

Investigating Privacy Attitudes and Behavior in Relation to Personalization

Evelien van de Garde-Perik

Panos Markopoulos

Eindhoven University of Technology, the Netherlands

Boris de Ruyter

Philips Research Eindhoven, the Netherlands

Berry Eggen

Wijnand IJsselsteijn

Eindhoven University of Technology, the Netherlands

This article presents an experimental study of privacy-related attitudes and behaviors regarding a music recommender service based on two types of user modeling: personality traits and musical preferences. Contrary to prior expectations and attitudes reported by participants, personality traits are frequently disclosed to the system and other users, indicating that embedded modeling of user personality does not represent an acceptance barrier. Discrepancies between privacy attitudes and behaviors have been reported before in the context of e-commerce applications, but the corresponding studies could not exclude several conflicting hypotheses, such as participants expressing attitudes outside the context of specific privacy dilemmas and contact with researchers, which may have mitigated perceived privacy risks. Arguably, these are fundamental problems in empirical investigations into privacy that apply to most published works relating to privacy and user modeling. Measures to control these factors in this study are discussed, and methodological suggestions for future research are presented.

Keywords: *privacy; personalization; music recommender; user profile*

Adaptive systems and personalized systems in particular rely on having appropriate and sufficient information about their users to operate optimally. This could, for example, include information about the identity of the user, earlier usage of a service, the user's preferences and dislikes, and many other types of data (Kobsa, 2001). The collection and processing of such information can conflict with privacy concerns (Kobsa, 2002). More specifically, it has been suggested that privacy-related concerns are contingent on the kind of information collected (Ackerman, Cranor, & Reagle, 1999; Adams & Sasse, 2001; Bellotti & Sellen, 1993), the degree of control users have over disclosure (Bellotti & Sellen, 1993; Günther & Spiekermann, 2005; Olivero & Lunt, 2004), the degree of accessibility (Adams & Sasse, 2001; Bellotti & Sellen, 1993), or the way the information is used (Adams & Sasse, 2001; Bellotti & Sellen, 1993). Information may be collected by the user's explicit and intentional input, or information may be collected implicitly by monitoring the users as they go about their own tasks or activities. Implicit collection of information is an

essential element of scenarios relating to Ambient Intelligence (Aarts, Harwig, & Schuurmans, 2001) and can lead to privacy issues because of a lack of awareness and control by the people concerned (Cranor, 2004; Kobsa & Schreck, 2003; Nguyen & Mynatt, 2002).

In the context of this study, the term *privacy* refers to a boundary control process in which individuals regulate when, how, and to what extent information about them is communicated to others. This interpretation is based on the definitions by Westin (1967) and Altman (1975). In this study, the emphasis is on information privacy rather than interpersonal or social privacy. *Information privacy* refers to the claim by individuals that data about themselves should generally not be made available to other individuals and organizations and that, where data are possessed by another party, the individual must be able to exercise a substantial degree of control over that data and their use (Clarke, 1999).

Several technological solutions to guarantee people's privacy have been proposed (for a relatively recent survey, see Langheinrich, 2005). Other researchers have focused on providing design guidance to designers of Internet-based applications (e.g., Ackerman et al., 1999; Good & Krekelberg, 2003) and ubiquitous computing applications (Bellotti & Sellen, 1993; Langheinrich, 2002; Lederer, Hong, Dey, & Landay, 2004; Nguyen & Mynatt, 2002).

In contrast, there is still little known about how users of personalized systems experience disclosure of information, what motivates them, and how conscious they are of the choices they make concerning the disclosure of information.

Acquisti (2004) has applied behavioral economics to understand the individual decisionmaking process with respect to privacy in e-commerce. He has described how incomplete information and psychological distortions may prevent an individual from behaving rationally with regard to privacy.

Privacy appears to be an inherently difficult concept to study. For example, it seems that the privacy concerns people report during surveys do not match their actual behavior (see Berendt, Günther, & Spiekermann, 2005). In their study, participants were first questioned about their privacy preferences and afterward exposed to an anthropomorphic shopping "bot." The shopping bot asked various questions about product attributes, but it also asked more personal questions than customers should be expected to answer. Contrary to the researchers' expectations, even participants who were classified as highly concerned about their privacy according to their reported attitudes displayed a strong tendency to self-disclose.

There can be several explanations for this discrepancy. In general, attitudes are not good predictors of the behavior that is mediated by social and environmental factors and the costs that need to be incurred to exhibit behavior compliant with the attitudes expressed (see Ajzen & Fishbein, 2005). Furthermore, as already known from experiments by Milgram (1974), participants can engage in behaviors they would not normally engage in under the authority of the experimenter. An alternative explanation is that participants perceived the research study as a safe context for disclosure, trusting researchers to moderate the impact of disclosure.

These methodological problems also occur in many studies of privacy in the domain of computer-mediated communications, Internet use, and Ambient Intelligence. Studies of privacy restricted to laboratory experimentation or surveying opinions lack external validity in the absence of actual risk and outside a realistic context of use. Field observations of the actual use of a system can disguise many of the privacy concerns people may have

despite demonstrated risk-taking behavior and do not offer sufficient controls for the privacy risks and the benefits of the disclosure.

This article presents an experimental study that was set up to provide users with realistic privacy dilemmas during the use of a personalized music recommender service that was created specifically for the purpose of this study. Participants were presented with disclosure choices that were as realistic as possible. Logging of their behavior in these situations was combined with self-report and attitude measurements regarding privacy. The study enabled us to test some of the assumptions regarding privacy and personalized systems and provided an opportunity to explore the motivation behind observed privacy behavior.

First of all, the methodology used will be described, then the findings will be presented, and this is followed by a discussion of the results in the light of other research. Finally, the conclusions of this study and implications for future work are addressed.

The Music Recommender Study

This study involved participants using a music recommender service accessed through a web application over the Internet. Participants used this service from their own computers at work or at home, similar to any other Internet-based music recommender service. The system confronted the participants with privacy dilemmas relating to two types of personal information: preferences for different music genres and information about their personality. Participants were asked to disclose this information to the recommender service and/or other users of this service. Rentfrow and Gosling (2003) have established a relationship between music preferences and personality traits, which makes it feasible to recommend music based on personality traits. For example, someone who scores high in terms of the personality trait “neuroticism” as well as “openness to new experiences” is likely to be interested in the “reflective and complex” music dimension. Music genres that fall into this music dimension are blues, jazz, folk, and classical music (Rentfrow & Gosling, 2003).

1. The interest in contrasting these music preferences and personality traits comes from the relative legitimacy of these two approaches to user profiling for the purpose of recommending music. According to Iachello and Abowd (2005), an application is legitimate if the interest in using it for a specific purpose justifies the burden on individuals’ rights. Music preference data are seemingly innocuous but directly relevant to the music recommender and are commonly offered by users to systems of this kind. However, music preferences can be used for direct marketing, so they could also be considered to be somewhat sensitive by some users.

Personality traits are more or less stable, internal characteristics of people, which make their behavior consistent from one time to another yet different from the behavior that other people would exhibit in comparable situations (Child, 1968). As a user modeling approach, personality traits have a lot of potential because they are domain independent and may be applied in a much wider range of contexts than music recommendation. Conversely, a personality profile of a user can be misused and misapplied in a variety of contexts (e.g., by a prospective employer or a medical insurance provider). Storage of personality traits seems to have little legitimacy in the context of music recommendation and does therefore not

adhere to the principle of proportionality. This means that it does not balance the usefulness of the application and its effect on privacy (Iachello & Abowd, 2005). During the experimental study, participants were confronted with choices regarding whether to disclose personal information and whether disclosure would be anonymous. It was anticipated that personality trait information would be regarded as more sensitive, resulting in comparatively less disclosure than for music preferences. It was also anticipated that showing profile information directly to other users would result in less disclosure than if the profile information were used only by the recommender system and not shown to anyone else. It is important to realize that this study is of an explorative nature, combining quantitative and qualitative methods. The underlying attitudes were set out to check whether existing instruments for assessing privacy attitudes would be good predictors of actual behavior. Furthermore, context-specific privacy preferences were surveyed and explanations for disclosure behavior considered; these were then analyzed qualitatively to explore the relationships between them.

