By Bettina Berendt, Oliver Günther, *and* Sarah Spiekermann

PRIVACY IN E-COMMERCE: Stated Preferences vs. Actual Behavior

In search of an intelligent privacy watchdog at the user's service.

In times of ubiquitous electronic communication and increasing industry pressure for standard electronic authentication, the maintenance of privacy, or "the right to be left alone" becomes a subject of increasing concern. The possibility of a "transparent human," whose vital information is up for grabs, can most easily be envisioned in the realm of e-commerce, due in part to the large amounts of data available, and in part to the high payoffs expected from using this data for marketing purposes.

Illustration by Serge Bloch

Questionnaire-based surveys suggest many people strongly oppose this trend. For example, 75% of German Internet users surveyed in [5] professed some fear their privacy may be compromised when surfing the Internet; 60% had avoided a Web site in order to protect privacy; and 47% sometimes provided false data. Similar results have been obtained in other countries: 82% of online users have refused to give personal information; and 34% have lied when asked about their personal habits and preferences [12]. The use of aliases, including obvious ones such as "Donald Duck" is commonplace. The integrity and efficiency of commercial Web sites' data protection measures are widely doubted.

Given these widespread concerns about personal privacy in a networked world, it is commonly assumed that online behaviors reflect such concerns. Privacy enhancing technologies (PETs) such as P3P build on the idea that Internet users study privacy statements, and restrict what they reveal to whom—in short, that they *act in accordance with their privacy preferences.* However, as our study will show, this is often not the case.

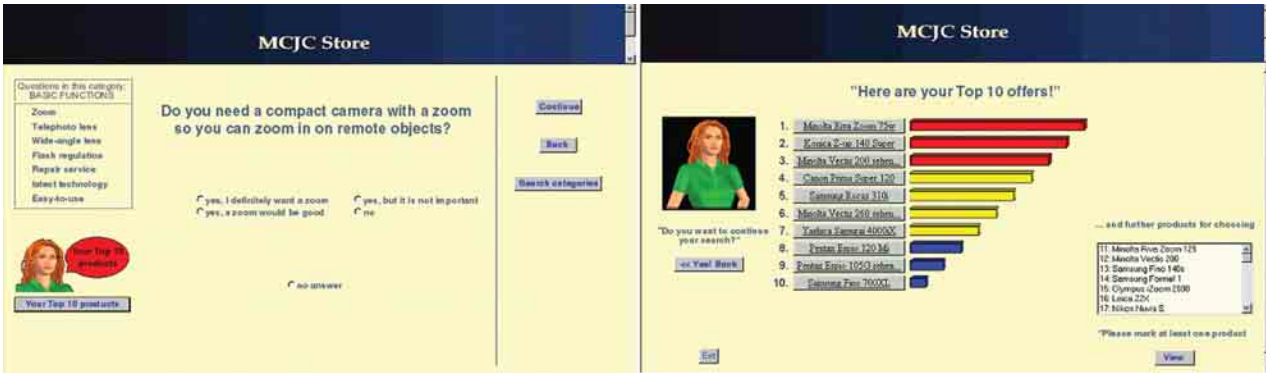

Figure 1. Agent questions and recommendations.

In this article, we describe results from a large-scale online shopping experiment. Findings suggest that, given the right circumstances, online users easily forget about their privacy concerns and communicate even the most personal details without any compelling reason to do so. This holds true in particular when the online exchange is entertaining and appropriate benefits are offered in return for information revelation—circumstances easily created by second-generation agent technologies and embodied interface agents. Privacy statements have no impact on most users' behavior. We also discuss possible reasons for this discrepancy between stated preferences and actual behavior, and suggest how to help users better align their actions with their goals.

An Experimental Investigation of Privacy Attitudes and Behavior

The experimental setting was an online store with agent recommendations. The initial goal of our study (for details, see [11]) was to investigate drivers and impediments of online interaction in general. Privacy concerns were suspected to be *one* major impediment of truthful and deep online interaction. In particular, our study focused on how self-reported privacy concerns relate to actual self-disclosing behavior, and on the ultimate impact of privacy statements.

