha: 0.71 AVE: 0.529 Inspired by [Xu et al. 2009; 2011]	The system helped me to make a tradeoff between privacy and usefulness	0.851
	I felt clueless about what information to disclose	
	The system made me more cautious than usual disclosing this type of information	
	The system helped me to protect my privacy	0.909
	The system has too much information about me	-0.608
The system does not know anything I would be uncomfortable sharing with it The system made me disclose several things that I normally would not disclose to an app like this I felt tricked into disclosing more information than I wanted		0.626

⁵ In our questions to participants, we consistently referred to the mockup of the app recommender (Fig. 2) as "the system".

Trust in the company (EXP) Alpha: 0.93 AVE: 0.845 Based on [Jarvenpaa and Tractinsky 1999; Metzger 2004])	I believe the company providing this software is trustworthy in handling my information	0.927
	I believe this company tells the truth and fulfills promises related to the information I provide	0.917
	I believe this company is predictable and consistent regarding the usage of my information	0.886
	I believe this company is honest when it comes to using the information I provide	0.945
	I think it is risky to give my information to this company	
	There is too much uncertainty associated with giving my information to this company	
(Self-anticipated)⁶ satisfaction with the system (EXP) Alpha: 0.91 AVE: 0.722 Based on [Knijnenburg et al. 2012] and inspired by [Hui et al. 2006; Xu et al. 2009; 2011]	Providing this company my information would involve many unexpected problems	
	I feel safe giving my information to this company	
	The system has no real benefit to me	
	Using the system is annoying	-0.811
	The system is useful	0.885
	Using the system is a pleasant experience	
	Using the system makes me happy	0.841
	Overall, I am satisfied with the system	0.923
	I would recommend the system to others	0.870
	I would use this system if it were available	
	I would pay \$2 to use this system	
I would quickly abandon using this system	-0.759	
It would take a lot of convincing for me to use this system		

4.5 Results

We subsequently subjected the disclosure behaviors, the subjective evaluations, and the manipulated system aspects to Structural Equation Modeling (SEM), which simultaneously fits the measurement model and the structural relations between measured constructs. A structural equation model can be seen as a series of linear regressions that together describe a path of effects between the behaviors, evaluations, and manipulations. Essentially, in the graphical representation of our structural equation model (Fig. 3) each set of incoming arrows can be seen as a linear regression between manipulations (rectangles) and factors (ovals)⁷. Numbers on the arrows (as well as their thickness) represent the β coefficients (and standard error) of the regression effect represented by the arrow. Factors are scaled to have a standard deviation (SD) of 1. For any arrow $A \rightarrow B$, a 1 SD difference in A thus causes a β SD difference in B. The χ^2 values test the effect of all justifications simultaneously; the β coefficients below the χ^2 values represent the effect (in β SD difference) of each justification tested against the baseline of ‘no justification’. The β coefficients on the ‘demographics first’ arrow represent the effect of the ‘demographics first’ request order (in β SD difference) compared to the ‘tracking first’ request order. For a more detailed introduction to SEM as applied in this paper, see the electronic appendix.

⁶ As explained, users did not use the actual system, but only its information disclosure part. We asked users to evaluate their satisfaction with the system and to thereby assume that it would provide actual app recommendations as users were shown in the introduction. These items thus relate to the users’ own *self-anticipation* of their satisfaction.

⁷ The indicators of the factors are left out for clarity; Table II and Table III explain how factors are measured.

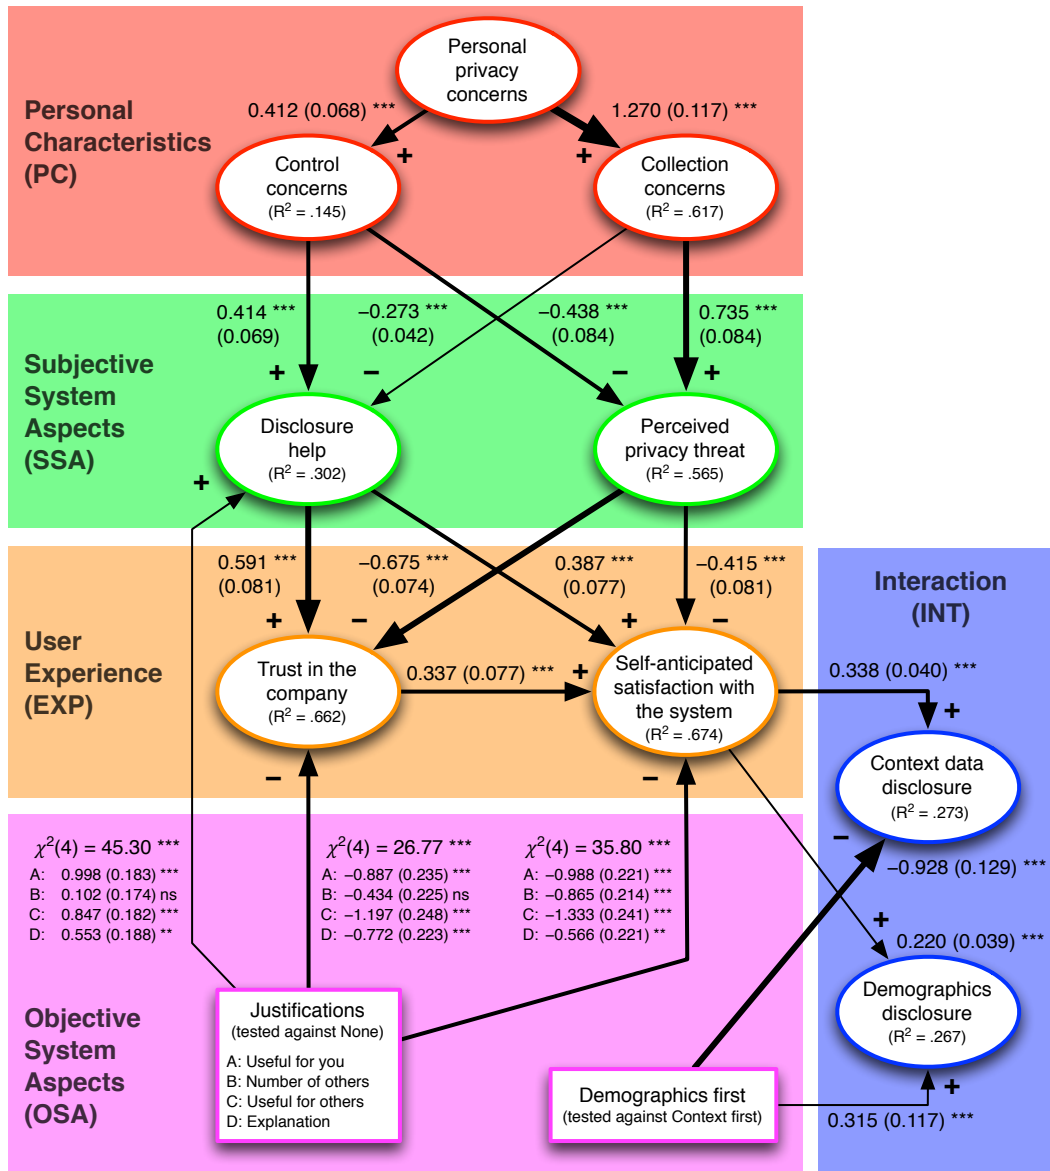


Fig. 3. The Structural Equation Model (SEM) for the data of the experiment. The model shows the subjective factors behind users' information disclosure decisions when using a recommender system, and the effect of request order and different types of justifications (Significance levels: *** $p < .001$, ** $p < .01$, 'ns' $p > .05$)

To avoid over-fitting, we constructed the structural equation model on our exploratory sample, tested it on the confirmatory sample, and then pruned any effects that were not consistently significant between the two samples. We then fitted the resulting model to the entire dataset. This procedure gives us additional confidence that the effects in the model extend beyond our current sample. Moreover, because our model is based on the validated user experience framework of Knijnenburg et al. [2012], we assert that the findings hold true beyond the context of the current experiment, and more generally describe users' information disclosure decisions when using a recommender system. The final model (Fig. 3) has a good

model fit⁸: $\chi^2(912) = 1540$, $p < .001$; $RMSEA = 0.037$, 90% CI: [0.034, 0.041], $CFI = 0.977$, $TLI = 0.976$.

The model shows that aside from the request order, all effects on context and demographics disclosure are mediated by users' self-anticipated satisfaction with the system (request order, satisfaction \rightarrow context data, demographics disclosure). The higher their satisfaction, the more inclined users are to disclose information. Satisfaction is higher for participants who trust the company, feel helped in their disclosure, and perceive a low level of privacy threat (trust, help, threat \rightarrow satisfaction). In other words, privacy-related aspects have a significant influence on the users' satisfaction with the system⁹. Trust in the company itself is also higher for participants that feel helped in their disclosure and perceive a low level of privacy threat (help, threat \rightarrow trust), indicating that the effects of privacy extend beyond the system to the company providing the system.