Experimental Design

The study followed a within-subjects design whereby participants were exposed to profiling of both music preferences and personality traits. The order in which participants used the two recommender systems was counterbalanced. Participants could use the recommender systems at home or at work.

The study involved four disclosure choice moments. Two disclosures related to profile information and the other two related to personality traits. In both cases, participants were first asked to share information for comparison by the system (as in collaborative filtering) and to then share the current profile information with other users. In each case, participants could choose between three levels of disclosure at these choice moments: no disclosure, anonymous disclosure, or disclosure including identity information.

Data of both a qualitative and quantitative nature were collected: Actual disclosure choices were monitored by system logs, explanations for these choices were gathered through questionnaires and interviews, and attitudes toward privacy and the use of personal information were measured by means of a questionnaire.

Participants

Participants were recruited by e-mail announcements via secretaries and bulletin boards within a technical university and an industrial research lab. In view of the database of music available, recruitment was aimed at participants ages 18 to 50. Forty-eight participants took part in the study and completed the online questionnaire. As a reward, participants were promised a music CD selected on the basis of the songs recommended to them. This was done to motivate participants to be honest about their music preferences and personality traits; it served to increase the realism of both the costs and benefits of the disclosure they would make during the experiment.

The ages of participants ranged from 17 to 49, with an average age of 26. Half of participants were 23 years old or younger. Twenty-one participants were interviewed. Interviewee ages ranged from 19 to 49, with an average age of 27.

Apparatus and Materials

The music recommendations were provided via a web-based application created for the purposes of this study. The service offered personalized playlists of songs. Streaming technology was used to make these songs available for participants to play on their computers. The experimental recommender service is built on a database of nearly 6,000 songs spread evenly over 14 different music genres.

While using the personality-based recommender, participants saw a screen like the one shown in Figure 1. At the top of the screen, a status bar showed the number of times the user had logged into the service, the number of playlists requested, the source of information on which the recommendation was based, and the progress within the study. On the left-hand side of the screen, an overview of the current profile information based on personality traits was shown. The right side of the screen displayed the current playlist. Below this playlist, a button was displayed; this could be used to request a new playlist. The screen of the recommender based on music preferences was similar to the screen shown in Figure 1, except that the status part stated, "The recommender is using your preferences to generate playlists." On the left side of the screen, an overview of the music preference profile was shown.

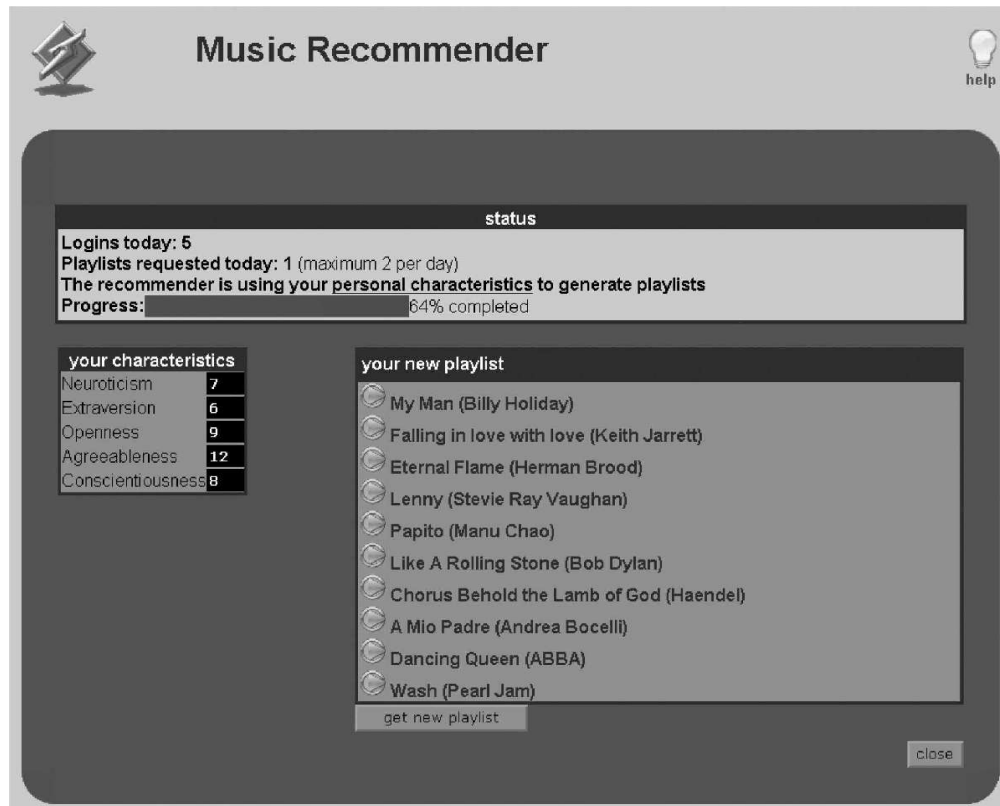
Procedure

The procedure of the experimental study is shown schematically in Figure 2. People who were interested in participating were sent instructions by e-mail. They were not informed in advance that the research related to privacy. Participants had to register on the music recommender Web site by providing their e-mail address. This was their business/university e-mail address, which consisted of their name and company. At the time of registration, participants were assigned randomly to an experimental group (linked to a specific order in which they would experience the two recommender systems) and received a personal access code.

Participants were asked to access the portal site of the music recommender on six separate days to ensure that the participants would not rush through all phases of the study. They were asked to listen to at least one playlist per day. If a participant had listened to a playlist the previous day, the participant was offered the next phase. The music recommender therefore served two purposes: It operated as a music recommender service, and it also implemented the experimental procedure. The participants were asked to use the music recommender for 6 days within a period of 2 to 3 weeks.

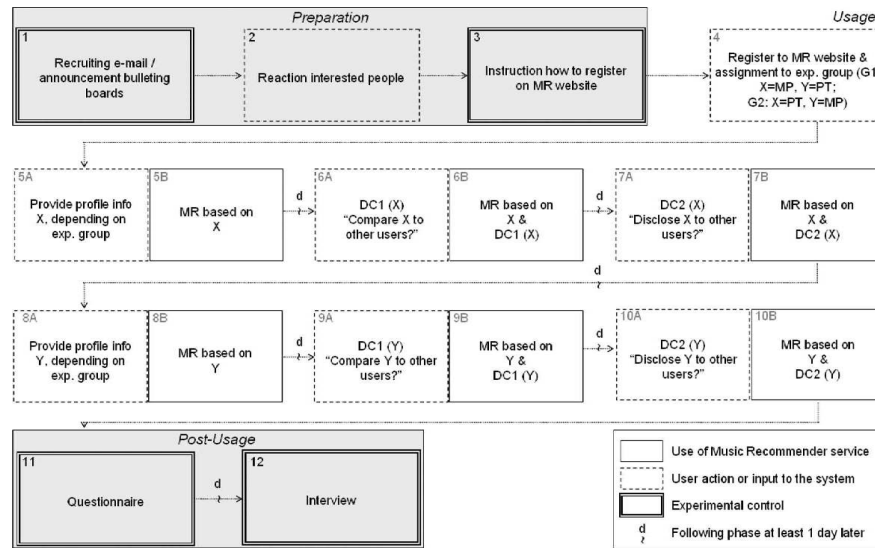
As indicated earlier, the order in which participants experienced the two recommender systems was counterbalanced. The first phase of both music recommender systems (5/8 in Figure 2) consisted of a default situation in which no additional disclosure of information was required (local use of the information only). In the second phase (6A/9A in Figure 2), participants were asked to disclose their profile information for the purpose of collaborative filtering (for the participants, this was referred to as "comparing preferences to those of others users" or "comparing personality traits to those of others users"). In the third phase (7A/10A in Figure 2), participants were asked to disclose their profiles directly to other users. In this way, participants were asked to reveal an increasing amount of information.

