

Impacts of Contextualized Communication of Privacy Practices and Personalization Benefits on Purchase Behavior and Perceived Quality of Recommendation

Alfred Kobsa

Donald Bren School of
Information and Computer Sciences
Irvine, CA 92697-3425
+1 949 824-3007
kobsa@uci.edu

Max Teltzrow

Institute of Information Systems
Humboldt University
Spandauer Str. 1, 10178 Berlin, Germany
+49 30 2093 5730
teltzrow@wiwi.hu-berlin.de

ABSTRACT

Consumer surveys have consistently demonstrated that privacy statements on the web are ineffective in alleviating users' privacy concerns. We investigated a new user interface design approach in which the privacy practices of a website are explicated in a contextualized manner, and users' benefits from providing personal data clearly explained. To test the merits of this approach, we conducted a user experiment with two versions of a web store that allegedly provided personalized book recommendations: one with a traditional global disclosure and one that additionally provides contextualized explanations of privacy practices and personalization benefits. We found that subjects in the second condition were significantly more willing to share personal data with the website, rated the perceived benefit resulting from data disclosure significantly higher, and also made considerably more purchases. We discuss the implications of these results and point out open research questions.

Keywords

Privacy disclosure, personalization, user benefit, trust, recommendation, perceived quality, adoption, purchases

INTRODUCTION

Privacy plays a major role in the relationship between companies and Internet users. More than two third of the respondents in [3] indicated that knowing how their data will be used would be an important factor in their decision on whether or not to disclose personal data. It seems though that the communication of privacy practices on the Internet has so far not been very effective in alleviating consumer concerns: 64% of Internet users surveyed in [10]

indicated having decided in the past not to use a website, or not to purchase something from a website, because they were not sure about how their personal information would be used.

The current predominant way for websites to communicate how they handle users' data is to post comprehensive privacy statements (also known as "privacy policies" or "privacy disclosures"). 76% of users find privacy policies very important [11], and 55% stated that a privacy policy makes them more comfortable disclosing personal information [13, 19]. However, privacy statements today are usually written in a form that gives the impression that they are not really supposed to be read. And this is indeed not the case: whereas 73% of the respondents in [1] indicate having viewed web privacy statements in the past (and 26% of them claim to always read them), web site operators report that users hardly pay any attention to them¹. [2] criticizes that people are turned off by long, legalistic privacy notices whose complexity makes them wonder what the organization is hiding.

Relegating the communication of privacy policies to merely publishing comprehensive privacy disclosures also disregards the situational nature of privacy [18].² Users seem to make privacy decisions much more consistently in concrete situations than upfront. In fact, privacy preferences stated upfront and actual usage behavior often seem to differ significantly [4, 20].

Moreover, merely communicating a company's privacy policy is not sufficient. In situated interviews [5], users pointed out that "in order to trust an e-Commerce company, they must feel that the company is doing more than just protecting their data – it must also be providing

* This paper is an excerpt of [15]. Results from a second experiment as well as a proposed explanatory model have been added.

¹ For example, [16] indicates that less than 0.5% of all users read privacy policies.

² This criticism also applies to P3P [8] that is intended to alleviate current problems with privacy statements.