In terms of personal characteristics, general personal privacy concerns drive control and collection concerns (general privacy concerns \rightarrow control, collection concerns), but these concerns have the opposite influence on perception of disclosure help and privacy threat: people's collection concerns cause a decrease in the perceived value of disclosure help and an increase in perceived threat, but controlling for disclosure concerns, people's control concerns cause an increase in perceived value of disclosure help and a decrease in perceived threat (control, collection concerns \rightarrow help, threat). Compared to other systems, our system provides rather detailed control over users' information disclosure, which may have caused this effect.

The justifications have a significant impact on perception of disclosure help, trust in the company, and self-anticipated satisfaction with the system (justification type \rightarrow help, trust, satisfaction). The 'useful for you', 'useful for others' and 'explanation' justifications each significantly increase the perceived value of disclosure help. However, this positive effect is canceled out by a negative effect on trust in the company and on self-anticipated satisfaction with the system. Fig. 4 shows that the total (direct plus mediated) effects of the justifications on trust in the company are essentially zero¹⁰ (i.e. the justifications fall within 2 standard errors of the baseline ('None') except for the 'useful for others' justification), and that the total effects on self-anticipated satisfaction with the system and disclosure behavior are negative (i.e., most justifications have a disclosure rate that is more than 2 standard errors lower than the baseline). In other words, the justifications do not work; in fact they rather *decrease* users' satisfaction.

Finally, the request order has a direct impact on disclosure behavior. Basically, disclosure of each type of data is higher when it is requested first. Requesting demographics first increases demographics disclosure but decreases context disclosure, and vice versa. The effect is stronger on context disclosure though ($\beta = 0.315$ vs. $\beta = -0.928$), which are the more sensitive data.

⁸ A good model is not statistically different from the fully specified model ($p > .05$). However, this statistic is commonly regarded as too sensitive, and researchers have therefore proposed other fit indices [Bentler and Bonett 1980]. Based on extensive simulations, Hu and Bentler [1999] propose cut-off values for other fit indices to be: $CFI > .96$, $TLI > .95$, and $RMSEA < .05$, with the upper bound of its 90% CI falling below 0.10.

⁹ Note that this finding may in part be caused by the fact that participants only tested the information disclosure part of the system. Still, the effects are large enough (they explain a hefty 67.4% of the total variance in self-anticipated satisfaction) to assert that privacy is an important aspect of user satisfaction.

¹⁰ See the electronic appendix for an example of how to calculate total effects.

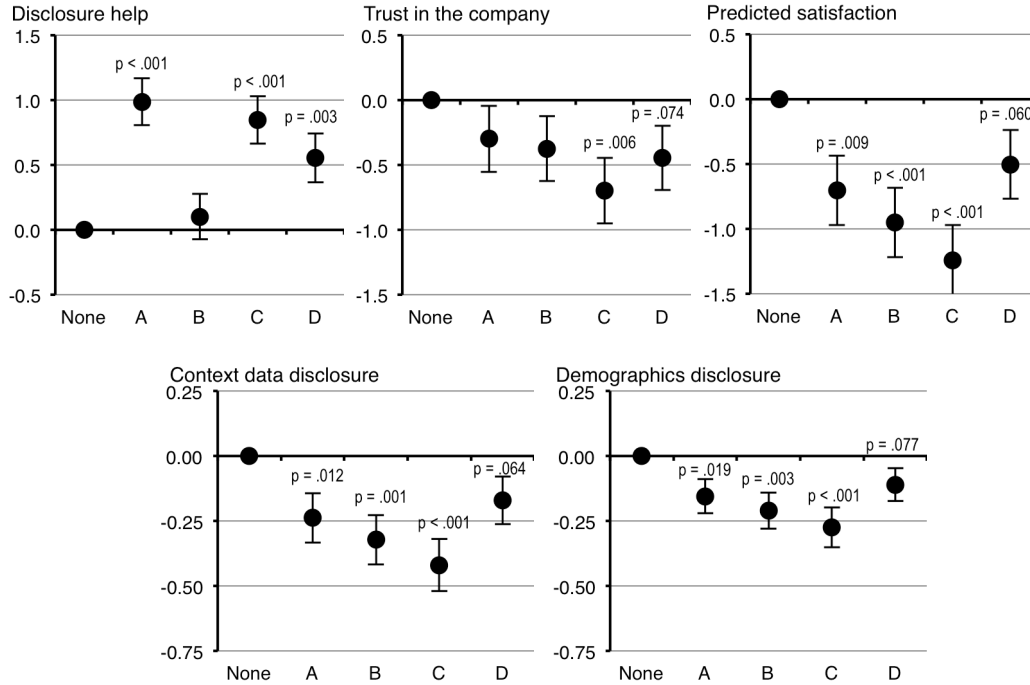


Fig. 4. The total effects of the justifications (A: Useful for you, B: Number of others, C: Useful for others, D: Explanation) on the different outcomes, tested against the baseline condition (No justification). Vertical axes are in sample standard deviations of the measured factor (i.e. 95% of the participants have a value between -2 and $+2$). Error bars are ± 1 standard error of the measurement.

5. DISCUSSION

The results of the experiment provide interesting insights into users' information disclosure decisions when using a recommender system. They also demonstrate how personal privacy concerns and manipulated system aspects influence the decision process. In this section we reflect on these results by integrating them with findings from an interview study with 17 participants who used a paper prototype of our recommender system. Details of the interview study are presented in the electronic appendix.

5.1 The Cognitive Process Behind Information Disclosure Decisions

Based on the results of the experiment and the findings of our interview study, we can make the following claims about the cognitive process involved in making information disclosure decisions:

The disclosure decision is first and foremost the outcome of an assessment of the self-anticipated satisfaction with the system.

In the experimental model, almost all effects on information disclosure are mediated by users' self-anticipated satisfaction with the system. This is in line with the interview findings: 16 out of 17 participants mentioned at least once the potential usefulness of providing the information as a reason for disclosure. The experiment shows that the effect of self-anticipated satisfaction on context disclosure is about 1.5

times higher than the effect on demographics disclosure ($\beta = 0.338$ vs. $\beta = 0.220$), despite the fact that context data is more privacy-sensitive (see Table II).

User's self-anticipated satisfaction with the system is strongly impacted by their trust in the company, the perceived value of disclosure help, and the perceived privacy threats entailed by the disclosure.

In the experiment, these three factors are the main subjective determinants of satisfaction with the system. Moreover, part of the effect of perceived value of disclosure help and perceived privacy threat is mediated by participants' trust in the company, indicating that users' evaluation of the system can have a lasting effect on the reputation of the company.

In the interview study, 15 out of 17 participants also mentioned the reputation of or trust in the company as a reason to disclose information or rather not. This result is in line with a prior study by Teo et al. [2004]. Most of our participants were able to recall a recent privacy scandal, and had lowered their evaluation of the involved company as a result.

Interview participants also mentioned privacy threat as a determinant of their disclosure decisions. Privacy threats were mentioned in a positive as well as negative sense: 14 participants said they would disclose something because it did not pose a privacy threat; 13 participants said they would not disclose something because it did pose a threat. For 9 participants, privacy concerns occasionally trumped their initial sense of the usefulness of the information; they would deem disclosing the information "not worth the risk". Typical threats mentioned were unwanted advertisements (mentioned 11 times), the company selling their information (mentioned 12 times) and security concerns or other unintended breaches of confidentiality (mentioned by all 17 participants).

5.2 Effects of Personal Privacy Concerns

Privacy concerns influence the disclosure decision via perceived value of disclosure help and perceived privacy threat, but the effects of different types of privacy concerns vary considerably:

Users' collection concerns decrease the perceived value of disclosure help, and increase the perceived privacy threat.

According to the experimental model, collection concerns decrease the valuation of disclosure help. Users with high collection concerns may feel that the system has ulterior motives to 'help' them in their disclosure. For instance, 9 participants in the interview study were skeptical about the veracity of the stated percentage in the justification message. This is also in line with our findings in [Knijnenburg and Kobsa 2013] where we demonstrate that especially for males with a low disclosure tendency, it is best not to 'help' them with a justification message.

In contrast, users' control concerns increase the perceived value of disclosure help and reduce the perceived privacy threat.