Figure 1
Screenshot of Music Recommender System Based on Personality Traits



In each of the three phases of using the two recommender systems, the actual recommendation mechanism remained constant, but the percentage of recommendations that were offered according to the profile was increased if users chose to disclose their profile information. This was done to ensure that the recommender performance would improve predictably when the user chose a higher level of disclosure. This recommendation mechanism was not revealed to participants. Instead, participants were told in advance that the recommendations might improve as a result of their disclosure. Without more information about the recommendation mechanism, they would suppose that the data would actually be used for collaborative filtering. In the first phase of both recommender systems, 80% of the recommendations were generated according to the user's profile. Depending on the choices made by the participant in the second and third phase (regarding whether to disclose and whether disclosure would be anonymous), the recommendation could improve, so that 90% or 100% of the recommendations were generated according to the profile. After they had used the music recommender, participants were asked to complete an online questionnaire

Figure 2
Flowchart of Phases in Music Recommender Study



Note: MR = music recommender; G = experimental group; MP = music preferences; PT = personality traits; DC = disclosure choice; X = first source for recommendations; Y = second source for recommendations.

with a combination of open and multiple-choice questions. Again, it was not made clear to the participants that the focus of the study was privacy.

Finally, 21 out of 48 participants were contacted to arrange an interview appointment after the completion of the online questionnaire. The interview was set up in such a way as to achieve an open atmosphere in which participants would express their feelings freely. Interviewees were questioned in more detail about their choices during the study and encouraged to give thorough explanations about the answers they had given in the online questionnaire.

Measures

Before participants could actually receive personally recommended playlists, they had to provide their first type of profile information (which depended on the experimental group to which they were assigned). The second type of profile information was collected only after 3 days of using the first recommender system (see Figure 2). The profile information was collected by means of two short, validated measures: the Short Test of Music Preferences (STOMP) and the Ten Item Personality Measure (TIPI).

The STOMP test (Rentfrow & Gosling, 2003) inquires about the basic preference level for 14 different music genres on a scale from 1 (*strongly dislike*) to 7 (*strongly like*). It asks people to:

“please indicate your basic preference level for the genres listed using the scale provided.”

Example genres are country, jazz, or rock. The STOMP values were used as a basis for music recommendations in the music-preference-based application.

The TIPI (Gosling, Rentfrow, & Swann, 2003) encompasses a personal judgment of the extent to which 10 pairs of personality traits apply on a scale from 1 (*disagree strongly*) to 7 (*agree strongly*). It gives people the following task description:

“Here are a number of personality traits that may or may not apply to you. Indicate (. . .) the extent to which you agree or disagree with that statement. You should rate the extent to which the pair of traits applies to you, even if one characteristic applies more strongly than the other.”

Examples of pairs of personality traits used in TIPI are “reserved/quiet” or “sympathetic/warm.” The TIPI scores were converted to the Big Five personality dimensions (i.e., extraversion, agreeableness, conscientiousness, neuroticism, and openness to experiences) and used as a basis for music recommendations in the personality-traits-based condition. The Big Five personality dimensions are frequently used in personality assessments, although they are usually based on long questionnaires of sometimes more than 200 items that have to be rated. For the purpose of the personality-based music recommender system, it was felt that such a long questionnaire would take up too much time and effort for the participants. The short TIPI test was created for situations where very short measures are needed or where personality is not the primary focus of interest. Both these conditions hold for this study, so the TIPI was chosen as a personality inventory. Furthermore, because the length of TIPI and STOMP are quite comparable, participants were likely to have similar expectations about the use of the two types of profile information by the two different recommender systems (a 200-item questionnaire might raise the expectations of participants with regard to the accuracy of the music recommendations, and it would be likely to raise suspicion about other potential use of the information as well).

While the music recommender was being used, the actual level of disclosure chosen (no disclosure, anonymous disclosure, or disclosure including identity) in the four choice situations was recorded. This measure relates to the main aim of the research: to see whether participants would choose different levels of disclosure depending on the type of information involved (either personality traits or music preferences) and the use of the information (for comparison by the system or for showing directly to other users). Besides their choice of disclosure, participants were also asked to rate the quality of each playlist of recommendations on a 5-point scale. This quality rating was included to assess whether there was any relationship between the perceived benefits of the system and the disclosure behavior of participants.

After using the music recommender for 6 days, all participants were asked to complete an online questionnaire consisting of some open and some multiple-choice questions. This questionnaire was aimed at understanding the factors underpinning the disclosure behavior and it addressed these topics:

- General demographics and interest in music.
- General appreciation of the music recommender.
- Attitudes toward disclosing music preferences and personality traits.
- Explanations for the level of disclosure chosen before and after the study (open question).
These questions were posed to see whether people’s experience of using the recommender

would lead them to make different choices afterward (e.g., because of disappointing benefits or any privacy concerns they had).

- General privacy attitude measures were taken to see whether they would explain any differences in disclosure behavior. Privacy attitudes were measured by different tools: the Privacy Segmentation Index (PSI; Harris Interactive, 2002) and the Privacy Attitude Questionnaire (PAQ; Chignell, Quan-Haase, & Gwizdka, 2003). In addition to these tools, some general questions regarding privacy (e.g., “I like to get advance notice if information is collected about me,” or “I am willing to provide personal information in return for low-cost products or convenience”), and four questions about the worries concerning the disclosure of personal information in different situations were included.

The PSI consists of three statements about the use of personal information by organizations. Participants have to judge the extent to which they agree with each of these three statements. This tool was included because of its brevity and because this type of clustering is widely used in human-computer interaction (HCI) research (Ackerman et al., 1999; Berendt et al., 2005; Consolvo, Smith, Matthews, LaMarca, Tabert, & Powledge, 2005).

The PAQ consists of 36 statements about various behaviors relating to privacy that people may or may not exhibit. Example statements include the following: “No organization or person should disseminate personal information about me without my knowledge,” “I respond to telephone marketing surveys,” or “I like to change my passwords frequently.” Participants had to judge the extent to which they agree with each statement. The tool was originally developed to aid designers, as there is little information they can use as a basis for the design of new technologies and interfaces with privacy implications (e.g., personalization). The PAQ tool was included in this study because of its relevance to the domain of personalization. This also allowed a comparison of the two different tools for measuring privacy attitudes.

In-depth semistructured interviews were conducted, which varied in duration between half an hour and an hour. After the completion of the online questionnaires, participants were given an interview appointment. All pilot participants were contacted and interviewed (see Perik, de Ruyter, Markopoulos, & Eggen, 2004). The remaining participants were contacted at random. Twenty-one interviews were conducted. The aim of the interview was to gain a more thorough understanding of the factors on which disclosure decisions are based and to obtain additional information about some of the answers given in the online questionnaire. The interviews were semistructured and covered the following topics:

- Opinion on the music recommender.
- Expected goal/aim of the research. This topic was addressed in an open fashion without mentioning privacy at all. The topic was included to check whether people were aware that the research related to privacy and whether this influenced their disclosure behavior.
- Considerations for choosing a specific level of disclosure during and after the study.
- Some questions to address their understanding of the system (e.g., accessibility of data to other parties or expectations about changes in the system after each disclosure choice).
- Experience with the system after each disclosure choice. This was addressed to see whether participants noticed any benefits from disclosing information and whether this may have influenced their disclosure choices.
- Feelings about the disclosure of music preferences or personality traits profile information.