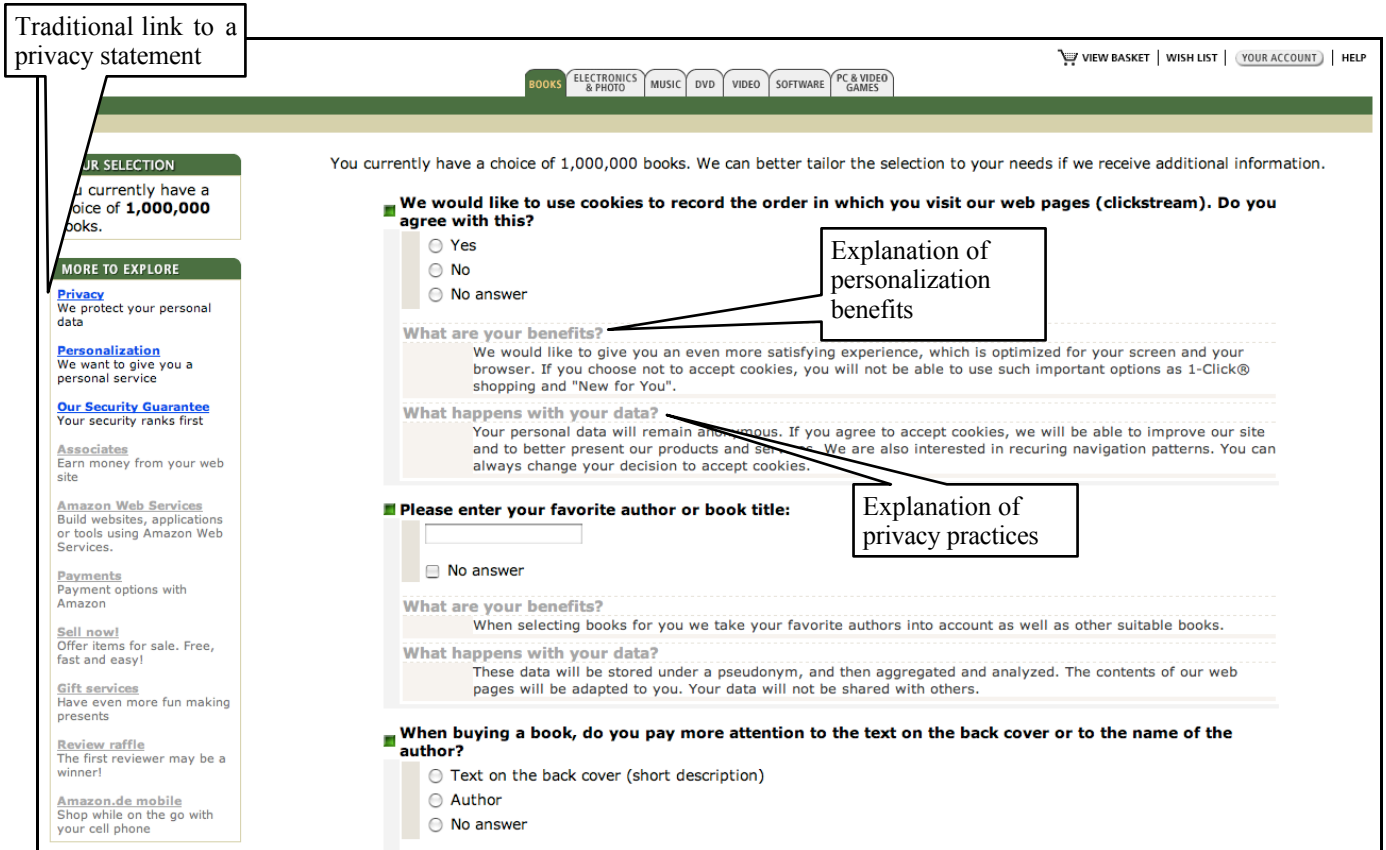


Figure 1: Global and contextual communication of privacy practices and personalization benefits

them with functionality and service that they value.” The way in which personal data is used for the provision of these services must be clearly explained. Current web privacy statements hardly address the connection between personal data and user benefits.

A DESIGN PATTERN FOR WEBSITES THAT COLLECT PERSONAL DATA

To adequately address privacy concerns of users of personalized websites, we investigate user interface design patterns that communicate the privacy practices of a site both at a global and a local (contextualized) level. Similar to design patterns in object-oriented programming, interface design patterns constitute descriptions of best practices within a given design domain based on research and application experience [22]. They give designers guidelines for the efficient and effective design of user interfaces.