The model shows that, controlling for collection concerns, control concerns actually have a positive impact on the perceived value of disclosure help and a negative impact on perceived privacy threat. This is in line with Nowak and Phelps' [1995] postulate that when users perceive to be in control of their information disclosure, this actually reduces the significance of privacy threats. Regardless of the justification or the request order, the Applause system allows users to control the disclosure

of each piece of information separately. 8 participants in the interview study noted that they liked this feature, and 7 of them believed that the system adequately protected their privacy by providing this level of control. However, 6 other participants did not feel in control, and consequently felt that the system was not helping them at all, and that the requests just had the purpose of invading their privacy.

Control concerns have more impact on the perception of disclosure help, and collection concerns have more impact on perceived privacy threats, but in the end they have a more or less equal and opposite effect on self-anticipated satisfaction and disclosure behavior.

In the model, control concerns have the largest impact on the perceived value of disclosure help ($\beta = 0.414$ vs. $\beta = -0.273$ for collection concerns), whereas collection concerns have a larger impact on perceived privacy threat ($\beta = 0.735$ vs. $\beta = -0.438$ for control concerns). Considering the total effects of control and collection concerns, their impact on self-anticipated satisfaction ($\beta = 0.524$, vs. $\beta = -0.632$), context data disclosure ($\beta = 0.177$, vs. $\beta = -0.214$), and demographics disclosure ($\beta = 0.116$, vs. $\beta = -0.139$)¹¹ are roughly equal but opposite.

General privacy concerns cause both control and collection concerns, but have a total negative impact on self-anticipated satisfaction and disclosure behavior.

In the model, general privacy concerns drive both control and collection concerns, but the effect on collection concerns is much stronger ($\beta = 0.412$ for control concerns, vs. $\beta = 1.270$ for collection concerns). The total effects of general privacy concerns on self-anticipated satisfaction ($\beta = -0.588$), context data disclosure ($\beta = -0.199$), and demographics disclosure ($\beta = -0.130$)¹² are therefore negative. In the interview study, 9 participants explained that their overall concern about privacy issues influenced the way they approached individual information disclosure decisions.

5.3 Effects of Strategies

As expected, the different types of justification messages and request order manipulations influenced participants' disclosure behavior. Our unified approach allows us to explain in detail how this influence plays out:

Except for 'number of others', our justifications increase users' valuation of disclosure help.

The results of the experiment show a direct effect of the justifications on the perceived value of disclosure help (see Fig. 4). More specifically, the 'useful for you', 'useful for others' and 'explanation' messages each increase the valuation of disclosure help compared to providing no justification. Likewise, 12 participants in the interview study mentioned that they appreciated the help that these messages provided. Interestingly, the 'number of others' justification provides no additional disclosure help. This is in line with the interview study results: 11 participants do not like this justification at all, and some even believe that it is worse than having no justification. Participants mentioned that the message "feels like peer pressure".

¹¹ See the electronic appendix for an example of how to calculate total effects. For the mentioned total effects, all p -values are $< .001$.

¹² For these total effects, all p -values are $< .001$.

Except for the ‘number of others’ justification, the justifications decrease users’ trust in the company, and all justifications decrease users’ self-anticipated satisfaction.

The experimental model shows that negative effects on trust in the company negate the positive effects of justifications on perceived value of disclosure help. The total effects (see Fig. 4) show that overall, the ‘useful for others’ and ‘explanation’ message reduce trust in the company. All justifications have a negative direct effect on users’ self-anticipated satisfaction, and as a result the total effects are also negative. Comparing these total effects with the total effects of personal privacy concerns (see Section 5.3), they are similar in size. In other words: privacy concerns and justifications have a roughly equal effect on users’ trust and satisfaction.

We find no parallel of these effects in our interview study. Interestingly, the ‘number of others’ justification is again the odd one out, in that it does not significantly decrease the trust in the company. Arguably, users regard the number of other users disclosing the requested data as a sufficiently neutral statistic.

Ultimately, the justifications lower users’ disclosure rates.

The total effects (see Fig. 4) of the justifications on disclosure are all negative, which means that the baseline system without justification actually results in the highest disclosure rates. These effects are again roughly the same size as the effects of privacy, collection and control concerns. The interview study reveals that users typically treat the justification message as a warning sign: 11 participants mentioned 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). It seems that the justification provide more inhibition than encouragement.

Elsewhere [Knijnenburg and Kobsa 2013], we show that these effects occur regardless of the percentage in the justification message (except for the ‘number of others’ justification, but even for that message a high percentage merely reduces the negative effect, and never actually increases disclosure). This is in line with the interview study results. Participants generally had a ‘cut-off’ percentage below which they would not disclose something. Moreover, 14 participants would at several occasions refuse to disclose something they deemed too private despite a high percentage.

In [Knijnenburg and Kobsa 2013] we demonstrate however that choosing the justification based on characteristics of the individual user may be a way to increase both disclosure and satisfaction.

Changing the request order increases the disclosure of the data requested first but decreases disclosure of data requested later in the interaction.

The model shows an inherent trade-off in the order of the requests: requesting demographical data first (as opposed to requesting context data first) increases demographics disclosure ($\beta = 0.315$) and decreases context data disclosure ($\beta = -0.928$), and vice versa. In other words, asking a certain type of data first increases its disclosure, but decreases the disclosure of the other type of data.

There are two possible explanations for this effect. One is that users become more wary of privacy threats as the data collected about them accumulates. In the interview study, 5 participants mentioned that at some point, they felt that the combination of several items they disclosed caused additional privacy concerns. An alternative explanation is that users get tired of answering so many disclosure requests. Support for this comes from the fact that 9 interviewees mentioned that

they had to answer too many requests, and 6 participants noted that they would decline disclosure if it would take too much effort to disclose the information.

There are reasons to believe that the former explanation holds more ground. If the latter explanation were correct, the effect should be most pronounced for demographics disclosure, because in the current system it takes more effort to disclose demographical data than context data (since demographics disclosure requires the user to key in the data, whereas context disclosure merely requires users to click a ‘yes’ or ‘no’ button). In fact, though, the effect is stronger for context data disclosure than for demographics disclosure (see Table II). This is in line with Acquisti et al. [2011] who also find that the order effect is strongest for more sensitive data.

6. CONCLUSION AND FUTURE WORK

The results of our unified approach successfully describe how users make information disclosure decisions in context-based recommender systems, in dependence on privacy concerns and manipulated disclosure justification strategies. Specifically, we demonstrate that users consider satisfaction, trust, perceived threats and system-provided disclosure help when making information disclosure decisions.

Our results leave room for future work. First, the experimental model is based on the results of an explorative effort with a single system (the core framework has however been successfully validated with four other recommender systems; see Knijnenburg et al. [2012]). Confirmatory studies with different systems in other domains can verify the generalizability of our model. Moreover, we merely studied a mockup system, albeit we made sure that participants had the impression that they were disclosing their data to a “real” company. Participants also did not receive any actual recommendations, and therefore had no chance to adjust their disclosure behavior to actual system results. 16 out of 17 participants in the interview study stated that this would be a strategy they would follow; future research should explore this reactive user behavior.

Regardless of these caveats, our unified approach provides a good platform for testing different system strategies to influence disclosure. Specifically, it was able to explain the unexpected negative effect of justifications on disclosure behavior: although justifications increase the perceived value of disclosure help that the system offers, this positive effect is canceled out by negative effects on trust and satisfaction. In the current experiment, justifications thus mainly made users more skeptical about the intentions of the system and the possible benefits it can provide.

We also found that early requests are more likely to receive answers than later requests. This effect is stronger on context disclosure (the more sensitive data) than on demographics disclosure. Acquisti et al. [2011] also found that asking the most sensitive questions first increases overall disclosure. Designers should therefore not sequence requests based on their usefulness for the recommendation quality only, but also on their privacy sensitivity giving priority to more sensitive questions. However, it still remains to be seen how far one can take this without offending users.

Our results indicate that additional research is needed to come up with the ‘best’ request order, and with justifications that are both convincing and trust-inducing. An alternative approach is to tailor the justifications and request order to the user and the usage situation, an approach we consider in [Knijnenburg and Kobsa 2013]. Following this approach, one could envision an adaptive system that takes into account the user’s request history, and dynamically selects the next request plus

justification based on this history, the context, and the goals of the system (e.g. increasing disclosure and/or increasing satisfaction).

At a more general level, our study demonstrates that in order to attain robust results and careful explanations of discovered effects, research on privacy decision-making should take a unified approach that considers personal privacy characteristics, information disclosure behavior and user experience.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library.

