ISR Institute for Software Research
University of California, Irvine

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

**Bart Piet Knijnenburg**
University of California, Irvine
bart.k@uci.edu

**Alfred Kobsa**
University of California, Irvine
kobsa@uci.edu

April 4, 2012
**ISR Technical Report # UCI-ISR-12-1**

**Institute for Software Research**
ICS2 221
University of California, Irvine
Irvine, CA 92697-3455
**www.isr.uci.edu**

# 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. However, 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 often do not hold up well across studies. This paper 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

**Technical Report:**

Released on April 4, 2012 by the Institute for Software Research.

Report number: UCI-ISR-12-1 http://www.isr.uci.edu/tech_reports/UCI-ISR-12-1.pdf

## 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 [Adomavicius and Tuzhilin 2011]. Privacy research however indicates that quite a few people feel uncomfortable disclosing demographal 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].

---

This work has been funded by Samsung US R&D Center.

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

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 2002a; 2002b; Polat and Du 2003; 2005]. 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].

However, decision-making is an inherently complex problem, 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 activity 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 lack knowledge about both the positive and the negative consequences of disclosure [Acquisti and Grossklags 2005; 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 ordering 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. Others study users' general privacy concerns, but disregard their impact on disclosure behavior. 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. INTEGRATING EXISTING APPROACHES TO PRIVACY

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] for more comprehensive surveys), but

to provide a foundation for the subsequently discussed study and 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]. The Westin Privacy Scale (WPS) uses four items 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. 1974]. In this light, Culnan [1993] used 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 process, 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] showed 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 [Kobsa and Teltzrow 2005], 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 given a reason for the request. 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 and Larose [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 as a personal trait (as discussed in section 2.1) is a surprisingly bad predictor of disclosure behavior [Spiekermann and Grossklags 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

Not many researchers on information disclosure consider perceptual and experiential aspects of the systems they evaluate (Hui et al. [2006] and Xu et al. [2009] are notable exceptions). Strategies aimed at influencing users' disclosure behavior may have unforeseen effects on their perceptions and experiences, and these effects may cancel out or even negate the intended effects on disclosure. Likewise, users' perceptions and experiences may mediate the effect of users' personal privacy preferences on disclosure. 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 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 (SC and PC) can influence the perceptions, experience, and interaction.

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. The second problem is that the influence of both personal traits and system characteristics on information disclosure varies extensively from system to system. The Knijnenburg et al. [2012] framework allows us to remedy both problems. First, it integrates the effects of privacy as a personal trait (PC) and of design characteristics (OSA) on information disclosure decisions (INT) in a single model. Moreover, it describes these effects as mediated by system-specific perceptions (SSA) and experiences (EXP). 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.

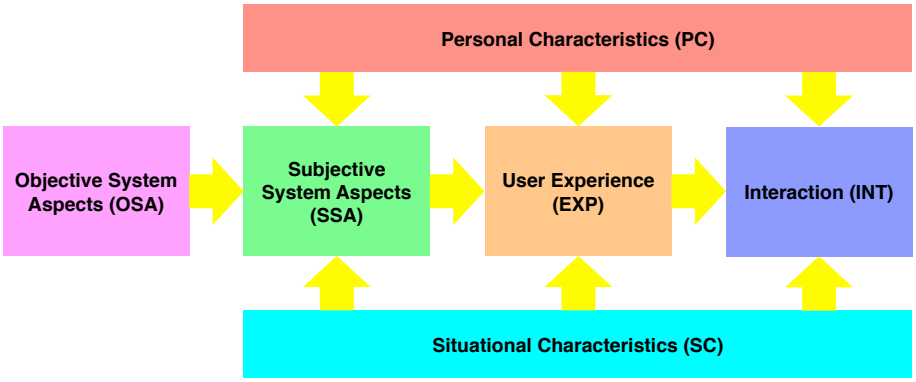


Fig. 1. Framework for user-centric evaluation of recommender systems (Knijnenburg et al. [2012]).