Results

The disclosure behavior of the users as it was logged and an analysis of the self-reported data are presented. First, the procedure used for the analysis of the qualitative data is explained. If necessary, the raw data were translated into English (all interviews were conducted in Dutch, recorded on tape and transcribed verbatim; the questionnaire questions were posed in English, but a few participants chose to answer in Dutch). Both the questionnaire and interview included open questions, where participants were free to express their feelings in their own words. The questionnaire and interview data were analyzed by means of open coding (Strauss & Corbin, 1990). Some numerical summaries of the coded data are given if such an overview is considered to help promote clear presentation of the data, but it should be taken into account that the interview and questionnaire data are of a qualitative nature and should be interpreted as such.

Recommendations Appreciated More When Based on Music Preferences Than When Based on Personality Traits

During the use of the recommender, participants were asked to rate the quality of each playlist on a 5-point scale ranging from 1 (*very bad*) to 5 (*very good*). The average quality rating of all playlists was 3.02. Playlists that were based on music preferences were rated higher on average (3.41) than playlists that were recommended to participants on the basis of personality traits (2.63, $p < .001$). This suggests that based on the inventories used for user profiling (TIPI and STOMP), personality traits are less suitable for providing music recommendations than preferences for music genres.

A similar image arose from the questionnaire data where participants were asked to list “things they liked or disliked.” Four participants said they liked the music that was recommended to them on the basis of their preferences for music genres, and only one said they disliked these recommendations. In contrary, with regard to the system based on personality traits, four participants were not satisfied with the recommendations, and only one participant was happy with the outcome.

Participants Felt They Participated in a Study of Music Recommendation and not of Privacy

As “things they liked or disliked,” most participants mentioned the music that was recommended to them. Hardly any privacy-related topics were mentioned. Only one participant mentioned, among other things, that “there is very little information about how your decisions will affect your privacy.” Another participant said he or she disliked “the comparison of my data with that of other persons” without further explanation or reference to privacy. Five participants said they disliked the lack of control over the user profile after the initial information was provided. Even though this does not necessarily imply a lack of privacy, it should be noted that user control is frequently mentioned in relation to privacy (see the definition in the introduction and, for example, Bellotti & Sellen, 1993; Günther & Spiekermann, 2005; Margulis, 2003; W3C, 1998). The fact that the vast majority of participants

did not address the topic of privacy explicitly suggests that they experienced the system as a music recommender service and not as a study in privacy.

In the interviews with 21 of the participants, their initial reaction to the music recommender was discussed to see if they would mention any privacy issues. Overall, 10 interviewees were positive about the music recommender, 5 were negative, and 6 were more or less neutral. None of the interviewees raised privacy issues at this point. Many participants commented on the concept of the music recommender or the type of music that was recommended to them (e.g., “I thought it was very surprising. Nice and surprising. (. . .) I got to hear music that I didn’t already know but liked nevertheless. It was nice for once to hear some new music instead of the music I have in my own playlist.” This implies two things: Participants perceived the study as being about music recommendations, and the music recommender provided actual benefits to them.

Participants were asked via a questionnaire about their initial reaction to having to disclose profile information to the system. Most participants said they had no problem with providing either music preferences ($n = 21$) or personality traits ($n = 19$). Many participants indicated that they expected to be asked about their music preferences ($n = 18$), but the inquiries about personality traits surprised quite a few participants ($n = 6$). No privacy-related comments were made with regard to the collection of music preferences or personality traits. So although personality traits were not readily associated with music recommender services, it seems that participants were quite open about disclosing them.

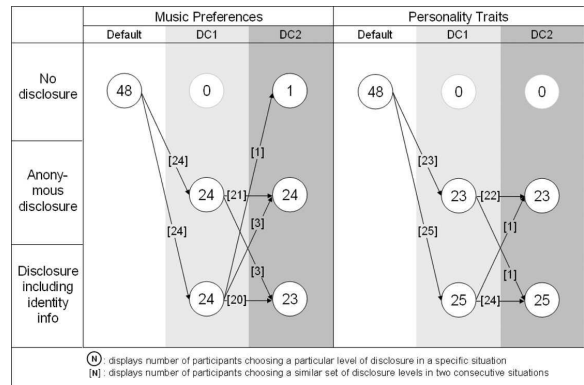
Interviewees Unaware of the Study’s Focus on Privacy

The goal of the research was discussed at the beginning of most interviews. There were some interviewees who expected only the development or improvement of the music recommender to be the goal of the research. Although there were also some interviewees who expected privacy to be of interest to the research, it was never mentioned as a single expected goal. On the whole, interviewees turned out to be unaware of the actual aim of the research. It seems that interviewees who expected that privacy was the focus of the study tended to disclose less information than those who did not mention it. Interviewees also tended to disclose less information if they did not expect the focus of the research to be on music preferences (compared to those who did or those who did not mention it). It seems that participants feel more comfortable disclosing information if they are under the impression that they are using a music service and if privacy is not mentioned. It seems that thoughts about privacy make people more aware of the risks involved.

Disclosure Behavior Consistent Across Situations Yet Divided Between Participants

An overview of the disclosure behavior of all participants is shown in Figure 3. The number of participants who chose a particular level of disclosure per disclosure situation is shown in a circle. The number between brackets on the arrow between two circles refers to the number of participants who chose a similar set of disclosure levels in two consecutive situations. For example, [3] refers to the 3 participants who chose anonymous disclosure in

Figure 3
Overview of Disclosure Choices (DC) by All Participants During the Study



Note: The three rows represent the different disclosure options participants were given at each of the four disclosure choices. The columns represent the three different phases (of which one default situation where no disclosure choice could be made) for the recommender systems based on music preference and personality traits. The order in which participants used the two recommender systems was counterbalanced.

the situation relating to comparison of music preferences by the system and chose disclosure including identity information in the situation concerning showing music preferences directly to other users. Figure 3 shows that participants were divided in their disclosure behavior and maintained the same level of disclosure (44% chose anonymous disclosure and 42% chose disclosure including identity information) throughout the four choice situations. Only 15% of the participants varied their chosen level of disclosure at different phases of the study. These variations did not show a clear trend toward increasing or decreasing levels of disclosure. Similar findings were reported earlier for the pilot of this study (Perik et al., 2004).

Anonymous Disclosure Motivated by Privacy Concerns, Full Disclosure Motivated by Expected Benefits

In the questionnaire and interviews, participants were asked to explain why they chose a specific level of disclosure.

Motivation for anonymous disclosure. Many participants who chose anonymous disclosure said their choice was based on privacy concerns ($n = 19$). One of them said, “People may see the information I provided; I do not know who reads it and therefore prefer to have the information anonymous.” Some chose anonymous disclosure just to be on the safe side ($n = 4$) (e.g., “I felt that others did not need to know my name . . .”).

Motivation for disclosure including identity information. However, most of the participants who disclosed identity information based their decision on the benefit they might gain in return ($n = 12$). As expressed by one participant, “I want the system to perform best.

That's why I gave full permission." Another large group ($n = 10$) said they chose disclosure including identity information simply because they had "no problem" disclosing the information. A number of participants said that they wanted to support the research by disclosing their information ($n = 7$), and some of them added that they didn't care about the information disclosure either. The online questionnaire and the interview data together indicate that participants felt quite safe disclosing personal information within the context of this experiment, even though they were actually allowing the system to show their personal information to other users. It can therefore be argued that participation in the research did have some (but not much) influence on the behavior of participants.

Motivation mentioned by interviewees. All interviewees said they were influenced to some extent in their disclosure decisions by the costs and benefits involved. The costs and benefits expected as a result of the disclosure of information were often mentioned simultaneously (e.g., "I first wanted to see what the quality would be like, without giving full permission immediately"). Another participant stated, "You have to weigh up (. . .) the benefits and the costs. And yes, of course, it is difficult to estimate the cost of the information you provide. (. . .) If it is clear that it is going to be of benefit, then I will do it."