REFERENCES

- ACKERMAN, M.S., CRANOR, L.F., AND REAGLE, J. 1999. Privacy in e-commerce: examining user scenarios and privacy preferences. In *Proceedings of the 1st ACM conference on Electronic commerce*, Denver, CO, November 1999, ACM Press, 1–8. DOI:<http://dx.doi.org/10.1145/336992.336995>
- ACQUISTI, A. AND GROSSKLAGS, J. 2005. Privacy and Rationality in Individual Decision Making. *IEEE Security & Privacy* 3, 1, 26–33. DOI:<http://dx.doi.org/10.1109/MSP.2005.22>
- ACQUISTI, A. AND GROSSKLAGS, J. 2008. What Can Behavioral Economics Teach Us About Privacy? In *Digital Privacy: Theory, Technologies, and Practices*, A. ACQUISTI, S. DE CAPITANI DI VIMERCATI, S. GRITZALIS AND C. LAMBRINOUDAKIS, EDS. Taylor & Francis, 363–377.
- ACQUISTI, A., JOHN, L., AND LOEWENSTEIN, G. 2009. What is privacy worth? In *Twenty First Workshop on Information Systems and Economics*, Phoenix, AZ, December 2009.
- ACQUISTI, A., JOHN, L.K., AND LOEWENSTEIN, G. 2011. The Impact of Relative Standards on the Propensity to Disclose. *Journal of Marketing Research*, 1–15. DOI:<http://dx.doi.org/10.1509/jmr.09.0215>
- ADOMAVICIUS, G. AND TUZHILIN, A. 2011. Context-Aware Recommender Systems. In *Recommender Systems Handbook*, F. RICCI, L. ROKACH, B. SHAPIRA AND P.B. KANTOR, EDS. Springer US, Boston, MA, 217–253. DOI:http://dx.doi.org/10.1007/978-0-387-85820-3_7
- ARCHER, R.L. AND BERG, J.H. 1978. Disclosure reciprocity and its limits: A reactance analysis. *Journal of Experimental Social Psychology* 14, 6, 527–540. DOI:[http://dx.doi.org/10.1016/0022-1031\(78\)90047-1](http://dx.doi.org/10.1016/0022-1031(78)90047-1)
- BÉLANGER, F. AND CROSSLER, R.E. 2011. Privacy in the Digital Age: A Review of Information Privacy Research in Information Systems. *MIS Quarterly* 35, 4, 1017–1045.
- BENTLER, P.M. AND BONETT, D.G. 1980. Significance Tests and Goodness of Fit in the Analysis of Covariance Structures. *Psychological Bulletin* 88, 3, 588–606.
- BESMER, A., WATSON, J., AND LIPFORD, H.R. 2010. The impact of social navigation on privacy policy configuration. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, Redmond, Washington, July 2010. DOI:<http://dx.doi.org/10.1145/1837110.1837120>
- BÖHMER, M., BAUER, G., AND KRÜGER, A. 2010. Exploring the Design Space of Context-aware Recommender Systems that Suggest Mobile Applications. In *2nd Workshop on Context-Aware Recommender Systems*, Barcelona, Spain, September 2010.
- BUCHANAN, T., PAINE, C., JOINSON, A.N., AND REIPS, U.-D. 2007. Development of Measures of Online Privacy Concern and Protection for Use on the Internet. *Journal of the American Society for Information Sciences and Technology* 58, 2, 157–165. DOI:<http://dx.doi.org/10.1002/asi.20459>
- CANNY, J. 2002a. Collaborative Filtering with Privacy. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, Oakland, CA, May 2002, IEEE Press, 45–57. DOI:<http://dx.doi.org/10.1109/SECPRI.2002.1004361>
- CANNY, J. 2002b. Collaborative Filtering with Privacy via Factor Analysis. In *25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Tampere, Finland, August 2002, ACM Press, 238–245. DOI:<http://dx.doi.org/10.1145/564376.564419>
- 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, 181–202. DOI:<http://dx.doi.org/10.1007/s10799-005-5879-y>
- 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. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Portland, OR, April 2005, 81–90. DOI:<http://dx.doi.org/10.1145/1054972.1054985>
- CULNAN, M.J. 1993. “How Did They Get My Name?”: An Exploratory Investigation of Consumer Attitudes toward Secondary Information Use. *MIS Quarterly* 17, 3, 341–363. DOI:<http://dx.doi.org/10.2307/249775>
- DAVIDSSON, C. AND MORITZ, S. 2011. Utilizing implicit feedback and context to recommend mobile applications from first use. In *Proceedings of the 2011 Workshop on Context-awareness in Retrieval and*

- Recommendation*, Palo Alto, California, February 2011, ACM Press, 19–22. DOI:<http://dx.doi.org/10.1145/1961634.1961639>
- EGELMAN, S., TSAI, J., CRANOR, L.F., AND ACQUISTI, A. 2009. Timing is everything?: the effects of timing and placement of online privacy indicators. In *Proceedings of the 27th international conference on Human factors in computing systems*, Boston, MA, April 2009, ACM Press, 319–328. DOI:<http://dx.doi.org/10.1145/1518701.1518752>
- FITZSIMONS, G.J. AND LEHMANN, D.R. 2004. Reactance to Recommendations: When Unsolicited Advice Yields Contrary Responses. *Marketing Science* 23, 1, 82–94. DOI:<http://dx.doi.org/10.1287/mksc.1030.0033>
- VAN DE GARDE-PERIK, E., MARKOPOULOS, P., DE RUYTER, B., EGGEN, B., AND IJSSELSTELJN, W. 2008. Investigating Privacy Attitudes and Behavior in Relation to Personalization. *Social Science Computer Review* 26, 1, 20–43. DOI:<http://dx.doi.org/10.1177/0894439307307682>
- GIGERENZER, G. AND GOLDSTEIN, D.G. 1996. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* 103, 4, 650–669. DOI:<http://dx.doi.org/10.1037/0033-295X.103.4.650>
- GIRARDELLO, A. AND MICHAHELLES, F. 2010. AppAware: which mobile applications are hot? In *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services*, Lisbon, Portugal, September 2010, ACM Press, 431–434. DOI:<http://dx.doi.org/10.1145/1851600.1851698>
- HARRIS, L. AND ASSOCIATES. 1990. *The Equifax Report on Consumers in the Information Age*. Equifax Inc., Atlanta, GA.
- HARRIS, L. AND ASSOCIATES. 1991. *Harris-Equifax Consumer Privacy Survey 1991*. Equifax Inc., Atlanta, GA.
- HARRIS, L., ASSOCIATES, AND WESTIN, A.F. 1998. *Personalized Marketing and Privacy on The Net: What Consumers Want*. Privacy and American Business Newsletter.
- HARRIS, L., WESTIN, A.F., AND ASSOCIATES. 2003. *Consumer Privacy Attitudes: A Major Shift Since 2000 and Why*. Harris Interactive, Inc.
- HASSENZAHN, M. 2005. The thing and I: understanding the relationship between user and product. In *Funology, Human-Computer Interaction Series*, M.A. BLYTHE, K. OVERBEEKE, A.F. MONK AND P.C. WRIGHT, EDS. Springer Netherlands, 31–42. DOI:http://dx.doi.org/10.1007/1-4020-2967-5_4
- HU, L. AND BENTLER, P.M. 1999. Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural Equation Modeling: A Multidisciplinary Journal* 6, 1, 1–55. DOI:<http://dx.doi.org/10.1080/10705519909540118>
- HUI, K.-L., TAN, B.C.Y., AND GOH, C.-Y. 2006. Online information disclosure: Motivators and measurements. *ACM Transactions on Internet Technology* 6, 4, 415–441. DOI:<http://dx.doi.org/10.1145/1183463.1183467>
- HUI, K.-L., TEO, H.H., AND LEE, S.-Y.T. 2007. The Value of Privacy Assurance: An Exploratory Field Experiment. *MIS Quarterly* 31, 1, 19–33.
- IACHELLO, G. AND HONG, J. 2007. End-User Privacy in Human-Computer Interaction. *Foundations and Trends® in Human-Computer Interaction* 1, 1, 1–137. DOI:<http://dx.doi.org/10.1561/1100000004>
- JARVENPAA, S. AND TRACTINSKY, N. 1999. Consumer Trust in an Internet Store: A Cross-Cultural Validation. *Journal of Computer Mediated Communication* 5, 2. DOI:<http://dx.doi.org/10.1111/j.1083-6101.1999.tb00337.x>
- 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, 858–873. DOI:<http://dx.doi.org/10.1086/656423>
- KAHNEMAN, D., KNETSCH, J.L., AND THALER, R.H. 1990. Experimental Tests of the Endowment Effect and the Coase Theorem. *Journal of Political Economy* 98, 6, 1325–1348. DOI:<http://dx.doi.org/10.1086/261737>
- KAHNEMAN, D., SLOVIC, P., AND TVERSKY, A. 1982. *Judgment under uncertainty: heuristics and biases*. Cambridge University Press, Cambridge; New York.
- KAHNEMAN, D. AND TVERSKY, A. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47, 2, 263–292. DOI:<http://dx.doi.org/10.2307/1914185>
- KEITH, M.J., BABB, J.S., PAUL BENJAMIN LOWRY, FURNER, C.P., AND ABDULLAT, A. 2011. The Roles of Privacy Assurance, Network Effects, and Information Cascades in the Adoption of and Willingness to Pay for Location-Based Services with Mobile Applications. In *2011 Dewald Roode Information Security Workshop*, Blacksburg, VA, September 2011. <http://ifip.byu.edu/2011/Keith%20et%20al.%20Mobile%20Applications.pdf>
- KITTUR, A., CHI, E.H., AND SUH, B. 2008. Crowdsourcing user studies with Mechanical Turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Florence, Italy, April 2008, ACM Press, 453–456. DOI:<http://dx.doi.org/10.1145/1357054.1357127>
- KNIJNENBURG, B.P. AND KOBSA, A. 2013. Helping users with information disclosure decisions: potential for adaptation. In *Proceedings of the 2013 ACM international conference on Intelligent User Interfaces*, Santa Monica, CA, March 2013, ACM Press, 407–416. Forthcoming, 10.1145/2449396.2449448