## 3. ONLINE EXPERIMENT

Based on the aforementioned prior work we present a unified approach to studying privacy decisions, which integrates privacy-related concepts into the Knijnenburg et al. [2012] framework. In terms of personal characteristics (PC), our main interest lies in general privacy concerns, collection concerns and control concerns [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 [Kobsa and Teltzrow 2005; Besmer et al. 2010; Patil et al. 2011; Acquisti et al. 2011] and request order [Acquisti et al. 2011]. In order to improve the robustness of the effect of these factors on information disclosure, we consider the experience (EXP) variables 'trust in the company' (cf. [Jarvenpaa and Tractinsky 1999; Metzger 2004]) and 'satisfaction with the system' (cf. [Xu et al. 2009; Xu et al. 2011; Hui et al. 2006]), as well as 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]).

### 3.1 System

We apply these manipulations and measurements to a study of an online mockup of a mobile app recommender system with the working title 'Applause'. The system was inspired by app recommenders 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). We conducted a preliminary analysis of the effects in [Knijnenburg and Kobsa 2012], with a focus on personalizing the strategies for improving information disclosure and/or system satisfaction. For the current paper we gathered an additional sample of data, and we shifted our focus to an integrative model of privacy decision-making, using our unified approach.

The Applause 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). 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 Appy[1]. We reinforced this belief by ostensibly transferring users to the Appy website (with its own URL and branding) for the disclosure part of the experiment (Fig. 2).

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

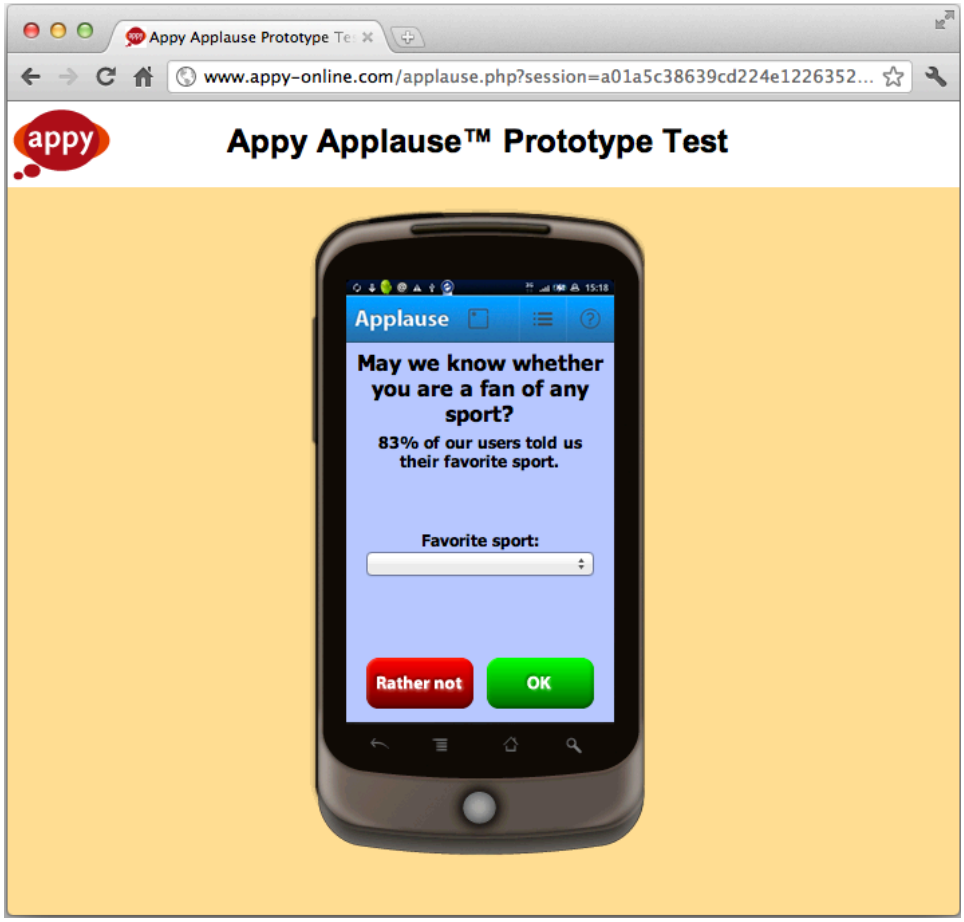


Fig. 2. The website of Appy with the Applause system mockup. On this website the participants do the disclosure part of the experiment

### 3.2 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 that provided 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 dataset, 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 dataset. The optimal model found in the exploration phase was tested on this set, and any inconsistent effects were removed from the model. The final model is based on the data from all 491 participants. 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 system[2]. 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 would be transferred back to the experimenters' website, where they were asked about their personal privacy concerns and their subjective and experiential evaluation of the system.