In a laboratory experiment, 206 participants took a virtual shopping trip for cameras and jackets. As an incentive to participate, these high-value goods were offered at a 60% discount compared to local store prices. The buying decision was assisted by an anthropomorphic shopping bot. Participants had to spend their own money if they decided to buy.

Before shopping, participants filled out a questionnaire. More than a quarter of the questions were privacy-related, addressing respondents' willingness to reveal certain types of private data, their general trust in privacy statements, the value of privacy, and their intended reactions to various privacy scenarios.

Participants were asked to sign the store's privacy statement, agreeing to the sale of their data to an anonymous project sponsor. One group received a "cordial" privacy statement, which told them their navigational data would be handed over to the sponsor, a reputable European company. This statement advised them of their rights under the European Union Directive on Data Protection (95/46/EC): the rights to be informed about who processes the data for which purpose; to inspect one's data; to enforce the amendment if incorrect; and to refuse to consent to specific types of usage. The other group received a

Findings suggest that, given the right circumstances, *online users easily forget about their privacy concerns and communicate even the most personal details without any compelling reason to do so.* This holds true in particular when the online exchange is entertaining and appropriate benefits are offered in return for information revelation.

"terse" privacy statement, which did not mention the EU Directive but told them it was unknown how the sponsor would use their data.

The navigation opportunities in the store were similar to those in current online shops. At the beginning, the anthropomorphic shopping bot named Luci introduced herself and her purpose to the user. Before entering the store, users were given the possibility to leave their home address, but no reason or requirement to do so.

Users were then invited to answer any of up to 56 agent questions (see Figure 1). At any time, the agent could be asked to determine a personalized top 10 list of products, based on the answers given so far. The user could request information on each product and choose to purchase it. Unlike current Web shopping agents, the bot not only focused on product attributes, but also asked soft questions typically found in offline sales conversations. The goal was to include more questions, and more personal questions, than one would expect customers to answer. In addition to product attribute questions regarding the desired features of a camera zoom lens, for example, we asked about the intended use of the product (for example: "At which occasions do you usually take photos?"). Other questions, such as "How important are trend models to you?" were designed to influence product recommendations. The selection of personal questions included several that were unrelated to the product but related to the sales context, and that would usually be considered inappropriate. For example, people were asked how "photogenic" or "conceited" they considered themselves to be. A previous, independent evaluation of bot questions found about half were considered non-legitimate or unimportant in [3].

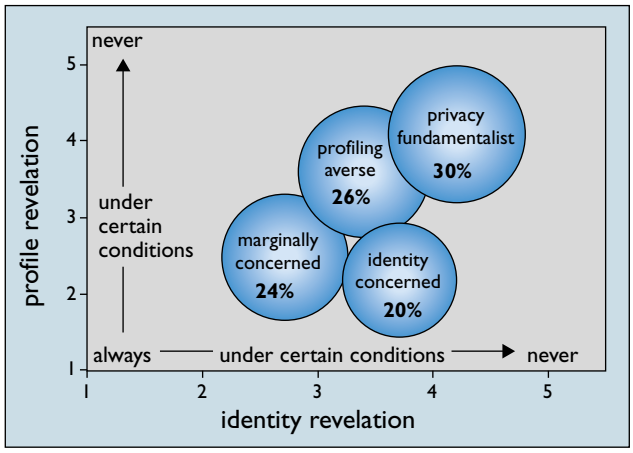

Figure 2. Four clusters of privacy attitudes.

The final analysis, based on 171 valid cases, compared users' attitudes with their actual behavior. Based on the Web log data, behavior was described in terms of the information provided. The quantity of that information was measured by the proportion of bot questions the user had answered. To measure information quality, we developed an index called *personal consumer information cost (PCIC)* that considered each answered question's legitimacy and importance in the sales context, as well as the difficulty of answering it. A PCIC of zero means the user would have no problem answering the question truthfully. A high PCIC implies that users would be highly reluctant to give this type of information. Regression analysis confirmed that PCIC is strongly negatively correlated with legitimacy and importance, and moderately positively correlated with difficulty [3]. A participant's PCIC index was computed as the sum of PCIC indexes of all questions he or she answered, grouping values into "high," "medium," and "low." A large number of answers in response to mostly irrelevant or non-legitimate questions thus leads to a high PCIC.