- KNIJNENBURG, B.P., WILLEMSSEN, 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, 441–504. DOI:<http://dx.doi.org/10.1007/s11257-011-9118-4>
- KOBSA, A. 2007. Privacy-Enhanced Web Personalization. In *The Adaptive Web: Methods and Strategies of Web Personalization*, P. BRUSILOVSKY, A. KOBSA AND W. NEJDJL, EDS. Springer Verlag, Berlin Heidelberg New York, 628–670. DOI:http://dx.doi.org/10.1007/978-3-540-72079-9_21
- KOBSA, A. AND TELTZROW, M. 2005. Contextualized communication of privacy practices and personalization benefits: Impacts on users' data sharing and purchase behavior. In *Privacy Enhancing Technologies: Revised Selected Papers of the 4th International Workshop, PET 2004, Toronto, Canada, May 26-28, 2004*, D. MARTIN AND A. SERJANTOV, EDS. Springer Berlin Heidelberg, 329–343. DOI:<http://dx.doi.org/10.1007/b136164>
- KOLTER, J. AND PERNUL, G. 2009. Generating User-Understandable Privacy Preferences. In *2009 International Conference on Availability, Reliability and Security*, Fukuoka, Japan, March 2009, IEEE Computer Society, 299–306. DOI:<http://dx.doi.org/10.1109/ARES.2009.89>
- LAUFER, R.S., PROSHANSKY, H.M., AND WOLFE, M. 1973. Some Analytic Dimensions of Privacy. In *Proceedings of the Lund Conference on Architectural Psychology*, Lund, Sweden, 1973, R. KÜLLER, ED. Dowden, Hutchinson & Ross.
- 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, 22–42. DOI:<http://dx.doi.org/10.1111/j.1540-4560.1977.tb01880.x>
- LEDERER, S., MANKOFF, J., AND DEY, A.K. 2003. Who wants to know what when? privacy preference determinants in ubiquitous computing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Ft. Lauderdale, FL, April 2003, ACM Press, 724–725. DOI:<http://dx.doi.org/10.1145/765891.765952>
- LEE, H.J. AND PARK, S.J. 2007. MONERS: A news recommender for the mobile web. *Expert Systems with Applications* 32, 1, 143–150. DOI:<http://dx.doi.org/10.1016/j.eswa.2005.11.010>
- LEE, J. AND LEE, J. 2007. Context Awareness by Case-Based Reasoning in a Music Recommendation System. In *Proceedings of the 4th International Symposium on Ubiquitous Computing Systems*, Tokyo, Japan, November 2007, H. ICHIKAWA, W.-D. CHO, I. SATOH AND H. YOUN, EDS. Springer Berlin Heidelberg, 45–58. DOI:http://dx.doi.org/10.1007/978-3-540-76772-5_4
- LI, H., SARATHY, R., AND XU, H. 2010. Understanding situational online information disclosure as a privacy calculus. *Journal of Computer Information Systems* 51, 1, 62–71.
- LI, H., SARATHY, R., AND XU, H. 2011. The role of affect and cognition on online consumers' decision to disclose personal information to unfamiliar online vendors. *Decision Support Systems* 51, 3, 434–445. DOI:<http://dx.doi.org/10.1016/j.dss.2011.01.017>
- LI, X. AND SANTHANAM, R. 2008. Will it be Disclosure or Fabrication of Personal Information? An Examination of Persuasion Strategies on Prospective Employees. *International Journal of Information Security and Privacy* 2, 4, 91–109. DOI:<http://dx.doi.org/10.4018/jisp.2008100105>
- MABLEY, K. 2000. *Privacy vs. Personalization: Part III*. Cyber Dialogue, Inc.
- MALHOTRA, N.K., KIM, S.S., AND AGARWAL, J. 2004. Internet Users' Information Privacy Concerns (UIIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research* 15, 4, 336–355. DOI:<http://dx.doi.org/10.1287/isre.1040.0032>
- METZGER, M.J. 2004. Privacy, trust, and disclosure: Exploring barriers to electronic commerce. *Journal of Computer-Mediated Communication* 9, 4.
- METZGER, M.J. 2006. Effects of Site, Vendor, and Consumer Characteristics on Web Site Trust and Disclosure. *Communication Research* 33, 3, 155–179. DOI:<http://dx.doi.org/10.1177/0093650206287076>
- NISSENBAUM, H.F. 2010. *Privacy in context: technology, policy, and the integrity of social life*. Stanford Law Books, Stanford, CA.
- NOWAK, G.J. AND PHELPS, J. 1995. Direct marketing and the use of individual-level consumer information: Determining how and when “privacy” matters. *Journal of Direct Marketing* 9, 3, 46–60. DOI:<http://dx.doi.org/10.1002/dir.4000090307>
- OH, J.-M. AND MOON, N. 2012. User-selectable interactive recommendation system in mobile environment. *Multimedia Tools and Applications* 57, 2, 295–313. DOI:<http://dx.doi.org/10.1007/s11042-011-0737-x>
- PATIL, S., PAGE, X., AND KOBSA, A. 2011. With a little help from my friends: can social navigation inform interpersonal privacy preferences? In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, Hangzhou, China, March 2011, ACM Press, 391–394. DOI:<http://dx.doi.org/10.1145/1958824.1958885>
- POLAT, H. AND DU, W. 2005a. Privacy-Preserving Collaborative Filtering. *International Journal of Electronic Commerce* 9, 4, 9–35.
- POLAT, H. AND DU, W. 2005b. SVD-based Collaborative Filtering with Privacy. In *Proceedings of the 2005 ACM symposium on Applied computing*, Santa Fe, New Mexico, March 2005, ACM Press, 791–795. DOI:<http://dx.doi.org/10.1145/1066677.1066860>