### 3.3 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 have 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.

The justification messages (see Table I) are tested against the baseline system with no justification message. The 'useful for you' and 'useful for others' justifications explain the benefits of disclosure (cf. [Wang and Benbasat 2007]) in two different ways. The 'number of others' justification appeals to the social norm (cf. [Besmer et 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. [Kobsa and Teltzrow 2005]). Note that the percentages in the messages were randomly chosen from 5% to 95% (in Knijnenburg and Kobsa [2012] we show that this percentage has barely any effect, so 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 first[3] (see Table I).

Table I. Experimental manipulations: strategies to influence information disclosure

| Manipulation | Conditions | Description |
|---|---|---|
| Justification type | 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]" |
| Message 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. |

[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.

[3] To facilitate users' decision-making, we grouped similar items within each category of requests. Within the demographical data requests we furthermore manipulated the order of four subgroups of items (leisure, personal, family, background), but this manipulation had no effect on disclosure. Furthermore, this grouping did not lead to a better factor structure than our two-factor solution.

### 3.4 Measures

The main dependent variable in the 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[4] (CFA) with dichotomous indicators and a weighted least squares estimator, estimating two factors: 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 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.

Table II. Items used to measure participants' disclosure behavior

| Type of data | Items | Level of disclosure | Factor loading |
|---|---|---|---|
| **Context** | Recommendation browsing | 87.0% | |
| | Location | 84.8% | 0.767 |
| Alpha: 0.79 | Phone model | 84.6% | 0.659 |
| AVE: 0.652 | App usage | 82.2% | 0.749 |
| | App usage time | 73.2% | |
| Factor correlation: 0.432 | 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** | Gender | 94.9% | |
| | Amount of reading | 93.5% | |
| Alpha: 0.86 | Age | 93.3% | |
| AVE: 0.784 | Education | 92.7% | |
| | News interests | 92.7% | |
| Factor correlation: 0.432 | 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

[4] For an introduction to Confirmatory Factor Analysis as applied in this paper, see Appendix A of [Knijnenburg et al. 2012]

in Table III. We subjected the items to a factor analysis with 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)<br><br>Alpha: 0.76<br>AVE: 0.774<br><br>Based on [Malhotra et al. 2004; Smith et al. 1996] | All things considered, the Internet causes serious privacy problems | |
| | Compared to others, I am more sensitive about the way online companies handle my personal information | 0.785 |
| | To me, it is the most important thing to keep my privacy intact from online companies | 0.708 |
| | I believe other people are too concerned with online privacy issues | |
| | I am concerned about threats to my personal privacy today | 0.824 |
| **Collection concerns** (PC)<br><br>Alpha: 0.86<br>AVE: 0.815<br><br>Based on [Malhotra et al. 2004; Smith et al. 1996] | It usually bothers me when online companies ask me for personal information | 0.860 |
| | When online companies ask me for personal information, I sometimes think twice before providing it | |
| | It bothers me to give personal information to so many online companies | 0.829 |
| | Online companies may collect any information about me because I have nothing to hide (new) | -0.749 |
| | I'm concerned that online companies are collecting too much personal information about me | 0.855 |
| | I'm not bothered by data collection, because my personal information is publicly available anyway (new) | 0.860 |
| **Control concerns** (PC)<br><br>Alpha: 0.58<br>AVE: 0.526<br><br>Based on [Malhotra et al. 2004] | 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) | |
| | 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) | |
| **Perceived value of disclosure help** (SSA)<br><br>Alpha: 0.75<br>AVE: 0.581<br><br>Inspired by [Wang and Benbasat 2007] | The system helped me to decide what information I should disclose | 0.788 |
| | The system explained how useful providing each piece of information was | 0.633 |
| | The system helped me to make a tradeoff between privacy and usefulness | 0.851 |
| | I felt clueless about what information to disclose | |

| | | |
|---|---|---|
| **Perceived privacy threats** (SSA)<br><br>Alpha: 0.71<br>AVE: 0.529<br><br>Inspired by [Xu et al. 2009; 2011] | The system made me more cautious than usual disclosing this type of information | |
| | The system helped me to protect my privacy | |
| | The system has too much information about me | 0.909 |
| | The system does not know anything I would be uncomfortable sharing with it | -0.608 |
| | 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 |
| **Trust in the company** (EXP)<br><br>Alpha: 0.93<br>AVE: 0.845<br><br>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 | |
| | Providing this company my information would involve many unexpected problems | |
| | I feel safe giving my information to this company | |
| **Satisfaction with the system** (EXP)<br><br>Alpha: 0.91<br>AVE: 0.722<br><br>Based on [Knijnenburg et al. 2012] and inspired by [Hui et al. 2006; Xu et al. 2009; 2011] | 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 | |