A clustering of the answers to privacy-related questions revealed four different groups of users (Figure 2). We clearly distinguished a group of *privacy fundamentalists* and another group of only *marginally concerned* users as found elsewhere [1]. We were able to differentiate the remaining participants by the focus of their privacy concerns: *identity concerned* users are more concerned about revealing information like their name, email, or mailing address, while *profiling averse* users are more concerned about disclosing such information as their interests, hobbies, and health status.

To investigate whether users' interaction behavior was consistent with their privacy attitudes, we examined whether participants voluntarily gave their address to Luci before entering the question-answer cycle, and how many and what types of her questions they answered.

Figure 3 shows the proportion of participants who disclosed their address (dashed line), and the distribution of PCIC index values. The solid red line shows the proportion with high PCIC; the dotted gray line shows the proportion with high or medium PCIC; and the outermost diamond shows 100%, that is, the proportion with high, medium, or low PCIC.

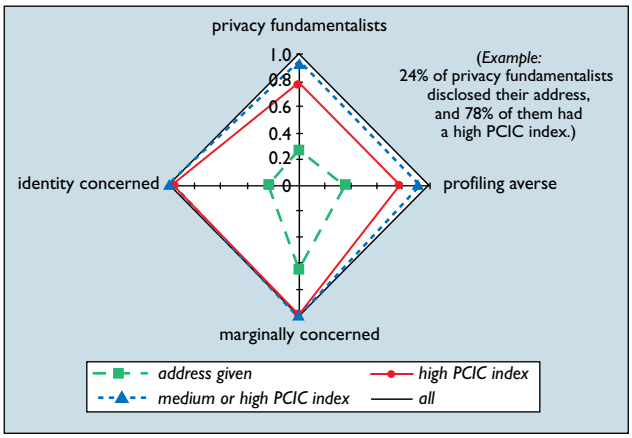

Figure 3. Attitude clusters and disclosing behavior.

As expected, disclosure rates increased from privacy fundamentalists to marginally concerned users. Identity concerned and profiling averse users showed intermediate disclosure rates and acted in *relative* accordance with their stated preferences: the former withheld their address more often, and the latter had lower PCIC index values. However, contrary to our expectations, the absolute level of disclosure was alarmingly high across all clusters, belying the previously expressed reluctance to disclose information online.

Neither the product category nor the type of privacy statement had a statistically significant impact. However, the cordial privacy statement (which referred to the EU Directive) induced slightly more participants to provide their address. This is a cause for concern: it suggests that the more people believe in the effectiveness of existing jurisdiction, the less they control their personal behavior.

In the debriefing questionnaire, most participants indicated that they appreciated the communication employed and they felt "personally addressed" and "supported" by agent Luci. This was stated even by those individuals who had previously expressed privacy concerns and were not too fond of the quality of Luci's recommendations. In the debriefing discussions, participants showed no sign of recognizing any link of the experiment to privacy research, and did not comment on a discrepancy between privacy preferences and behavior.

Why Does Behavior Diverge From Attitudes?

Our study demonstrates that Web users welcome a rich interactive environment, where they are willing to talk about themselves, thus creating the basis for efficient customer relationships. The other important news is they do not always act in line with their stated privacy preferences, giving away information about themselves without any compelling reason to do so.

This disparity may be disadvantageous not only for customers, but also for the e-commerce companies that may welcome the data initially. If customers are later confronted with the discrepancy between their actions and their ideals—for example, because the company uses information on customer preferences—they may react with resentment, which can damage the customer relationship [2].

Inconsistencies between people's behavior and their self-reports are a well-known phenomenon, with explanations emphasizing cognitive and/or social aspects of decision making and behavior [10].