In summary, disclosure choices were influenced by people's perception of privacy risks and the expected benefits. This was in line with the initial expectation that participants would balance their privacy against expected benefits from personalization and would adjust their disclosure behavior accordingly.

Interviewees Worry About Unclear Purpose of Disclosure and Accessibility of Information to Others

Many interviewees said they were influenced in their disclosure choices by other parties possibly having access to their data. For example, one participant explained that "the most important reason not to choose full permission but anonymous was because I didn't know the other people who could see the information." Other factors mentioned were the fact that they were participating in a research study or that they wanted to try out the system. Also, the specific features of the system (especially the lack of information about the purpose or consequences of information disclosure and the accessibility of information to other people) did influence participants' disclosure behavior. For example, one participant stressed the importance of knowing the purpose:

I would like to know what the purpose is of releasing information. (. . .) [Providing information] anonymously is not such a problem for me. If my information is published together with my name, then in the case of the music recommender, I have to question what purpose that serves. I did not see the benefit of that.

Another participant mentioned various desirable system features: "It should be really clear why you need to provide certain information. And you should know in what domain the information is used and who gets to see the information."

In addition to the information provided in questionnaires, interviewees stressed the importance of knowing the purpose for which information should be disclosed and expressed worries about other people gaining access to their information.

Table 1
Percentage of Participants (Split According to Gender) Choosing a Specific Level of Disclosure in Each of the Four Disclosure Choice Moments

	Music Preferences				Personality Traits			
	System Comparison		Sharing With Others		System Comparison		Sharing With Others	
	NA	ID	NA	ID	NA	ID	NA	ID
Men ($n = 40$)	47.5	52.5	50	50	45	55	45	55
Women ($n = 8$)	62.5	37.5	62.5	37.5	62.5	37.5	62.5	37.5
FET	$p = .701$		$p = .703$		$p = .454$		$p = .454$	

Note: NA = no disclosure or anonymous disclosure; ID = disclosure including identity information. Bottom row shows p value for two-sided Fisher's Exact Test (FET).

Table 2
Correlation (Spearman's rho) Between Age and Level of Disclosure for Each of the Four Disclosure Choices

Spearman's rho	Music Preferences		Personality Traits	
	System Comparison	Sharing With Others	System Comparison	Sharing With Others
Correlation coefficient	.152369	.022593	.188739	.049827
Significance (two-tailed)	.301199	.878853	.198885	.736629
N	48	48	48	48

Possible Alternative Explanations for Disclosure Behavior

Besides the factors mentioned earlier, such as the type of information involved or what the information is used for, it could very well be that disclosure behavior is influenced by other factors, such as gender, age, experience of using technology, or personality traits.

For gender, no significant differences in disclosure behavior were found in each of the four disclosure situations ($p = .701$ for comparing music preferences, $p = .703$ for sharing music preferences, $p = 0.454$ for both comparing and sharing personality traits, all two sided Fisher's Exact Test; see Table 1).

Similarly, no correlation was found between age and disclosure behavior in each of the four disclosure situations (see Table 2).

Because the participants in this study were recruited from a technical university and an industrial research lab, it was assumed that most participants would be experienced users of technology. Therefore, experience in using technology was not measured, and it cannot be used in an attempt to explain individual results.

With regard to the personality traits, only openness to new experiences could explain disclosure behavior. Participants who chose disclosure including identity information in all

four choice situations scored significantly higher on the “openness to new experiences” trait than participants who chose only anonymous disclosure ($U = 105.0, p = .005$).

Motives for Disclosure Not Consistent With Behavior

Comparison of the chosen levels of disclosure and the reported explanations of participants in questionnaires did not lead to a clear explanation of why some participants varied their level of disclosure whereas others did not. Sometimes, participants provided a consistent motivation for different disclosure choices and varying reasons for identical levels of disclosure. For example, a participant who chose disclosure including identity in the situations concerning showing information directly to other users and anonymous disclosure in the other situations provided the same explanation for these varying levels of disclosure, namely “I thought it would improve the recommender.” Another participant who chose anonymous disclosure in all four situations gave different explanations, namely “because I thought it would give me better playlists” in the situations relating to showing information directly to other users and “I want to stay anonymous in things I do over the Internet” in the other situations. The expectations with regard to the changes in the system after making a specific disclosure choice were discussed with nine interviewees. All but one said they had expected there would be some improvement in the recommendations afterward.

Participants Less Open to Disclosure Postexperiment and Expected More Benefits

In the questionnaire, participants were asked what level of disclosure they would choose for the same four situations they were asked about during the experiment. A different picture arose for the level of disclosure participants would choose after the study compared with the level they chose during the study, although most participants said they would choose exactly the same level of disclosure in all four situations as they did while using the recommender. Quite a number of participants (27%) said they would choose lower levels of disclosure in all or some of the situations (e.g., those situations relating to personality traits or to showing profile information directly to other users). More or less, the same reasons were mentioned for the disclosure choices during and after use of the recommender. However, after using the recommender, fewer participants mentioned privacy concerns or their wish to support research as a reason for their chosen level of disclosure. Yet more participants mentioned that there was no privacy risk or private information involved, and some participants mentioned that they expected better recommendations and more substantial improvements as a benefit from the disclosure of personal information. Similarly, most of the participants who chose different levels of disclosure during and after use of the recommender also explained their changes by pointing at privacy-related issues or the lack of benefits. This finding was confirmed in interviews with eight of these participants.

Sample Quite Representative in Terms of Privacy Attitudes

The PSI was used as a measure for privacy attitudes. According to the PSI, 38% of participants were privacy fundamentalists (very high privacy concern), 50% were privacy

pragmatists (balanced attitudes), and 13% were privacy unconcerned (very low or no concern). This matches well with the Harris Interactive (2002) sample, where the segments were 25%, 55%, and 20%, respectively.

The PAQ was used as a second measure for the privacy attitudes of 43 participants. These participants had a fairly neutral attitude toward the disclosure of personal information and were generally willing to be monitored. Analysis of the results shows that the participants were, on average, willing to expose their images to the public and they had an interest in protecting against unwanted intrusions.

Participants' opinions about some general privacy issues were assessed. It turned out that they like to receive advance notice or a clear description of the purpose of the information collected. Furthermore, participants said they valued being able to check and correct the personal information held by a system. Nonetheless, participants do little to protect themselves; they rarely read privacy policies, and they do not use encryption of e-mail. Participants did tend to be more willing to provide personal information in return for low-cost products or convenience. However, participants also said they provided fictitious data in some cases.

Personality Traits Perceived as More Sensitive Than Music Preferences

In the interviews, the sensitivity of the various types of information involved (music preferences, personality traits, and identity information) were discussed. Most interviewees indicated that identity information was regarded as most sensitive, followed by personality trait information. None of the interviewees said they regarded music preference information as sensitive.

In the questionnaire, participants were also asked how they felt about disclosing information either to other people or to a music content provider. It turns out that more participants worry about disclosing personality traits than about disclosing music preferences (44% vs. 10%, respectively, with regard to disclosure to other people and 50% vs. 19%, respectively, with regard to disclosure to a music content provider). Both of these ratios are significant ($p = .012$, and $p = .002$, two-sided Fisher's exact test).

PAQ and PSI Poor Predictors of Behavior; Context-Specific Privacy Concerns Are Better Predictors

The relationship between the actual disclosure behavior of participants and their privacy attitudes based on the PAQ or PSI were investigated. The PAQ scores did not provide a good indication of the disclosure behavior in this context, except for the "personal information" factor. The higher score on this factor relates to a higher level of disclosure. For the "exposure," "monitoring," and "protection" factors, there is no clear relationship with disclosure behavior. As for the PSI, one would expect that people who are "privacy unconcerned" would choose high levels of disclosure, and "privacy fundamentalists" would choose low levels of disclosure. However, this is not the case in the current study (see Table 3). There is very little difference in the level of disclosure chosen between the three PSI segments. Some participants who can be characterized as "privacy fundamentalists" chose disclosure including identity and said they had "no problem" with the information disclosure.