- PU, P., CHEN, L., AND HU, R. 2011. A user-centric evaluation framework for recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*, Chicago, IL, October 2011, ACM Press, 157–164. DOI:<http://dx.doi.org/10.1145/2043932.2043962>
- RICCI, F. 2011. Mobile Recommender Systems. *Information Technology & Tourism* 12, 3, 205–231. DOI:<http://dx.doi.org/10.3727/109830511X12978702284390>
- RIFON, N.J., LAROSE, R., AND CHOI, S.M. 2005. Your Privacy Is Sealed: Effects of Web Privacy Seals on Trust and Personal Disclosures. *Journal of Consumer Affairs* 39, 2, 339–360. DOI:<http://dx.doi.org/10.1111/j.1745-6606.2005.00018.x>
- SMITH, H.J., DINEV, T., AND XU, H. 2011. Information Privacy Research: An Interdisciplinary Review. *MIS Quarterly* 35, 4, 989–1015.
- SMITH, H.J., MILBERG, S.J., AND BURKE, S.J. 1992. Concern for Privacy Instrument. *Working document*, School of Business, Georgetown University, Washington, DC.
- SMITH, H.J., MILBERG, S.J., AND BURKE, S.J. 1996. Information Privacy: Measuring Individuals' Concerns about Organizational Practices. *MIS Quarterly* 20, 2, 167–196. DOI:<http://dx.doi.org/10.2307/249477>
- SOLOVE, D.J. 2006. A Taxonomy of Privacy. *University of Pennsylvania Law Review* 154, 3, 477–564.
- SPIEKERMANN, S., GROSSKLAGS, J., AND BERENDT, B. 2001. E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus actual Behavior. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, Tampa, FL, October 2001, ACM Press, 38–47. DOI:<http://dx.doi.org/10.1145/501158.501163>
- STEWART, K.A. AND SEGARS, A.H. 2002. An Empirical Examination of the Concern for Information Privacy Instrument. *Information Systems Research* 13, 1, 36–49. DOI:<http://dx.doi.org/10.1287/isre.13.1.36.97>
- TAYLOR, D., DAVIS, D., AND JILLAPALLI, R. 2009. Privacy concern and online personalization: The moderating effects of information control and compensation. *Electronic Commerce Research* 9, 3, 203–223. DOI:<http://dx.doi.org/10.1007/s10660-009-9036-2>
- TEO, H.H., WAN, W., AND LI, L. 2004. Volunteering Personal Information on the Internet: Effects of Reputation, Privacy Initiatives, and Reward on Online Consumer Behavior. In *Proceedings of the 37th Hawaii International Conference on System Sciences*, Big Island, HI, January 2004, IEEE Press, 181–190. DOI:<http://dx.doi.org/10.1109/HICSS.2004.1265435>
- THALER, R. 1980. Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization* 1, 1, 39–60. DOI:[http://dx.doi.org/10.1016/0167-2681\(80\)90051-7](http://dx.doi.org/10.1016/0167-2681(80)90051-7)
- TSAI, J.Y., EGELMAN, S., CRANOR, L., AND ACQUISTI, A. 2010. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research*. DOI:<http://dx.doi.org/10.1287/isre.1090.0260>
- TUROW, J., KING, J., HOOFNAGLE, C.J., BLEAKLEY, A., AND HENNESSY, M. 2009. Americans Reject Tailored Advertising and Three Activities That Enable It. Manuscript posted on the Social Science Research Network: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1478214
- 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, 217–246.
- WENNING, R. AND SCHUNTER, M. 2006. *The Platform for Privacy Preferences 1.1 (P3P1.1) Specification*. W3C Working Group Note.
- WESTIN, A.F., HARRIS, L., AND ASSOCIATES. 1981. *The Dimensions of privacy: a national opinion research survey of attitudes toward privacy*. Garland Publishing, New York.
- XU, H., LUO, X. (ROBERT), CARROLL, J.M., AND ROSSON, M.B. 2011. The personalization privacy paradox: An exploratory study of decision making process for location-aware marketing. *Decision Support Systems* 51, 42–52. DOI:<http://dx.doi.org/10.1016/j.dss.2010.11.017>
- XU, H., TEO, H.-H., TAN, B., AND AGARWAL, R. 2009. The Role of Push-Pull Technology in Privacy Calculus: The Case of Location-Based Services. *Journal of Management Information Systems* 26, 3, 135–174. DOI:<http://dx.doi.org/10.2753/MIS0742-1222260305>
- ZHANG, N., WANG, C., AND XU, Y. 2011. Privacy in Online Social Networks. In *Proceedings of the 2011 International Conference on Information Systems*, Shanghai, China, December 2011, Paper 3. <http://aisel.aisnet.org/icis2011/proceedings/ISsecurity/3>
- ZHENG, V.W., ZHENG, Y., XIE, X., AND YANG, Q. 2012. Towards mobile intelligence: Learning from GPS history data for collaborative recommendation. *Artificial Intelligence* 184–185, 17–37. DOI:<http://dx.doi.org/10.1016/j.artint.2012.02.002>

Online Appendix to: Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems

BART P. KNIJNENBURG, University of California, Irvine

ALFRED KOBSA, University of California, Irvine

A. STATISTICAL METHODOLOGY

The statistical analysis of the experiment presented in this paper consists of two parts. We first validate the measured latent concepts in a series of Confirmatory Factor Analyses (CFAs), one for disclosure behavior and one for the subjective measures. Subsequently, we test the structural relations between the manipulations and the subjectively and behaviorally measured latent concepts using Structural Equation Modeling (SEM). In this section, we provide a detailed explanation of these statistical methods using the data from this paper as an example.

A.1 Confirmatory Factor Analysis (CFA)

The first step is to confirm whether users' behavior and subjective evaluations measure the predicted latent factors. A latent factor captures the common variance (or 'shared essence') of the items that are used to measure it. A robust factor will only emerge if the items have enough common variance.

The post-experimental questionnaire items are created in such a way that we have a clear pre-existing notion of which items measure which factor (see Tables II and III of the main text). We thus set up our model in accordance to this hypothesis, as shown in Fig. A.1 for the factors 'general privacy concerns', 'collection concerns', and 'control concerns'. In this model, each of the items $\{gpc1\dots ctrl5\}$ is essentially a regression outcome¹³, predicted by the unobserved latent variables $\{gpc, coll, ctrl\}$. $\{I_{gpc,gpc1}\dots I_{ctrl,ctrl5}\}$ are the loadings of the items on the factors. The model tries to estimate these loadings so that the paths match the covariance matrix of the items $\{gpc1\dots ctrl5\}$ as closely as possible (e.g. $cov_{gpc1,gpc2} \approx I_{gpc,gpc1} * I_{gpc,gpc2}$, and $cov_{gpc1,coll1} = I_{gpc,gpc1} * w_{gpc,coll} * I_{coll,coll1}$). Intuitively speaking, factor analysis tries to model the "overlap" between items. The part of the variance that does not overlap (the 'uniqueness') is excluded (and represented by the arrows at the bottom of Fig. A.1). The more overlap gets extracted, the more reliable the factors are (the reliability of each factor is represented by the slanted arrows pointing to the factors from the top). The factors may be correlated with each other (e.g. $w_{gpc,coll}$). The solution has no standard scale, so we give the factors a standard deviation of 1 and a mean of 0.

If the model was specified correctly, it has a good least-squares fit, and the paths $\{I_{gpc,gpc1}\dots I_{ctrl,ctrl5}\}$ are significantly larger than zero. However, the following problems may arise:

- A certain item does not have enough in common with the other items to load significantly on its factor: In this case the item has "low communality" and should be removed from the analysis. We typically aim for at least 30% of the variance of an item to be captured by the factor.

© 2010 ACM 1539-9087/2010/03-ART39 \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

¹³ Since these outcomes are measured on a 5-point scale, this regression model is an ordinal response model.

- A certain item loads strongly on the wrong factor (this can be discovered by inspecting the ‘modification indices’ of the model): In this case, the item has a high ‘cross-loading’, and unless we can come up with a good reason for this to occur, it should be removed from the analysis.
- Two or more items may have a high residual correlation (discovered by inspecting the modification indices): In this case, the factors do not represent the amount of overlap between those items well enough. We can choose to either remove the items, or to ‘split up’ or redefine the existing factors in a way that better captures the common variance of the items.
- The correlation between two factors may be higher than the average variance extracted (AVE) from each item: In this case, we lack discriminant validity, and we must conclude that these two factors essentially measure the same concept.
- A certain factor has only low-loading items, with an AVE of < 0.50 : In this case, the factor is not measured robustly, and should be reconsidered. Another indicator of low robustness is Cronbach’s alpha, which is acceptable at $.7 < \alpha < .8$, good at $.8 < \alpha < .9$, and excellent at $\alpha > .9$.

In the current paper, we conducted this iterative process separately for the disclosure behaviors (yes/no items) and questionnaires (5-point scale items). Tables II and III display the resulting factors. Once an adequate factor solution is established, the remaining items are used in the second step of the analysis.

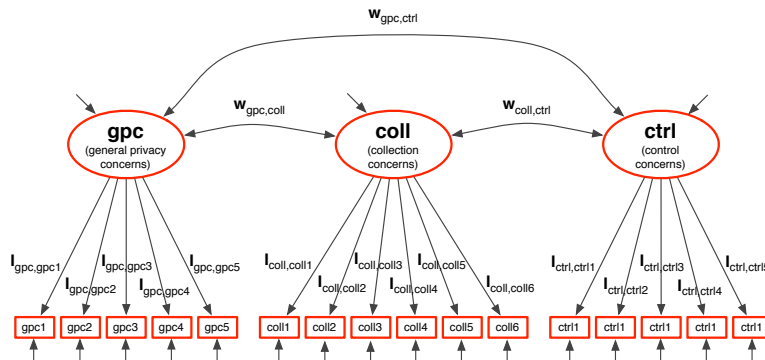


Fig. A.1. An example CFA, limited to the factors ‘general privacy concerns’, ‘collection concerns’, and ‘control concerns’

A.2 Structural Equation Modeling (SEM)

The second step is to test the structural relations between the experimental manipulations (justification type and request order) and the latent concepts. In Fig. 3 of the main text the structural relations are represented graphically as paths (arrows) between manipulations (rectangles) and latent concepts (ovals). Each set of incoming arrows essentially constitutes a linear regression. For example, the arrows pointing to ‘perceived privacy threat’ specify the regression:

$$y_i = \alpha + \beta_1 \eta_{1i} + \beta_2 \eta_{2i} + \varepsilon_i$$

where y_i is the value of ‘perceived privacy threat’ for participant i , η_{1i} is the value of ‘control concerns’ for participant i , η_{2i} is the value of ‘collection concerns’ for participant i , and ε_i is the residual (error). β_1 and β_2 are the regression coefficients (-0.438 and 0.735 respectively). Since the solution has no fixed intercept, α can be

ignored. Note that all latent concepts are still measured by their respective indicators such as in Fig. A.1 (these indicators are hidden in Fig. 3 to improve the clarity of the representation).