### 3.5 Results

We subsequently subjected the disclosure behaviors, the subjective evaluations, and the manipulated system aspects to Structural Equation Modeling[5] (SEM), which simultaneously fits the measurement model and the structural relations between measured constructs. To avoid over-fitting, we constructed the 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. The final model (Fig. 3) has a good[6] model fit: $\chi^2(912) = 1540$, $p < .001$; *RMSEA* = 0.037, 90% CI: [0.034, 0.041], *CFI* = 0.977, *TLI* = 0.976.

[5] For an introduction to Structural Equation Modeling as applied in this paper, see Appendix A of [Knijnenburg et al. 2012].

[6] Theoretically, a good model is not statistically different from the fully specified model ($p > .05$). However, this statistic is commonly regarded as too sensitive [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.

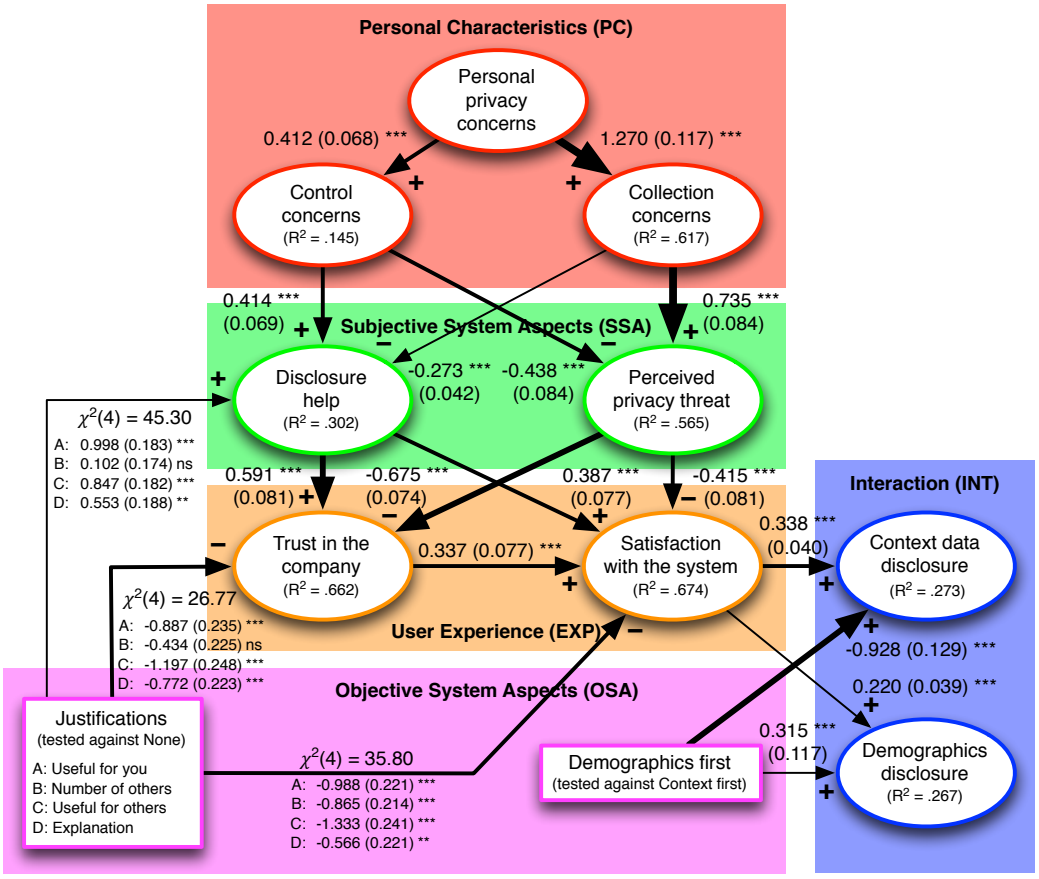


Fig. 3. The Structural Equation Model (SEM) for the data of the experiment.[7] (Significance levels: *** $p < .001$, ** $p < .01$, 'ns' $p > .05$)