Global Communication of Privacy Practices and Personalization Benefits

Global communication of privacy practices currently takes place by posting privacy statements on a company’s homepage or on all its web pages. Privacy policies are carefully crafted by legal council since they are legally binding and enforceable in many jurisdictions. Rather than completely replacing them by something new whose legal impact is currently unclear at best, our approach keeps current privacy statements in the “background” for legal reference and protection. However, we argue to enhance this kind of disclosure by additional information that explains

privacy practices and user benefits, and their relation to the requested personal data, in the given local context.

Local Communication of Privacy Practices and Personalization Benefits

We expect that tailored in-context explanation of privacy practices and personalization benefits will address users’ privacy concerns much better than global contextless disclosures. This approach breaks long privacy policies into smaller, more understandable pieces, refers concretely to the current context, and thereby allows users to make situated decisions regarding the disclosure of their personal data considering the explicated privacy practices and the explicated personalization benefits.

It seems safest to communicate privacy practices and personalization benefits at the level of each individual entry field for personal data. If a number of such fields form a visually separate sub-context on a page, compiled explanations may be given if the explanations for each individual field are not very different (due to legal differences, different sensitivity levels, privacy practices or personalization benefits). A page is the highest possible level at which compiled contextual explanations may be given (again, only if the field-level explanations are relatively similar). Visually separate sub-contexts on a page should be preferred though, due to the (cognitive) closure that they require.

An Example Website with Global and Contextual Communication of Privacy Practices and Personalization Benefits

Fig. 1 shows the application of the proposed interface design pattern to a web bookstore that gives personalized recommendations. The top three links in the left-hand frame lead to the global disclosures (to facilitate comprehension, we decided to split the usual contents of current privacy statements into three separate topics: privacy, personalization benefits, and security). The main frame contains input fields and checkboxes for entering personal data. Each of them is accompanied by an explanation of the site's privacy practices regarding the respective personal data (which focuses specifically on usage purposes), and the contribution to personalized recommendations that these data afford.

As in the theoretical model of [17], a user achieves an understanding of the privacy implications of the displayed situation both intuitively (taking the overall purpose of the site and page into account) and through adequate contextual notice. The traditional link to a privacy policy can still be accessed if so desired.

A COMPARATIVE EXPERIMENT

Materials

To evaluate the merits of our proposal, we developed a mock book recommendation and sales website whose interface was designed to suggest an experimental future version of a well-known online bookstore. Two variants of this system were created, one with contextual explanations of privacy practices and personalization benefits, and one without. Figure 1 shows an excerpt of the first variant, translated from German into English. The contextual explanations are given for each entry field (which is the safest of the strategies discussed above), under the headings "What are your benefits?" and "What happens with your data?" In the version without contextual explanations, these explanations are omitted.

In both conditions, the standard privacy policy of the web retailer is used. The three left-hand links labeled "Privacy", "Personalization" and "Our Security Guarantee" lead to the original company privacy statement (we split it into these three topics though and left out irrelevant text). In the condition with contextual explanations, the central policies that are relevant in the current situation are explained under "What happens with your data?" Such explanations state, for instance, that the respective piece of personal data will not be shared with third parties, or that some personal data will be stored under a pseudonym and then aggregated and analyzed. The explanation of the usage purpose is concise and kept in the spirit of P3P specifications [8].

A counter was visibly placed on each page that purported to represent the size of the currently available selection of books. Initially the counter is set to 1 million books. Data entries in web forms (both via checkboxes and radio buttons and through textual input) decrease the counter after each page by an amount that depends on the data entries made. The web forms ask a broad range of questions relating to users' interests. A few sensitive questions on

users' political interests, religious interests and adherence, their literary sexual preferences, and their interest in certain medical subareas (including venereal diseases) are also present. All questions "make sense" in the context of filtering books in which users may be interested. For each question, users have the option of checking a "no answer" box or simply leaving the question unanswered. The personal information that is solicited in the web forms was chosen in such a way that it may be relevant for book recommendations and/or general customer and market analysis. Questions without any clear relation to the business goals of an online bookstore are not being asked. A total of 32 questions with 66 answer options are presented. Ten questions allow multiple answers, and seven questions have several answer fields with open text entries (each of which we counted as one answer option).