There are three types of regression paths in Fig. 3: paths emanating from a manipulation with two conditions, paths emanating from a manipulation with more than two conditions, and paths emanating from a latent concept.

The request order manipulation has two conditions: demographics first and context first. As is common in linear regressions, this manipulation is represented in the regressions by a single dummy variable: ‘demographics first’. This means that the baseline model is measured for ‘context first’, and that the effect of the dummy variable represents the difference between ‘context first’ and ‘demographics first’. Because all latent factors are standardized, the effect of this difference is measured in sample standard deviations. For example, the number on the arrow from ‘demographics first’ to ‘context disclosure’ (Fig. 3) indicates that context disclosure is 0.928 standard deviations lower in the ‘demographics first’ condition than in the ‘context first’ condition. The number in parentheses (0.129) is the standard error of this coefficient. This standard error can be used in a z-test to test whether the path is significant ($p = z[-.928/.129] < .001$).

The justification type manipulation has five conditions, with ‘no justification’ as the natural baseline condition. In the regression equations, a dummy variable is created for each of the four justifications. In Fig. 3, the effects labeled A, B, C and D represent the differences between the justifications and the baseline condition. Moreover, the χ^2 -value preceding these effects represents an omnibus test, which verifies whether the effect of all justifications combined is zero (for a significant omnibus effect, this should not be the case).

Finally, there are effects between two latent variables. These are also standardized regression effects, for example, a difference of one standard deviation in ‘disclosure help’ causes a 0.591 SD difference in ‘trust in the company’.

An *identifiable* structural equation model has to be a directed acyclic graph of causal paths. Such a graph can be constructed by assuming causal ‘order’ between variables. For instance, assuming that $\{A>B>C>D\}$ means that A can cause B, C and D (but B, C and D do not cause A), B can cause C and D (but C and D do not cause B) and C can cause D (but D does not cause C). An initial factor model is set up by defining all possible paths between the (latent) variables that respect the assumed causal order (i.e. the transitive closure). In a typical model, not all initially specified paths are significant, and non-significant paths can be removed. If a certain factor has multiple remaining incoming paths, the R^2 -value represents the fraction of variance explained by its predictors. For example, our model explains 27.3% of the variance in context data disclosure (due to the request order manipulation, and the effect of ‘self-anticipated satisfaction with the system’).

When there are only indirect and no direct effects between two variables, such as for ‘control concerns’ and ‘trust in the company’, we call this effect ‘fully mediated’. When there are both direct and indirect effects, such as for ‘perceived help’ and ‘satisfaction’, the effect is ‘partially mediated’. Finally, it is possible that the direct effect is positive, but the indirect effects are negative, or vice versa, such as for ‘justification type’ and ‘trust in the company’. This is called an ‘inconsistent mediation’. In the case of inconsistent mediations, total effects provide an evaluation of the sum of all positive and negative paths. The total effects can be calculated by adding all effects from a justification to a factor. For instance, the effect of

justification A on Trust is the direct effect $A \rightarrow \text{Trust}$, plus the effect $A \rightarrow \text{Help} \rightarrow \text{Trust}$. Using the values in Fig. 3: $-0.887 + .998 \cdot .591 = -0.303$.

The final model (after excluding non-significant effects) can be tested as a whole. The model χ^2 -statistic tests the difference in explained variance between the proposed model and a fully specified model. A good model is not statistically different from the fully specified model ($p > .05$). However, this statistic is commonly regarded as too sensitive, and researchers have therefore proposed other fit indices [Bentler and Bonett 1980]. Based on extensive simulations, Hu and Bentler [1999] propose cut-off values for other fit indices to be: $CFI > .96$, $TLI > .95$, and $RMSEA < .05$, with the upper bound of its 90% CI falling below 0.10.

A.3 Limitations of CFA and SEM

Using confirmatory factor analysis, one may run the epistemological risk of confirming a certain model when a better model exists unknowingly. The number of factors, and which items load on them, is not always straightforward. A look at the correlations between the measured items can give insights into the selection of a factor model; another option is to compare several factor models in terms of model fit. In our case, for example, we explicitly explored various factor solutions for disclosure behavior. The two-factor solution in Table II of the main text, which contrasts demographics disclosure and context data disclosure, had the best fit though.

Like any regression model, structural equation models make assumptions about the direction of causality in the model. From a modeling perspective, an effect ('disclosure help' \rightarrow 'trust in the company') and its reverse ('trust in the company' \rightarrow 'disclosure help') are equally plausible. By including manipulations into our model, we are able to 'ground' the causal effects: participants are randomly assigned to a condition, so condition assignment cannot be caused by anything in the model (i.e. 'justification type' \rightarrow 'disclose help' is possible, but not 'disclosure help' \rightarrow 'justification type'). Furthermore, the Knijnenburg et al. [2012] framework provides hypotheses for the directionality of causal effects (since in the framework we hypothesize that $SSA \rightarrow EXP$ and not $EXP \rightarrow SSA$, we can limit ourselves to testing 'disclosure help' \rightarrow 'trust in the company').

Furthermore all modeled structural relations are linear, although it is possible to include transformed variables, such as quadratic effects. In practice, however, such transformations are rarely needed, because latent factors are modeled to have a normal distribution.

B. OVERVIEW OF INTERVIEW FINDINGS

We conducted semi-structured interviews with 17 participants. The interviews lasted between half an hour and one hour, and consisted of two parts. In the first part, we showed participants paper versions of the Applause information disclosure screens. Participants were asked whether they would disclose the requested information, and to elaborate on this decision (i.e. to give a personal reason for disclosing or not disclosing the information). We showed participants a combination of the different justification types (see Table I in the original paper; the 'explanation' justification was added to the online experiment based on comments from the first 11 interview participants, and was therefore only tested in the last 6 interviews). At one point during the interview, we put screenshots of the different justification types next to each other, and explicitly asked participants to compare the justifications and to indicate which one they liked best.

After participants completed their disclosure decisions, we asked them to reflect on the application. We asked them to freely elaborate on their experience, but where appropriate, we primed them with specific questions on satisfaction, trust, threat, perceived help, and general privacy concerns, in line with the questions in the online study (Table III of the main text).

The interviews were then transcribed and structured using an ‘open-coding’ process, in which statements from participants are grouped and emerging groups are merged and split until a coherent and useful grouping remains. Below we do not strive to capture all the insights that emerged from the interview study, but rather give an overview of the discussed topics and the spectrum of participants’ answers.

B.1 Security issues

- Users think about security issues when faced with the decision to disclose personal information. All participants mentioned at least once that they would not disclose a certain piece of information due to security concerns.
- Recent security breaches can influence users’ perception of the security of online services. 4 participants mentioned the recent hacking of the PlayStation Network, and 1 participant said he would have answered the questions differently after he was told about this recent event.
- Users believe that nothing is 100% secure. Although 2 participants specifically stated that they believed the system to be secure, 6 participants expressed the belief that anything can be hacked.
- Users believe that mobile phones can be lost or stolen, and this poses additional security concerns. 3 participants said that they were afraid to lose their phone and that their data could thereby fall into the wrong hands.
- Companies are usually not blamed for security breaches. 3 participants mentioned that as long as the company does everything in their power to protect the data, they are not to blame for eventual breaches.
- Disclosure of credit card purchases was most influenced by beliefs about security. 8 participants gave security as a reason not to disclose credit card information, because they almost instantly linked this to identity theft.
- Location tracking was also influenced by security beliefs. 4 participants mentioned security in this context. These participants were afraid that an unauthorized third party could use the information to find out if they were at home or not (which could help in planning a burglary attempt).
- Email tracking was also influenced by security beliefs. 3 participants mentioned security in this context. These participants mentioned that this was because their email contained passwords and financial information.
- Some users try to cope with security issues by providing less detailed information to the system. 2 users mentioned that due to security issues, they would rather disclose less detailed information than what the system requested.