The model shows that aside from the request order, all effects on context and demographics disclosure are mediated by users' satisfaction with the system. The higher their satisfaction, the more inclined they are to disclose information. Satisfaction is higher for participants that trust the company, feel helped in their disclosure, and perceive a low level of privacy threat. Trust in the company itself is also higher for participants that feel helped in their disclosure and perceive a low level of privacy threat.

In terms of personal characteristics, general personal privacy concerns drive control and 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.

The justifications have a significant impact on perception of disclosure help, trust in the company, and satisfaction with the system. The 'useful for you', 'useful for

[7] Factors are represented as circles (the indicators are left out for clarity; please refer to Table II and Table III). $R^2$ values represent the proportion of variance explained by the model. Numbers on the arrows (as well as their thickness) represent the $\beta$ coefficients (and standard error) of the effect represented by the arrow. Factors are scaled to have an SD of 1. For any arrow A → B, one SD difference in A thus causes a $\beta$ SD difference in B. The $\chi^2$ values test the effect of all justifications simultaneously; the $\beta$ coefficients below the $\chi^2$ values represent the effect (in $\beta$ SD difference) of each justification tested against the baseline of 'no justification'. The $\beta$ coefficients on the 'demographics first' arrow represent the effect of the 'demographics first' request order (in $\beta$ SD difference) against the 'tracking first' request order.

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 satisfaction with the system. Fig. 4 shows that the total (direct plus mediated) effects of the justifications on trust in the company are essentially zero, and that the total effects on satisfaction with the system and disclosure behavior are negative.

Finally, the request order has a direct impact on disclosure behavior. 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.

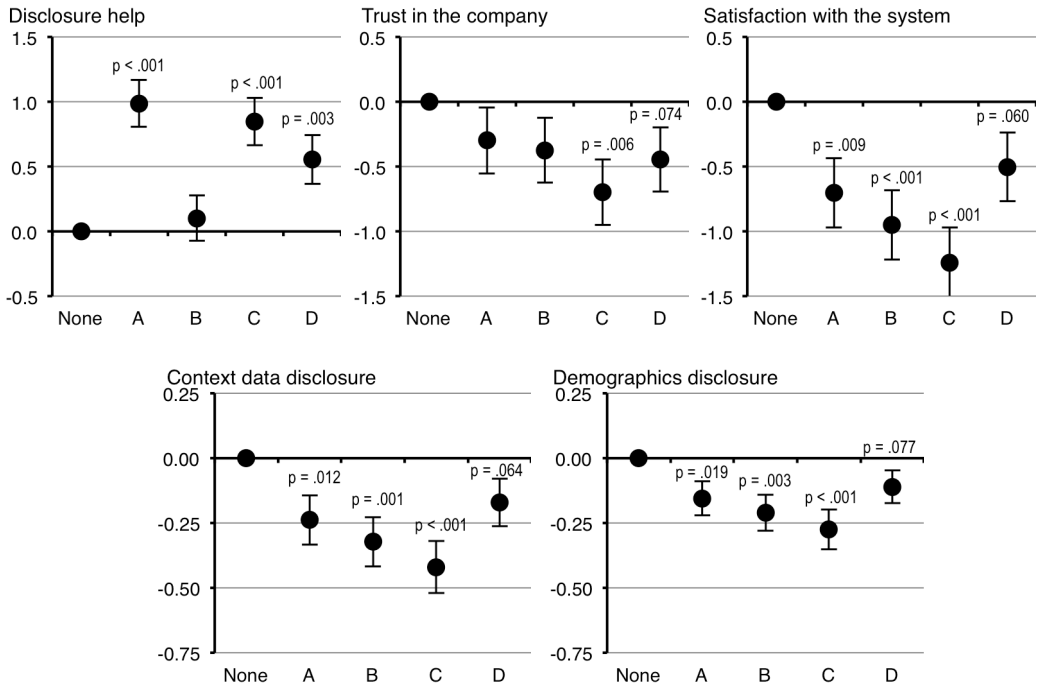


Fig. 4. The total effects of the justifications (A: Useful for you, B: Number of others, C: Useful for others, D: Explanation) on the different outcomes, tested against the baseline condition (No justification). Error bars are ±1 SE.