After nine pages of data entry (with a decreased book selection count after each page), users are encouraged to review their entries and then to retrieve books that purportedly match their interests. Fifty predetermined and invariant books are then displayed that were selected based on their low price and their presumable attractiveness for students (book topics include popular fiction, politics, tourism, and sex and health advisories). The prices of all books are visibly marked down by 70%, resulting in out-of-pocket expenses between €2 and €12 for a book purchase. For each book, users can retrieve a page with bibliographic data, editorial reviews, and ratings and reviews by readers.

Users are free to choose whether or not to buy one single book. Those who do are asked for their shipping and payment data (a choice of bank account withdrawal and credit card charge is offered). Those who do not buy may still register with their postal and email addresses, to receive personalized recommendations in the future as well as newsletters and other information.

Subjects and Procedures

58 subjects participated in the experiment. They were students of Humboldt University in Berlin, Germany, mostly in the areas of Business Administration and Economics. The data of 6 subjects were eventually not used, due to a computer failure or familiarity with the student experimenters. Participants were promised a € 6 coupon for a nearby popular coffee shop as a compensation for their participation, and the option to purchase a book with a 70% discount. Prospective participants were asked to bring their IDs and credit or bank cards to the experiment.

When subjects showed up for the experiment, they were reminded to check whether they had these credentials with them, but no data was registered at this time. Paraphernalia that are easily associated with the web book retailer, such as book cartons and logos, were casually displayed.

In the instructions part of the experiment, subjects were told that they would test an experimental new version of the online bookstore with an intelligent book recommendation engine inside. Users were advised that the more and the better data they provided, the better would be

the book selection. They were also told that their data would be given to the book retailer after the experiment. It was explicitly pointed out though that they were not required to answer any question. Subjects were asked to work with the prototype to find books that suited their interests, and to optionally pick and purchase one of them at a 70% discount. They were instructed that payments could be made by credit card or by withdrawal from their bank accounts.

A between-subjects design was used for the subsequent experiment, with the system version as the independent variable: one variant featured non-contextual explanations of privacy practices and personalization benefits only, and the other additionally contextualized explanations. Subjects were randomly assigned to one of the two conditions (we will abbreviate them by “-expl” and “+expl” in the following). They were separated by screens, to bar any communication between them. After searching for books and possibly buying one, subjects filled in two post-questionnaires, one online and one on paper. Finally, the data of those users who had bought a book or had registered with the system were compared with the credentials that subjects had brought with.

RESULTS

Data Sharing

We analyzed the data of 26 participants in the conditions “-expl” and “+expl”. We first dichotomized their responses by counting whether a question received at least one answer or was not answered at all. Whereas on average 84% of the questions were answered in condition -expl, this rose to 91% in the second condition (see Table 1). A Chi-Square test on a contingency table with the total number of questions answered and not answered in each condition showed that the difference between conditions was statistically significant ($p < 0.001$).

The two conditions also differed with respect to the number of answers given (see Table 1). In condition “-expl”, subjects gave 56% of all possible responses on average (counting all options for multiple answers), while they gave 67% of all possible answers in condition “+expl”. A Chi-Square contingency test showed again that the difference between the two conditions is highly significant ($p < 0.001$). The relative difference between the number of answers provided in the two conditions is even higher than in the dichotomized case (19.6% vs. 8.3% increase).