Table 3
Percentage of Participants (Split According to PSI Segmentation
and Shown for All Participants) Choosing a Specific Level of Disclosure
in Each of the Four Disclosure Choice Moments

		Unconcerned (<i>n</i> = 6)			Pragmatists (<i>n</i> = 24)			Fundamentalists (<i>n</i> = 18)			Total (<i>N</i> = 48)		
		N	A	ID	N	A	ID	N	A	ID	N	A	ID
MP	System comparison	0	67	33	0	46	54	0	50	50	0	50	50
	Sharing with others	0	67	33	0	46	54	6	50	44	2	50	48
PT	System comparison	0	67	33	0	42	58	0	50	50	0	48	52
	Sharing with others	0	67	33	0	42	58	0	50	50	0	48	52

Note: N = no disclosure; A = anonymous disclosure; ID = disclosure including identity information; MP = music preferences; PT = personality traits; PSI = Privacy Segmentation Index.

Furthermore, some of the “privacy unconcerned” participants argued that they valued their anonymity when they were asked to explain their disclosure behavior.

The questionnaire items relating to worries about disclosing personal information to other people or to a music content provider do give a better indication of actual disclosure behavior. The participants who said they were not worried about disclosing music preferences or personality traits tended toward a higher level of disclosure than participants who said they were worried about these types of disclosure.

Discussion

When setting up the study presented here, we had several concerns and expectations. The concerns were mostly about creating realistic privacy dilemmas that participants would experience in a realistic context of use. The development of a purpose-built application enabled the study to provide a personalized service and to deploy the experimental protocol for collecting empirical data. Although this approach is very laborious, it has a lot of potential for privacy research. The results of this study complement similar findings from surveys or experiments conducted in the artificial setting of the laboratory.

Regarding the initial concern about the realism of the privacy dilemmas, several precautions were taken and verified post hoc. The incorporation and verification of the following precautions was fundamental in ensuring the validity of empirical results relating to privacy:

- Ensuring the participants were not aware of the study’s focus on privacy.
- Avoiding sampling bias.
- Ensuring that benefits and costs from disclosure were actually experienced as such.
- Providing rewards for participants to encourage honest disclosure (as was done in this study).
- A purpose-built application should have a look and feel analogous to current services and should not appear minimal or scientific (software made for experimentation normally looks different from a commercial service).

Throughout the experiment, a number of measures were taken to avoid priming participants as to the study's interest in privacy because this might have induced normative reactions. On the whole, this was successful; for example, none of the participants spontaneously mentioned privacy issues when discussing the system. This indicates that the setup succeeded in not sensitizing them to privacy and that their behavior during the study would be a good representation of actual behavior in such a situation.

One concern in terms of how representative privacy experiments are is the natural tendency of privacy-concerned individuals to decline to participate in research. Clearly, when consent is obtained or when the nature of the study is described, individuals who are more concerned about privacy may refrain from participating. Possible strategies to encourage these privacy-concerned individuals to participate include the following:

- Recruit participants for an experiment with a non-privacy-related topic (e.g., music recommendations, as in the current study).
- Recruit participants for a survey study, as this may be considered more anonymous and less threatening.
- Use field observation and ask people for permission to use their data for research purposes afterward. However, this leads to obvious ethical concerns, and privacy-concerned individuals may still have a tendency to decline.

To prevent and check for a potential sampling bias in this study, potential participants were not informed about the focus on privacy, and finally, participants were questioned about their general and context-specific privacy attitudes. According to the PSI and PAQ, the participants varied in their level of concern for privacy; some were unconcerned about their privacy, whereas others had a high level of concern about their privacy.

There was a spread between participants in the extent to which they perceived privacy risks while using the music recommender. Participants reported differences in perceived privacy risks in questionnaires and interviews and disclosure behavior was also divided. Guaranteeing a spread in privacy attitudes among study participants and assessing the actual level of perceived risk that participants experience during a study should be a standard procedure for privacy research. The assessment of perceived risk in particular is often omitted in privacy research, thus constituting a serious threat to the validity of the results.

Despite the fact that the participants did not notice that the purpose of the study related to privacy, their nuanced behaviors and comments relating to the reasons why they chose a specific level of disclosure show that they were conscious of and influenced by costs and benefits relating to privacy and personalization. This is consistent with published results relating to disclosure behavior (Chellappa & Sin, 2005; Teltzrow & Kobsa, 2004), showing that the privacy issues they faced were very realistic and representative for this application area.

Initially, it was anticipated that participants would consider personality to be more personal than music preferences and would be less inclined to disclose the information concerned. At the very least, this reluctance toward disclosure was expected from participants with a high level of concern about privacy. This expectation was indeed consistent with the opinions expressed by participants, but this difference in sensitivity did not translate into differences in their disclosure behavior. One could draw two different conclusions from this finding: Storage of a model of users' personalities is less sensitive with respect to privacy than was initially expected—removing one of the most serious barriers for its acceptance

as a basis of personalization (this was reported before about the pilot of this study, see Perik et al., 2004). However, users may need to be protected from disclosing personal information too easily in contexts where it does not fulfill the legitimacy criterion (Iachello & Abowd, 2005). For personality traits to be viable as a user-modeling approach, future research should provide a thorough understanding of privacy risks relating to misuse or leaks of personality profiles.

Another expectation during the setting up of the experiment was that participants would balance privacy costs against expected benefits from personalization and would vary their behavior through the experiment accordingly. For example, appreciation of music should encourage them to become more open to disclosure, or at least some individuals should modify their disclosure choices according to the recipient of the information they have disclosed. It is surprising that such a tradeoff did not take place. Participants selected a specific level of disclosure throughout the experiment and kept it constant throughout the experimental conditions. Furthermore, the quality of the recommendations based on personality traits was perceived to be lower than those based on music preferences. So the personality-based recommender system seemed to involve a higher risk and to provide lower benefits, yet participants still chose similar levels of disclosure for music preferences and personality traits. This could be because the difference in sensitivity between music preferences and personality traits is small or because the benefits they experienced did not justify changing disclosure. The questionnaire and interview data reported do not support this latter explanation, however. An alternative explanation, given the novelty of profiling the personality of users, could be that curiosity about the effect of personality traits on their recommendations drove participants to experiment and explore this feature despite their privacy concerns.

Prior to the study, it was expected that showing information directly to other users might be considered to involve more risk than the mere comparison of data with that of other users. However, no difference in disclosure was found between these situations, and participants' comments did not support this expectation either. Participants said they were somewhat hesitant or cautious in their disclosure choices because they did not know exactly what would happen with the information involved or who would see the information.

When studying privacy, it is important that the privacy dilemmas are actually experienced as such (as was the case in the study described here). Also, the system should provide benefits to participants that measure up to current offerings. In practice, this could mean it is necessary to carry out a pilot study to confirm the quality of the system itself before using it to study privacy. Depending on the hypothesis tested, it is necessary to check that the privacy dilemmas introduced as a manipulation are evaluated accordingly by the participants and that they do also produce the expected range of behaviors. For example, in the study presented, a check was carried out to verify that personality traits and music preferences are perceived to be sensitive. Furthermore, the quality of the recommendations was assessed. However, the tradeoff between costs and benefits in the study was not evident, probably because the variations were not large enough to motivate participants to adapt their disclosure during the experiment. To study dynamic modification of disclosure preferences, a pilot study should be carried out first to check that the variation in costs and benefits is sufficient to motivate disclosure behavior that varies across the choice situations of the experiment.