In many situational contexts, decisions are based on heuristics rather than on rational consideration of all factors for or against all possible courses of action (for an overview, see [4]). In our case, the shopping context may in particular have drawn attention to the potential *gains* from disclosing personal information: product recommendations and the chance to obtain a discount. In addition, the wealth of choices in this store interface may have led to a certain decision aversion and the accompanying wish to collect all possible information.

While many users have strong opinions on privacy and do state privacy preferences, they are unable to act accordingly. ***Once they are in an online interaction, they often do not monitor and control their actions sufficiently***; privacy statements seem to have no impact on behavior.

In contrast, both the questionnaire items and the direct evaluation of bot questions in [3] may have framed information disclosure in terms of a *loss* of privacy, and the possible lack of legitimacy and importance. Moreover, the first impression of Luci may have engendered a positive mood, which makes positive memories of productive (and harmless) interactions with shop assistants more *available* and leads to the expectation that the current interaction will be like this too.

Heuristics that specifically simplify *communication* are likely to have played a role as well. Since the first anecdotal evidence of interactions with ELIZA, Joseph Weizenbaum's 1960s-era computer "psychotherapist," it has repeatedly been observed that people tend to treat interactive software like a trustworthy human communication partner. Human communication is usually characterized by adherence to the "Gricean maxims of cooperativity," which involve saying things that are true and relevant to the conversation. This generally holds for agent communication as well. (The Gricean maxims are in fact a popular guideline for agent communication design.) At least as pervasive as the *actual* adherence to these maxims, however, is the often-implicit *expectation* of adherence. Even in surveys and experiments, people assume their dialogue partners ask questions that are relevant [8]. This expectation often leads people to re-frame their perception of something said that, at first sight (or viewed in isolation [3]), seems to violate the cooperativity maxims.

These results and their interpretation must be substantiated and investigated further. Our sample was self-selected, relatively well-educated, young, and with considerable online experience: 92.7% were students; 44.2% were females; and 98.5% (91.7%) were (regular) Internet users. A more diverse population should be investigated. It was not possible to control the different roles that the financial incentives and the specific shop interface may have played in the sample. Also, it could not be ruled out that despite our efforts to the contrary, the university setting may have made participants more trustful. Nonetheless, the findings appear significant enough to warrant better measures of protection.

Protecting Privacy

An important result of this study is that while many users have strong opinions on privacy and do state privacy preferences, they are unable to act accordingly. Once they are in an online interaction, they often do not monitor and control their actions sufficiently; privacy statements seem to have no impact on behavior. Users rely on legal protection, even though it is widely known that laws and regulations have difficulty responding to the fast changes in Internet communications. Given this discrepancy, software appears to be a better basis for effective privacy protection.

Currently, P3P is still the most prominent tool for privacy protection, as it can give automatic warnings if a Web site's privacy policy does not correspond to one's personal privacy preferences. Yet, beyond these warnings, the tool does not protect a user once a Web site is entered. Privacy preferences cannot be expressed on a per-service level; they are static across the Web. Furthermore, P3P is not scalable to meet the privacy needs of more intelligent infrastructures of the future.

We therefore advocate the development of a more elaborate privacy enhancing technology (PET) building on current research in identity management systems. To support the rich and service-dependent

interaction users desire, the following requirements for privacy protection should be met:

1. PET should always be under the full control of the user.
2. PET should monitor third-party services as P3P does today and bring potential problems to the user's attention. Yet in addition, it should *learn* users' privacy preferences by observation [7], and change settings dynamically and on a per-service level.
3. PET should record Web service interactions and create information-rich *client-side profiles* [9]. At the user's discretion, parts of that profile could be made available to marketers or peer networks.

To empower users and protect them against the described context effects:

4. PET should have an easy-to-use interface and privacy-friendly default settings.
5. PET should provide identity management [6], allowing users to adopt new pseudonyms whenever they (re-)enter a site and sheltering client-side profiles.
6. PET should decontextualize. Recognition and blocking of dangerous interactions could be done automatically if PET were able to understand all interactions. However, since automatic language understanding is anything but perfect, PET must employ its user interfaces to support users' thinking about their actions while they are acting, for example, with windows that disturb the flow of interaction popping up upon unclear information requests. Based on its learning capabilities, the agent should issue warnings selectively to avoid ineffectiveness. However, learning from (ineffectual) behavior is not enough. Rather, PET could, for example, cluster interactions and periodically submit them to a critical review by the user. Alternatively, "good" interaction histories could be pooled in a peer network, and used as a basis for individual PET agents' learning.