## 4. 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 this process. In this section we reflect on these results by integrating them with findings from an interview study with 17 participants that used a paper prototype of the recommender system. In the interviews we asked participants to elaborate on their disclosure decisions.

### 4.1 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 satisfaction with the system.*
In the experimental model, almost all effects on information disclosure are mediated by 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 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 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 (11x), the company selling their information (12x) and security concerns or other unintended breaches of confidentiality (all 17 participants).

### 4.2 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 2012] 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 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 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. $\beta = -0.139$)[8] are roughly equal but opposite.

*General privacy concerns cause both control and collection concerns, but have a total negative impact on 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 satisfaction ($\beta = -0.588$), context data disclosure ($\beta = -0.199$), and demographics disclosure ($\beta = -0.130$)[9] 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.

### 4.3 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. 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".

[8] For these total effects, $p$-values are $< .001$.
[9] For these total effects, $p$-values are $< .001$.

*Except for the 'number of others' justification, the justifications decrease users' trust in the company, and all justifications decrease users' 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' satisfaction, and as a result the total effects are also negative. 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. 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 2012], 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 2012] 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 that 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.

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.

## 5. 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 of a recommender system, even though 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 based on actual system results. 16 out of 17 participants in the interview study declared that this would be a strategy they would follow; future research should explore this dynamic behavior.

Regardless of these reservations, 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 explore in [Knijnenburg and Kobsa 2012]. 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.

## REFERENCES

ACKERMAN, M.S., CRANOR, L.F., AND REAGLE, J. 1999. Privacy in e-commerce: examining user scenarios and privacy preferences. *Proceedings of the 1st ACM conference on Electronic commerce*, 1–8.

ACQUISTI, A. AND GROSSKLAGS, J. 2005. Privacy and Rationality in Individual Decision Making. *IEEE Security & Privacy 3*, 1, 26–33.

ACQUISTI, A. AND GROSSKLAGS, J. 2008. What Can Behavioral Economics Teach Us About Privacy? *Digital Privacy: Theory, Technologies, and Practices*, 363–377.

ACQUISTI, A., JOHN, L., AND LOEWENSTEIN, G. 2009. What is privacy worth? *Twenty First Workshop on Information Systems and Economics*.

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.

ADOMAVICIUS, G. AND TUZHILIN, A. 2011. Context-Aware Recommender Systems. In: F. Ricci, L. Rokach, B. Shapira and P.B. Kantor, eds., *Recommender Systems Handbook*. Springer US, Boston, MA, 217–253.

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. *Proceedings of the Sixth Symposium on Usable Privacy and Security - SOUPS '10*.

BÖHMER, M., BAUER, G., AND KRÜGER, A. 2010. Exploring the Design Space of Context-aware Recommender Systems that Suggest Mobile Applications. *2nd Workshop on Context-Aware Recommender Systems*.

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.

CANNY, J. 2002a. Collaborative Filtering with Privacy via Factor Analysis. *25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2002)*, 238–245.

CANNY, J. 2002b. Collaborative Filtering with Privacy. *IEEE Symposium on Security and Privacy*, 45–57.

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.

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.

DAVIDSSON, C. AND MORITZ, S. 2011. Utilizing implicit feedback and context to recommend mobile applications from first use. *Proceedings of the 2011 Workshop on Context-awareness in Retrieval and Recommendation - CaRR '11*, 19–22.

EGELMAN, S., TSAI, J., CRANOR, L.F., AND ACQUISTI, A. 2009. Timing is everything?: the effects of timing and placement of online privacy indicators. *Proceedings of the 27th international conference on Human factors in computing systems*, 319–328.

VAN DE GARDE-PERIK, E., MARKOPOULOS, P., DE RUYTER, B., EGGEN, B., AND IJSSELSTEIJN, W. 2008. Investigating Privacy Attitudes and Behavior in Relation to Personalization. *Social Science Computer Review 26*, 1, 20–43.

GIGERENZER, G. AND GOLDSTEIN, D.G. 1996. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review 103*, 4, 650–669.

GIRARDELLO, A. AND MICHAHELLES, F. 2010. AppAware: which mobile applications are hot? *Mobile HCI 2010*, ACM Press, 431–434.