The results of this study show that the question of whether disclosure is anonymous is more important than the type of information disclosed or the situation involved. Regardless of the type of information or the way it is going to be used, some participants were particularly anxious to safeguard their anonymity. The study by Berendt et al. (2005) also identified a group of participants who were primarily concerned about their identity. Other studies indicate the influence of identification on information disclosure; however, they do not distinguish groups of users on the basis of this influence (see, e.g., Ackerman et al., 1999). Future research could explore the potential of segmenting users on the basis of their need for anonymity versus their general privacy preferences.

A common element in the results described is a discrepancy between the privacy attitudes stated in questionnaires and interviews and people's actual behavior. The most obvious difference is that personality traits were considered by more people to be more sensitive than music preferences, yet the extent to which these two types of information were disclosed was practically identical. Even some privacy-concerned individuals chose to disclose their profile and identity information despite their self-reported concerns about such disclosure.

As mentioned already, a similar discrepancy between privacy attitudes and behavior is found in the study by Spiekermann, Grossklags, and Berendt (2001) or Berendt et al. (2005) in the context of disclosing personal data to a shopping bot. Their findings could be challenged on four accounts. First, experimental tasks were conducted in the context of a laboratory, which may influence participants' perception of privacy. Second, attitudes were measured prior to the behavior and in the absence of a specific task context. Third, participants were explicitly shown privacy statements of the companies involved before starting their shopping experience, which may have raised their awareness of privacy issues. The present study is consistent in its findings (albeit in a different application domain) while addressing these threats. The music recommender appeared to be similar in every way and was used in similar situations to any Internet-based recommendation service. Furthermore, participants were on the whole not aware of the focus on privacy and experienced the privacy risks as real.

The study by Ackerman et al. (1999) has also found that some participants were quite willing to disclose personal data regardless of whether they reported a high level of concern about privacy. However, their study involved a survey in which participants did not experience the actual consequences of their stated disclosure behavior. The present study provides stronger evidence of this discrepancy as it relates to surveyed attitudes regarding a specific context after the relevant disclosure choices had been made.

Regarding instruments that exist for measuring privacy-related attitudes, the PAQ and PSI inventories did not give sufficient insight into actual disclosure behavior. It seems that the development of standardized and validated instruments for assessing general privacy attitudes would be a useful methodological advance. In contrary, the simple questions concerning the worries people had about disclosing music preferences and personality traits did form a good indication of actual disclosure behavior. This emphasizes that it is important to assess attitudes in a way that relates closely to the context of interest, as it is known that attitudes expressed outside a specific context are very poor predictors of actual behavior (Ajzen & Fishbein, 2005).

The interviews conducted turned out to be invaluable for understanding participants' motivation for disclosure. The qualitative data obtained allowed the experimenter to clarify ambiguous comments made in the questionnaire. Because participants were not informed

about the study's focus on privacy before answering the questionnaire, some of their answers were not straightforward in their implications for privacy.

The chosen setup provided great control over the experimental condition, although it took a lot of time and effort to build the application for the music recommender service. One could also question the external validity of such an experimental setup: as discussed earlier, the mere thought of participating in research may change some people's concerns about privacy. Potentially, these issues could be overcome by using an existing service and logging actual use. However, deception issues arise because data cannot be logged without notifying the users a priori. If permission is obtained to collect user data for research purposes, then this may influence user behavior in very similar ways to the purpose-built setup. Furthermore, analysis of the use of an existing application gives less control over the context in which disclosure choices are made. This may cause difficulties in eliminating conflicting variables, in ensuring a balanced sample in surveying opinions at appropriate points in time, and in obtaining the right logs. This difficulty of applying proper research methodologies to study privacy attitudes and behavior has been addressed in recent workshops (Patil, Romero, & Karat, 2006; Romero, Perik, & Patil, 2005).

Conclusion

The aim of this study has been to examine how specific types and uses of personal information would influence people's privacy decisions and attitudes. For this purpose, participants used two different music recommender systems: one based on music preferences and one based on personality traits. Participants were asked whether the system was allowed to use their personal information for either profile matching within the system or for direct disclosure to other users. Participants experienced four choice situations in which they could choose the desired level of disclosure. A combination of logging disclosure behavior, questionnaires, and interviews made it possible to study disclosure behavior, the sensitivity of personality as an element of user modeling, and the relationship between attitudes and behavior in this domain.

The type of information (music preferences, personality traits) and the intended use of the information (collaborative filtering and access by other users) did not affect disclosure behavior. On the other hand, it appears that identity information in particular is very important to some participants and less so to others.

The study suggests that modeling personality traits of users does not present an acceptance barrier relating to privacy concerns (see also Perik et al., 2004). However, the potential misuse of this information is not yet understood sufficiently. This lack of understanding of the potential risks of modeling the personality traits makes users unable to guard their privacy, which raises practical and ethical problems relating to the development of related services. Considering the ease with which users disclose information that they consider personal, safeguards may be needed to prevent disclosure in contexts in which this is not safe. Further research needs to be conducted to see if the findings relating to personality can be extended to other applications as well.

This research contributes to existing literature on personalized systems and privacy in several ways. The current study has provided strong evidence regarding the discrepancy

between stated attitudes and user behavior relating to privacy. Personality traits were claimed to be sensitive information, yet they were disclosed to much the same extent as the less sensitive music preferences, despite participants' claims that they balance costs against the benefits of disclosure. Even privacy-concerned individuals have disclosed their profile information including their identity information.

The current study also illustrates the difficulties of doing ethically responsible and ecologically valid experimental work in the domain of HCI and privacy. Furthermore, limitations of existing inventories for surveying privacy preferences have been proposed. Sensitivity toward disclosing one's identity seems to be more important to people than the other information they exchange. Consequently, the attitude toward disclosure of one's identity provides a better categorization of users.

Our current research aims to investigate whether adherence to legal privacy guidelines leads to a higher acceptance of systems by end users and to identify which guidelines are most valued by them. Furthermore, the possibility of classifying users with regard to their privacy preferences is explored.

Important implications for future empirical studies concerning privacy are triangulation of different data collection methods, representative sampling by profiling the pool of participants with an established privacy attitudes inventory, exposure to realistic privacy risks (unlike in privacy surveys where participants are not confronted with the consequences of their self-proclaimed behavior), making sure these risks are not mitigated by trust in the experimenter, and finally disguising the experimenter's interest in privacy (as this may influence participants' behavior). Using a purpose-specific application that provides partial control for the context of disclosure was an interesting but laborious approach; Compared to the alternative of logging existing services, it provided more control over the context of disclosure and allowed the sampling of user opinions to be timed very precisely with regard to the use of the system.

References

- Aarts, E., Harwig, H., & Schuurmans, M. (2001). Ambient Intelligence. In J. Denning (Ed.), *The invisible future* (pp. 235-250). New York: McGraw-Hill.
- Ackerman, M. S., Cranor, L. F., & Reagle, J. (1999, November). Privacy in e-commerce: Examining user scenarios and privacy preferences. In *Proceedings of the 1st ACM conference on electronic commerce EC '99* (pp. 1-8). New York: ACM Press.
- Acquisti, A. (2004). Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM conference on electronic commerce EC '04* (pp. 21-29). New York: ACM Press.
- Adams, A., & Sasse, M. A. (2001). Privacy in multimedia communications: Protecting users, not just data. In A. Blandford, J. Vanderdonck, & P. Gray (Eds.), *People and computers XV: Interaction without frontiers. Joint proceedings of HCI2001 and IHM2001* (pp. 49-64). London: Springer-Verlag.
- Ajzen, I., & Fishbein, M. (2005). The influence of attitudes on behavior. In D. Albarracín, B. T. Johnson, & M. P. Zanna (Eds.), *The handbook of attitudes* (pp. 173-221). London: Lawrence Erlbaum.
- Altman, I. (1975). *The environment and social behavior: Privacy, personal space, territory, crowding*. Monterey, CA: Brooks/Cole.
- Bellotti, V., & Sellen, A. (1993). Design for privacy in ubiquitous computing environments. In G. De Michelis & C. Simone (Eds.), *Proceedings of the third European conference on computer supported cooperative work ECSCW'93* (pp. 77-86). Dordrecht, the Netherlands: Kluwer Academic Publishers.