B.2 Other uses of users’ data

- All users worry about possible other uses of the information they disclose. Several participants mentioned at least once that they would not disclose a certain piece of information due to the fear that the information would be used in an unintended way. Others mentioned that they would not disclose certain information or even abandon the system if they found out that the information was being used in an unintended way.

- Most users dislike it when a company sells their personal information. 12 participants mentioned this explicitly, 5 of them were primed about it. However, 4 participants mentioned stated that they would not mind a company selling their personal information.
- Users think that advertisement is the most likely unintended use for their information. 11 participants explicitly mentioned advertisement, 2 of them were primed about it. Advertisement is however not a problem for all users. 6 of participants were okay with the system giving them advertisements based on their personal information. They said that this was just part of the deal. However, 4 participants mentioned that if they would pay for the app, they would expect the company not to sell their information or use it for advertisement.
- Surveillance is another unintended use mentioned by users. 5 participants mentioned this in the context of tracking the microphone. 2 other participants mentioned this in the context of tracking location.

B.3 Privacy concerns

- Users consider privacy when deciding what requests to answer. 13 participants would not disclose at least one piece of information because they deemed it too private, and 14 participants would disclose at least one piece of information because they did not deem it private. For 11 participants both situations occurred at least once. Only 1 participant did not use privacy as a reason for any of the disclosure decisions. 6 participants mentioned that some information is so private that the system should not even ask for it.
- Users talk about privacy in different ways. 12 participants mentioned that something was (not) “personal”. 6 participants referred to a gut feeling of discomfort. 5 participants mentioned the word “private” or “privacy”. 4 participants mentioned that something was “too much”, “too invasive”, or “too intrusive”. 4 participants said that they did not want the system to know a certain thing, or that they thought the system did not need to know a certain thing. 3 participants mentioned that something was (not) “secret” or “confidential”. 3 participants said that disclosing something “doesn’t bother me”.
- Users say that privacy has priority over usefulness. 8 participants disclosed at least one piece of information despite the fact that they thought it would not be useful for the system, because they did not deem it to be private. On the other hand, 9 participants felt that at least one piece of information was so private that they decided to not disclose it, even though they thought it would be useful for the system.
- The combination of several pieces of information can cause additional privacy concerns. 5 participants mentioned that at some point, they felt that the combination of several items they disclosed caused additional privacy concerns.
- Some users deal with privacy concerns by providing less detailed information. 7 participants mentioned that due to privacy issues, they would rather disclose less detailed information than what the system requested.
- Some users want more manual control over what they disclose to the system. 4 participants mentioned that they wanted to be able to control what the system would track.

B.4 Benefits of disclosure

- The benefits of disclosure are an important factor in deciding whether to disclose information, but not always. 13 participants explicitly stated that they would disclose something because it would be useful for the system to know, and 11 participants explicitly stated that they would not disclose something because it would not be useful for the system to know. However, 7 participants disclosed something despite thinking that it would not be useful, and 2 participants did not disclose something despite thinking that it would be useful.
- Users will occasionally give their intuitive idea of the benefits of providing certain information priority over the system-stated benefits (i.e. the justification message). 10 participants at least once gave their intuition priority over the system-stated benefits. 7 of them decided that something was useful despite the low percentage in the justification message, while 5 of them decided that something was not useful despite the high percentage in the system message. Only 1 participant stated that she changed her mind about her initial idea of the usefulness of allowing the system to track her accelerometer when the system showed a low percentage.

B.5 Control

- Users like to have the option to choose what to disclose. 8 participants liked the fact that they could decide what to disclose per specific piece of information, and 7 participants said that they would have problems with a system that uses their information without asking.
- Some users think that the system protects their privacy by giving them control over their settings. 7 participants say that the system protects their privacy because they can choose what to disclose. However, 6 participants did not think that the system protects their privacy at all.
- Giving users control makes some of them more cautious, but others less cautious. 8 participants mentioned that by asking specific questions, the system made them more cautious. On the other hand, 3 participants said that the fact that they could change their level of disclosure allowed them to be less cautious initially. 3 participants did not think that the system influenced their caution.
- Users may refuse to disclose something in fear that the system would make incorrect inferences based on the information. 7 participants did not disclose something because they were afraid it would lead to incorrect recommendations, 6 participants did not disclose something because they believed that the information misrepresented them, and 3 participants mentioned that some people might not disclose something because they could be ashamed of it. Conversely, 6 participants decided to disclose something because it represented them well.
- Many users like to be able to change their disclosure. 16 participants explicitly mentioned this, 9 of them said they would like to be able to change their answers after seeing the actual recommendations.
- Some users may lie as a way to control the system. 3 participants said they would lie on one of the requests.
- Some users want to manage their identity by giving information at a different level of detail. 4 participants mentioned this as a way to give the system a better representation of them.

B.6 Justifications

- Users may only skim the justification message. 8 participants said they did not completely read the message, and just looked at the percentage. 2 of them mentioned that the message was too wordy.
- Most users appreciate the justification message, but some are aware that it may trick them into disclosing more. When asked, 12 participants commented they liked to receive such a message. 5 participants believed that the messages could potentially trick them into disclosing more information, whereas 4 participants said that the system would not be able to trick them. 9 participants said that they would be skeptical about the veracity of the stated percentage.
- Most users are only occasionally influenced by the justification message. 11 participants at least once used a low percentage in the message as a reason not to disclose something. Conversely, 5 participants at least once used a high percentage in the message as a reason to disclose something. On the other hand, all participants at least once said they would disclose something despite a low percentage in the message, and 14 participants at least once said they would not disclose something despite a high percentage.
- The typical percentage at which users are convinced differs per user. 6 participants said that a percentage around or below 50% would be too low. 4 participants said that they would be okay with percentages around or somewhat below 50%.
- Which justification message works best differs per user. The “number of others” message is worst: 3 participants like it best, but 11 participants do not like it; some even find it worse than having no message at all. 10 participants prefer the “useful for you” or “useful for others” message, and these were often considered equal. 3 participants prefer the “explanation” message, but this message was only shown to the last 6 participants. In fact, 10 of the first 11 participants wanted an explanation justification, and 6 of them explicitly said they would prefer such an explanation over any other justification message.

B.7 Effort of disclosure

- Users do not want to answer too many requests. 9 participants mentioned this, and 5 of them wanted an option to skip certain requests and come back to them later.
- Users will not answer requests that take too much effort to answer. 6 participants decided to not answer a certain request because it was too difficult to answer quickly. 4 participants said that a certain request was superfluous (i.e., the system should already know the answer based on answers to previous questions).
- Whether users take the effort to answer a request may depend on their mood. 2 participants said that they would answer more or fewer requests depending on their mood.

B.8 Disclosure heuristics

- Users decide whether to disclose something by looking at what they normally do, or what they did before. 13 participants based their disclosure decision on what they normally do, or what they would disclose to other people. 11 participants based their disclosure decision on what they did in a specific

other situation (5 participants compare with Google, 3 with Amazon, 2 with Yelp, 1 with Facebook, and 1 with a credit card company). Similarly, 7 participants seemed to disclose something because the information was “already out there anyway”.

- Users’ decision depends on the company that provides the system. 15 participants mentioned the reputation of the company as a reason to disclose or not. 6 of them mentioned that they trust some companies more than others. 3 of them mentioned that they believe some companies would provide a higher security than others. 3 of them mentioned that they would trust a company that has this information already anyway (e.g. a bank). 7 participants believe that that manufacturer of the phone or the network provider would be the best company to provide the recommender service, because they have a clear motive to make the phone usage experience better (and they would have no ulterior motive). 4 participants trust the current company (Appy), whereas 3 participants do not.
- When a certain request does not apply to the user, they will skip the request. 6 participants skipped a request that did not apply to them, even when the request had a specific “does not apply to me” answer. They believe that skipping the request is equivalent to indicating this explicitly. However, 1 participant explicitly said she would answer a question about sports with “I do not follow any sports” in order to not get any sports recommendations.

B.9 Privacy and responsibility

- Most users believe that online companies have too much information about them. 10 participants think this is the case. However, 9 participants believe that they are more careful about their privacy than others.
- Most users believe that they are themselves ultimately responsible for their privacy. 7 participants said that by choosing to use the app, they were themselves responsible for their privacy. However, 3 participants made the side note that the company needed to be clear about its practices. Similarly, while 3 participants mention that they do not have enough control over their online privacy, 3 other participants find it unrealistic to expect further control over their privacy once they choose to use an app.

B.10 Satisfaction and usage intention

- Users are generally positive about the idea of a mobile app recommender system. 15 participants like the idea, and said that they would use it. 7 participants said that their current phone lacks a good way to find apps. 5 of them would recommend Applause to others if it worked well, and 3 of them would even pay for using it.