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., ASSOCIATES, AND WESTIN, A.F. 2003. *Consumer Privacy Attitudes: A Major Shift Since 2000 and Why*. Harris Interactive, Inc.

HASSENZAHL, M. 2005. The thing and I: understanding the relationship between user and product. *Funology*, 31–42.

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.

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.

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.

JARVENPAA, S. AND TRACTINSKY, N. 1999. Consumer Trust in an Internet Store: A Cross-Cultural Validation. *Journal of Computer Mediated Communication 5*, 2, 1–36.

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.

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.

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.
KITTUR, A., CHI, E.H., AND SUH, B. 2008. Crowdsourcing user studies with Mechanical Turk. ACM Press, 453–456.
KNIJNENBURG, B. AND KOBSA, A. 2012. Helping users with information disclosure decisions: potential for adaptation. *Submitted*, http://www.ics.uci.edu/~kobsa/papers/KnijnenburgKobsa–UMAP2012Submit.pdf.
KNIJNENBURG, B.P., WILLEMSEN, M.C., GANTNER, Z., SONCU, H., AND NEWELL, C. 2012. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction 22*, 4-5, forthcoming, http://www.usabart.nl/portfolio/KnijnenburgWillemsen–UMUAI2011_UIRecSy.pdf.
KOBSA, A. 2007. Privacy-enhanced web personalization. *The adaptive web*, 628–670.
KOBSA, A. AND TELTZROW, M. 2005. Contextualized communication of privacy practices and personalization benefits: Impacts on users' data sharing and purchase behavior. *Privacy Enhancing Technologies*, 329–343.
KOLTER, J. AND PERNUL, G. 2009. Generating User-Understandable Privacy Preferences. *2009 International Conference on Availability, Reliability and Security*, 299–306.
LAUFER, R.S., PROSHANSKY, H.M., AND WOLFE, M. 1974. Some Analytic Dimensions of Privacy. In: R. Kuller, ed., *Some Analytic Dimensions of Privacy*. Dowden, Hutchinson & Ross, Stroundsbourg, PA.
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.
LEDERER, S., MANKOFF, J., AND DEY, A.K. 2003. Who wants to know what when? privacy preference determinants in ubiquitous computing. ACM Press, 724–725.
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.
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 (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research 15*, 4, 336–355.
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.
NISSENBAUM, H.F. 2010. *Privacy in context : technology, policy, and the integrity of social life*. Stanford Law Books, Stanford, Calif.
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.
PATIL, S., PAGE, X., AND KOBSA, A. 2011. With a little help from my friends: can social navigation inform interpersonal privacy preferences? *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, 391–394.
POLAT, H. AND DU, W. 2003. Privacy-Preserving Collaborative Filtering. *International Journal of Electronic Commerce 9*, 4, 9–35.
POLAT, H. AND DU, W. 2005. SVD-based Collaborative Filtering with Privacy. *ACM Symposium on Applied Computing*, 791–795.
PU, P., CHEN, L., AND HU, R. 2011. A user-centric evaluation framework for recommender systems. *Proceedings of the fifth ACM conference on Recommender systems*, ACM Press, 157–164.
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.
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.
SOLOVE, D.J. 2006. A Taxonomy of Privacy. *University of Pennsylvania Law Review 154*, 3, 477–564.
SPIEKERMANN, S. AND GROSSKLAGS, J. 2001. E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus actual Behavior. *EC'01: Third ACM Conference on Electronic Commerce*, 38–47.
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.
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.
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. *Proceedings of the 37th Hawaii International Conference on System Sciences*, 181–190.
THALER, R. 1980. Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization 1*, 1, 39–60.

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*.

TUROW, J., KING, J., HOOFNAGLE, C.J., BLEAKLEY, A., AND HENNESSY, M. 2009. Americans Reject Tailored Advertising and Three Activities That Enable It. *SSRN eLibrary*.

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. AND LOUIS HARRIS AND ASSOCIATES. 1981. *The Dimensions of privacy : a national opinion research survey of attitudes toward privacy*. Garland Pub., 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.

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.

ZHANG, N., WANG, C., AND XU, Y. 2011. Privacy in Online Social Networks. *Proceedings of the 2011 International Conference on Information Systems*, Paper 3.