- Berendt, B., Günther, O., & Spiekermann, S. (2005). Privacy in e-commerce: Stated preferences vs. actual behavior. *Communications of the ACM*, 48(4), 101-106.
- Chellappa, R. K., & Sin, R. G. (2005). Personalization versus privacy: An empirical examination of the online consumer's dilemma. *Information Technology and Management*, 6(2-3), 181-202.
- Chignell, M. H., Quan-Haase, A., & Gwizdka, J. (2003). The Privacy Attitude Questionnaire (PAQ): Initial development and validation. In *Proceedings of the Human Factors and Ergonomics Society 47th annual meeting* (pp. 1326-1330). Santa Monica, CA: Human Factors and Ergonomics Society.
- Child, I. L. (1968). Personality in culture. In E. F. Borgatta & W. W. Lambert (Eds.), *Handbook of personality theory and research* (pp. 82-145). Chicago: Rand McNally.
- Clarke, R. (1999). Internet privacy concerns confirm the case for intervention. *Communications of the ACM*, 42(2), 60-67.
- Consolvo, S., Smith, I. E., Matthews, T., LaMarca, A., Tabert, J., & 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 CHI '05* (pp. 81-90). New York: ACM Press
- Cranor, L. F. (2004). "I didn't buy it for myself": Privacy and ecommerce personalization. In C.- M. Karat, J. O. Blom, & J. Karat (Eds.), *Designing personalized user experiences in eCommerce* (pp. 57-73). Norwell, MA: Kluwer Academic Publishers.
- Good, N. S., & Krekelberg, A. (2003). Usability and privacy: A study of KaZaA P2P file sharing. In *Proceedings of the SIGCHI conference on human factors in computing systems CHI '03* (pp. 137-144). New York: ACM Press.
- Gosling, S. D., Rentfrow, P. J., & Swann, W. B., Jr. (2003). A very brief measure of the Big Five personality domains. *Journal of Research in Personality*, 37, 504-528.
- Günther, O., & Spiekermann, S. (2005). RFID and the perception of control: The consumer's view. *Communications of the ACM*, 48(9), 73-76.
- Harris Interactive. (2002, February). *Privacy on and off the internet: What consumers want* (Study No. 15229). New York: Author. Retrieve January 10, 2007, from http://www.aicpa.org/download/webtrust/priv_rpt_21mar02.pdf
- Iachello, G., & Abowd, G. D. (2005). Privacy and proportionality: Adapting legal evaluation techniques to inform design in ubiquitous computing. In *Proceedings of the SIGCHI conference on human factors in computing systems CHI '05* (pp. 91-100). New York: ACM Press.
- Kobsa, A. (2001). Generic user modeling systems. *User Modeling and User-Adapted Interaction*, 11(1-2), 49-63.
- Kobsa, A. (2002). Personalized hypermedia and international privacy. *Communications of the ACM*, 45(5), 64-67.
- Kobsa, A., & Schreck, J. (2003). Privacy through pseudonymity in user-adaptive systems. *ACM Transactions on Internet Technology*, 3(2), 149-183.
- Langheinrich, M. (2002). A privacy awareness system for ubiquitous computing environments. In G. Borriello & L. E. Holmquist (Eds.), *Proceedings ubicomp 2002 international conference on ubiquitous computing* (pp. 237-245). Berlin/Heidelberg, Germany: Springer-Verlag.
- Langheinrich, M. (2005). *Personal privacy in ubiquitous computing tools and system support*. Unpublished doctoral dissertation, Swiss Federal Institute of Technology, Zurich, Germany.
- Lederer, S., Hong, J. I., Dey, A. K., & Landay, J. A. (2004). Personal privacy through understanding and action: Five pitfalls for designers. *Personal and Ubiquitous Computing*, 8(6), 440-454.
- Margulis, S. (2003). On the status and contribution of Westin's and Altman's theories of privacy. *Journal of Social Issues*, 59(2), 411-429.
- Milgram, S. (1974). *Obedience to authority: An experimental view*. New York: Harper and Row.
- Nguyen, D. H., & Mynatt, E. D. (2002, June). *Privacy mirrors: Understanding and shaping socio-technical ubiquitous computing systems* (Georgia Institute of Technology, Technical Report GIT-GVU-02-16). Retrieve January 10, 2007, from <http://www.erstwhile.org/writings/PrivacyMirrors.pdf>
- Olivero, N., & Lunt, P. (2004). Privacy versus willingness to disclose in e-commerce exchanges: The effect of risk awareness on the relative role of trust and control. *Journal of Economic Psychology*, 25(2), 243-262.
- Patil, S., Romero, N. A., & Karat, J. (2006). Privacy and HCI: Methodologies for studying privacy issues. In *CHI '06 extended abstracts on human factors in computing systems* (pp. 1719-1722). New York: ACM Press.

- Perik, E. M., de Ruyter, B. E. R., Markopoulos, P., & Eggen, J. H. (2004). The sensitivities of user profile information in music recommender systems. In *Second annual conference on privacy, security and trust PST 2004* (pp. 137-141). Fredericton, New Brunswick, Canada: University of New Brunswick.
- Rentfrow, P. J., & Gosling, S. D. (2003). The do re mi's of everyday life: The structure and personality correlates of music preferences. *Journal of Personality and Social Psychology, 84*, 1236-1256.
- Romero, N. A., Perik, E. M., & Patil, S. (2005). Appropriate methodology for empirical studies of privacy. In M. F. Costabile & F. Paternò (Eds.), *Proceedings human-computer interaction—INTERACT 2005* (pp. 87-89). Berlin/Heidelberg, Germany: Springer-Verlag.
- Spiekermann, S., Grossklags, J., & 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 EC '01* (pp. 38-47). New York: ACM Press.
- Strauss, A., & Corbin, J. (1990). *Basics of qualitative research: Grounded theory procedures and techniques*. Newbury Park, CA: Sage.
- Teltzrow, M., & Kobsa, A. (2004). Impacts of user privacy preferences on personalized systems. In C.-M. Karat, J. O. Blom, & J. Karat (Eds.), *Designing personalized user experiences in eCommerce* (pp. 315-332). Norwell, MA: Kluwer Academic.
- W3C. (1998, July). *P3P guiding principles: W3C NOTE 21 July 1998*. Available at <http://www.w3.org/TR/NOTE-P3P10-principles>
- Westin, A. F. (1967). *Privacy and freedom*. New York: Atheneum.

Evelien van de Garde-Perik is a PhD candidate at Eindhoven University of Technology. Her PhD focuses on privacy concerns of users of Ambient Intelligent systems and possible ways to address those concerns by means of system design. She may be reached at e.m.v.d.garde@tue.nl.

Panos Markopoulos is an associate professor in user-centered engineering at the Industrial Design Department of Eindhoven University of Technology. He may be reached at p.markopoulos@tue.nl.

Boris de Ruyter, after his graduation, worked as a research assistant in experimental psychology at the University of Antwerp. Since 1994, he has been with Philips Research, where his research focuses on user modeling and psychometrics. In 2006, he was appointed principal scientist. He may be reached at boris.de.ruyter@philips.com.

Berry Eggen is a full professor and head of the user-centered engineering group of the Industrial Design Department of Eindhoven University of Technology. He may be reached at j.h.eggen@tue.nl.

Wijnand IJsselsteijn is an assistant professor in human-technology interaction at the Technology Management Department of Eindhoven University of Technology. He may be reached at w.a.ijsselsteijn@tue.nl.