	<i>-expl</i>	<i>+expl</i>	<i>diff</i>	<i>p</i>
Questions answered	84%	91%	+ 8%	<.001
Answers given	56%	67%	+20%	<.001
Book buyers	58%	77%	+33%	.07
“Data allowed store to select better books”	2.85	3.40	+19%	.035

Table 1: Effect on data sharing, purchases and perceived benefit

The results demonstrate that the contextual communication of privacy practices and personalization benefits has a significant positive effect on users’ willingness to share personal data. The effect is even stronger when users can give multiple answers. We found no significant difference between questions that we regarded as more sensitive and less sensitive questions.

Purchases

Table 1 shows that the purchase rate in condition “+expl” is 33% higher than in condition “-expl” (note that all subjects saw the same set of 50 books in both conditions). A t-test for proportions indicates that this result approaches significance ($p < 0.07$).

Perceived quality of recommendation

The paper questionnaire that was administered at the end of the study included several Likert questions on subjects’ perception of the privacy practices of the website as well as its service quality. Possible answers ranged from “strongly agree” to “strongly disagree”. One of these questions was “Did you feel that the particulars that you gave helped <bookseller> to chose interesting books for you?” Table 1 shows the average responses in the two conditions after encoding them on a one to five scale. The difference between the two conditions is highly significant (one-tailed t-test, $p < 0.05$). Note again that all subjects were offered the same set of books.

DISCUSSION OF THE RESULTS AND OPEN RESEARCH QUESTIONS

Our experiment was designed so as to ensure that subjects had as much “skin in the game” as possible, and thereby to increase its ecological relevance. The incentive of a highly discounted book and the extremely large selection set that visibly decreased with every answer given was chosen to incite users to provide ample and truthful data about their interests. The perceptible presence of the web book retailer, the claim that all data would be made available to them, and the fact that names, addresses and payment data were verified (which ensured that users could not use escape strategies such as sending books to P.O. boxes or someone they know) meant that users really had to trust the privacy policy that the website promised when deciding to disclose their identities.

The results demonstrate that the contextualized communication of privacy practices and personalization benefits has a significant positive effect on users’ data sharing behavior, and on their perception of the website’s privacy practices as well as the perceived benefit resulting from data disclosure. The additional finding that this form of explanation also leads to more purchases approached significance. The adoption by web retailers of interface design patterns that contain such explanations therefore seems clearly advisable.

Our results support several of the assumptions underlying the model in Fig. 2, which is centrally based on the notion of trust. In condition “+expl”, users’ better understanding of the website’s privacy practices and of the contribution of disclosed data to resulting personalization benefits is likely

to have increased users' trust and alleviated their privacy concerns. This in turn led to more data disclosure.

The decision to buy a book was a significant step in our experiment since at this point users revealed personally identifiable information (name, shipment and payment data) and risk that previously pseudonymous information may be linked to their identities. We already reported above that users indicate in surveys to refrain from shopping if they are uncertain about the possible fate of their data. It seems that the increased trust of users in condition "+expl" due to contextualized privacy disclosure may have contributed to more users opting to reveal their identities.

We have no direct explanation for the higher perceived benefits from data disclosure in condition "+expl". One can speculate about positive transfer effects from higher perceived privacy standards via higher trust.

Other characteristics of our experiment are also in agreement with the literature. [14] found in their study of consumer privacy concerns that "in the absence of straightforward explanations on the purposes of data collection, people were able to produce their own versions of the organization's motivation that were unlikely to be favorable. Clear and *readily available* explanations might alleviate some of the unfavorable speculation" [emphasis ours]. [9] postulate that consumers will "continue to disclose personal information as long as they perceive that they receive benefits that exceed the current or future risks of disclosure. Implied here is an expectation that organizations not only need to offer benefits that consumers find attractive, but they also need to be open and honest about their information practices so that consumers [...] can make an informed choice about whether or not to disclose." The readily available explanations of *both* privacy practices and personalization benefits in our experiment meet the requirements spelled out in the above quotations, and the predicted effects could be indeed observed.