Finally, an additional desirable (but potentially distant) feature would be the use of knowledge about Web services and their privacy practices beyond privacy statements. For this purpose independent agencies or public review boards would have to maintain standardized metadata on a company's privacy reputation. PET could systematically check this reputation index before submitting data about its user. By combining these techniques, PET would represent a more timely privacy protection and trust tool for modern Web applications. Even though people can potentially still reveal everything about themselves, this PET would go far in ensuring identity protection and serving as a learning and intelligent watchdog at the user's service. C

References

1. Ackerman, M.S., Cranor, L.F., and Reagle, J. Privacy in e-commerce: Examining user scenarios and privacy preferences. In *Proceedings of the ACM Conference on Electronic Commerce EC'99* (Denver, CO, Nov. 1999), 1–8.
2. Adams, A. The implications of users' privacy perception on communication and information privacy. In *Proceedings of the Telecommunications Policy Research Conference* (Washington, DC, 1999); www.tprc.org/ABSTRACTS99/adamspap.pdf.
3. Annacker, D., Spiekermann, S., and Strobel, M. E-privacy: A new search cost dimension in online environments. In *Proceedings of the 14th Conference on Electronic Commerce* (2001).
4. Bettman, R., Luce, M.F., and Payne, J.W. Constructive consumer choice processes. *Journal of Consumer Research 25* (1998), 187–217.
5. IFAK GmbH & Co. Nur begrenztes Vertrauen der Verbraucher beim Einkaufen im Internet. Research study. Taunusstein. Jan. 18, 2002. www.ifak.de/presse_8.php.
6. Jendricke, U. and Gerd tom Markotten, D. Usability meets security—The Identity Manager as your personal security assistant for the Internet. In *Proceedings of the 16th Annual Computer Security Applications Conference* (New Orleans, LA, Dec.2000), 344–353.
7. Lieberman, H. and Maulsby, D. Instructible agents: Software that just keeps getting better. *IBM Systems Journal 35*, (1996), 539–556.
8. Schwarz, N. *Cognition and Communication: Judgmental Biases, Research Methods, and the Logic of Conversation*. Lawrence Erlbaum, Hillsdale, NJ, 1996.
9. Shearin, S. and Liebermann, H. Intelligent profiling by example. In *Proceedings of the ACM Conference on Intelligent User Interfaces* (Santa Fe, NM, Jan. 2001), 145–151.
10. Simonson, I., Carmon, Z., Dhar, R., Drolet, A., and Nowlis, S.M. Consumer research : In search of identity. *Annual Review of Psychology 52* (2001), 249–275.
11. Spiekermann, S., Grossklags, J., and Berendt, B. E-privacy in 2nd generation E-Commerce: Privacy preferences versus actual behavior. In *Proceedings of the ACM Conference on Electronic Commerce (EC'01)* (Tampa, FL, Oct. 2001), 38–47.
12. Teltzrow, M., Kobsa, A. Impacts of user privacy preferences on personalized systems: A comparative study. In Karat, C.M., Blom, J., Karat, J., eds., *Designing Personalized User Experiences in eCommerce.* Kluwer Academic Publishers, Dordrecht, Netherlands, 2004, 315–332.

Bettina Berendt (berendt@wiwi.hu-berlin.de) is an assistant professor of information systems at Humboldt-Universität zu Berlin, Germany.

Oliver Günther (guenther@wiwi.hu-berlin.de) is a professor of information systems at Humboldt-Universität zu Berlin, Germany, and Director of InterVal, the Berlin Research Center on Internet Economics.

Sarah Spiekermann (sspiek@wiwi.hu-berlin.de) is an assistant professor at the Institute of Information Systems at Humboldt-Universität zu Berlin, Germany.