Having said this, we would however also like to point out that additional factors may also play a role in users' data disclosure behavior, which were kept constant in our experiment due to the specific choice of the web retailer, its privacy policy, and a specific instantiation of our proposed interface design pattern. We will discuss some of these factors in the following.

Reputation of a website. We chose a webstore that enjoys a relatively high reputation in Germany (we conducted surveys that confirmed this). It is well known that reputation increases users' willingness to share personal data with a website (see e.g. [6, 12, 21]). Our high response rates of 84% without and specifically 91% with contextual explanation suggest that we may have already experienced some ceiling effects. In a more recent version of the experiment we therefore changed the name and logo of the website to ones that had received a medium reputation rating in the prior survey. We found indeed similar effects of contextualized disclosures as at the website with high reputation, but with smaller numbers for data disclosure and purchases in both conditions. There was no interaction between reputation and form of disclosure.

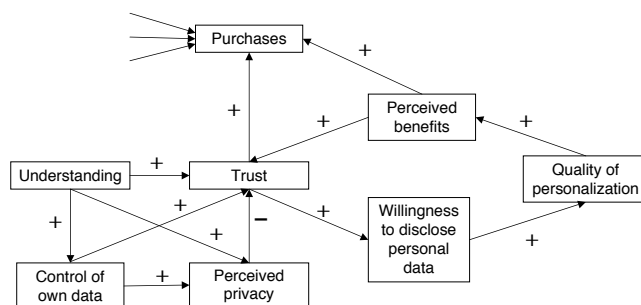


Figure 2: Suggested explanatory model

Stringency of a website's data handling practices. The privacy policy of the website that we mimicked is comparatively strict. Putting this policy upfront and explaining it in-context in a comprehensible manner is more likely to have a positive effect on customers than couching it in legalese and hiding it behind a link. Chances are that this may change if a site's privacy policy is not so customer-friendly.

Permanent visibility of contextual explanations. In our experiment, the contextual explanations were permanently visible. This uses up a considerable amount of screen real estate. Can the same effect be achieved in a less space-consuming manner, for instance with icons that symbolize the availability of such explanations? If so, how can the contextual explanations be presented so that users can easily access them and at the same time will not be distracted by them? Should this be done through regular page links, links to pop-up windows, or rollover windows that pop up when users brush over an icon?

References to the full privacy policy. As discussed above, privacy statements on the web currently constitute important and comprehensive legal documents. Contextual explanations will in most cases be incomplete since they need to be short and focused on the current situation, so as to ensure that users will read and understand them. For legal protection, it is advisable to include in every contextual explanation a proviso such as "This is only a summary explanation. See <link to privacy statement> for a full disclosure." Will users then be concerned that a website is hiding the juicy part of its privacy disclosure in the "small print", and therefore show less willingness to disclose their personal data?

Additional user experiments will be necessary to obtain answers or at least a clearer picture with regard to these questions.

ACKNOWLEDGMENTS

The work has been supported by the National Science Foundation (grant DST 0307504), Deutsche Forschungsgemeinschaft (DFG grant no. GRK 316/2), and by Humboldt Foundation (TransCoop program). We would like to thank Christoph Graupner, Louis Posern and Thomas Molter for their help in conducting the user experiment described herein. The comments of the anonymous reviewers are also appreciated.

REFERENCES

1. A Survey of Consumer Privacy Attitudes and Behaviors, Harris Interactive, 2001.
2. Abrams, M., Making Notices Work For Real People. in *25th International Conference of Data Protection & Privacy Commissioners*, (Sydney, Australia, 2003).
3. Ackerman, M.S., Cranor, L.F. and Reagle, J., Privacy in E-commerce: Examining User Scenarios and Privacy Preferences. *First ACM Conference on Electronic Commerce* (Denver, CO, 1999), 1-8.
4. Berendt, B., Günther, O. and Spiekermann, S. Privacy in E-Commerce: Stated Preferences vs. Actual Behavior. *Communications of the ACM* (to appear).
5. Brodie, C., Karat, C.-M. and Karat, J. How Personalization of an E-Commerce Website Affects Consumer Trust. In Karat, C.-M., Blom, J. and Karat, J. eds. *Designing Personalized User Experience for eCommerce*, Kluwer Academic Publishers, Dordrecht, Netherlands, 2004.
6. CG&I-R. Privacy Policies Critical to Online Consumer Trust, Columbus Group and Ipsos-Reid, 2001.
7. Chellappa, R.K. and Sin, R. Personalization versus Privacy: An Empirical Examination of the Online Consumer's Dilemma. *Information Technology and Management*, 6 (2-3), 2005.
8. Cranor, L., Langheinrich, M., Marchiori, M., Presler-Marshall, M. and Reagle, J. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. W3C Recommendation 16 April 2002, <http://www.w3.org/TR/P3P>
9. Culnan, M.J. and Bies, R.J. Consumer Privacy: Balancing Economic and Justice Considerations. *Journal of Social Issues*, 59. 323-353.
10. Culnan, M.J. and Milne, G.R., The Culnan-Milne Survey on Consumers & Online Privacy Notices: Summary of Responses. In: *Interagency Public Workshop Get Noticed: Effective Financial Privacy Notices*, (Washington, D.C., 2001). <http://www.ftc.gov/bcp/workshops/glb/supporting/culnan-milne.pdf>
11. Department for Trade and Industry. Informing Consumers about E-Commerce, Conducted by MORI, London: DTI, London, 2001. <http://www.consumer.gov.uk/ccp/topics1/pdf1/ecomfull.pdf>
12. Earp, J.B. and Baumer, D. Innovative Web Use to Learn About Consumer Behavior and Online Privacy. *Communications of the ACM Archive*, 46 (4). 81 - 83.
13. GartnerG2 Privacy and Security: The Hidden Growth Strategy.
14. Hine, C. and Eve, J. Privacy in the Marketplace. *The Information Society*, 14 (4). 253-262.
15. Kobsa, A. and Teltzrow, M. Contextualized Communication of Privacy Practices and Personalization Benefits: Impacts on Users' Data Sharing Behavior. In Martin, D. and Serjantov, A., eds. *Privacy Enhancing Technologies: Fourth International Workshop, PET 2004, Toronto, Canada*, Springer Verlag, Heidelberg, Germany, to appear. <http://www.ics.uci.edu/~kobsa/papers/2004-PET-kobsa.pdf>
16. Kohavi, R., Mining E-Commerce Data: the Good, the Bad, and the Ugly. in *Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, (San Francisco, CA, 2001), 8-13.
17. Lederer, S., Dey, A. and Mankoff, J. A Conceptual Model and Metaphor of Everyday Privacy in Ubiquitous Computing, Intel Research, 2002. http://www.intel-research.net/Publications/Berkeley/120520020944_107.pdf
18. Palen, L. and Dourish, P., Unpacking "Privacy" for a Networked World. in *CHI-02*, (Fort Lauderdale, FL, 2002), 129-136.
19. Roy Morgan Research. Privacy and the Community, Prepared for the Office of the Federal Privacy Commissioner, Sydney, 2001. <http://www.privacy.gov.au/publications/rcommunity.html>
20. Spiekermann, S., Grossklags, J. and Berendt, B., E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior. in *EC'01: Third ACM Conference on Electronic Commerce*, (Tampa, FL, 2001), 38-47.
21. Teo, H.H., Wan, W. and Li, L., Volunteering Personal Information on the Internet: Effects of Reputation, Privacy Initiatives, and Reward on Online Consumer Behavior. in *Proc. of the 37th Hawaii International Conference on System Sciences* (Big Island, HI, 2004).
22. van Duyne, D.K., Landay, J.A. and Hong, J.I. *The Design of Sites: Patterns, Principles, and Processes for Crafting a Customer-Centered Web Experience*. Addison-Wesley, Boston, 2002.