Contextualized Communication of Privacy Practices and Personalization Benefits: Impacts on Users' Data Sharing and Purchase Behavior[1]

Alfred Kobsa[1], Maximilian Teltzrow[2]

[1] School of Information and Computer Science
University of California,
Irvine, CA 92697-3425, U.S.A.
kobsa@uci.edu
http://www.ics.uci.edu/~kobsa

[2] Institute of Information Systems
Humboldt-Universität zu Berlin
Spandauer Str. 1, 10178 Berlin, Germany
teltzrow@wiwi.hu-berlin.de
http://www.wiwi.hu-berlin.de/iwi

Abstract. Consumer surveys demonstrated that privacy statements on the web are ineffective in alleviating users' privacy concerns. We propose a new user interface design approach in which the privacy practices of a website are explicated in a contextualized manner, and users' benefits in providing personal data clearly explained. To test the merits of this approach, we conducted a user experiment that compared two versions of a personalized web store: 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 its privacy practices and 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.

1 Introduction

Privacy plays a major role in the relationship between companies and Internet users. More than two third of the respondents in [1] 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:

[1] This research 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.

64% of Internet users surveyed in [2] 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.

Currently, the 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 [3], and 55% stated that a privacy policy makes them more comfortable disclosing personal information [4, 5]. 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 [6] 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 them.[2] [9] criticizes that people are turned off by long, legalistic privacy notices whose complexity makes them wonder what the organization is hiding. We clearly need better means for communicating corporate privacy practices than what is afforded by today's privacy statements on the web.

Communicating a company's privacy policy alone is not sufficient though. In situated interviews [10], 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 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.

Thus, websites need more advanced methods for communicating to users both their privacy practices and the benefits that users can expect by providing personal data. In this paper, we will discuss and analyze such methods in the context of personalized websites [11]. Privacy protection is particularly important in such sites as they require more detailed user information than regular sites and therefore pose higher privacy risks [12].

We first survey existing approaches to communicate privacy practices to web site visitors that go beyond the posting of privacy statements, and indicate their merits and shortcomings. We then propose a new contextualized strategy to communicate privacy practices and personalization benefits. In Section 4, we describe a between-subjects experiment in which we compare this approach with a traditional form of disclosure. We focus on differences between users' willingness to share personal data, differences in their purchase behavior, and differences in their perception of a site's privacy practices as well as the benefits they received by sharing their data. The final section discusses the results and outlines open research questions.

[2] For example, [7] reports that on the day after the company Excite@home was featured in a *60 Minutes* segment about Internet privacy, only 100 out of 20 million unique visitors accessed that company's privacy pages. [8] indicates that less than 0.5% of all users read privacy policies.

2 Existing Approaches and Their Shortcomings

The currently predominant alternative approach to communicating privacy practices to website visitors is the Privacy Preferences Protocol (P3P). It provides website managers with a standardized way to disclose how their site collects, uses, and shares personal information about users. However, the current P3P adoption rate stagnates at 30% for the top 100 websites, and only very slowly increases for the top 500 websites (currently at 22%) [13]. This relatively low adoption may be due to P3P's problematic legal implications [14], and the insufficient support to users in evaluating a site's P3P policy.

The latter problem is partly addressed by the AT&T Privacy Bird [15], which allows users to specify their own privacy preferences, compares them with a site's P3P-encoded privacy policy when users visit this site, and alerts them when this policy does not meet their standards. Upon request, the Privacy Bird also provides a summary of a site's privacy policy and a statement-by-statement comparison with the user's privacy preferences.

A few browsers also allow users to specify certain limited privacy preferences and to compare them with the P3P policies of visited websites. For example, *Internet Explorer 6* allows users to initially state a few privacy preferences and blocks cookies from sites that do not adhere to these preferences. The *Mozilla* browser goes one step further and allows users to enter privacy settings for cookies, images, popup windows, certificates and smart cards.

Finally, a simple non-technical approach is suggested by [9, 16]. The author correctly points out that the current lengthy and legalistic privacy statements "don't work". As an alternative, he suggests a "layered approach" which includes: one short concise notice with standardized vocabulary that is easy to follow and highlights the important information, and an additional long, "complete" policy that includes the details.

All these approaches suffer from the following major shortcomings though:

1. They require users to make privacy decisions upfront, without regard to specific circumstances in the context of a particular site or of individual pages at a site. This disregards the *situational nature* of privacy [17]. In fact, privacy preferences stated upfront and actual usage behavior often seem to differ significantly [18, 19].
2. The systems do not inform about the *benefits* of providing the requested data. For instance, respondents in [20] indicate to be willing to share personal data if the site offered personalized services.
3. They do not enhance users' *understanding* of basic privacy settings. For example, most users still do not know what a cookie is and what it can do.

Very recent work takes first steps to address some of these deficiencies. [21] aims at further enhancing the above-mentioned management of cookies and users' privacy in the *Mozilla* browser. Among other things, the authors study contextual issues such as how to enhance users' understanding of cookie settings, *at the time when cookie-related events occur* and in a form that is least distractive. [22] is concerned with the communication of privacy choices under the European Data Protection Directive [23]. From the privacy principles of this Directive, the authors derive four HCI guidelines for effective privacy interface design: (1) comprehension, (2) consciousness, (3) control, and (4) consent. Since single large click-through privacy policies or agreements do not meet the spirit of the Directive, the authors propose "*just-in-time* click-through agreements" on an *as-needed basis* instead of a large, complete list of service terms. These small agreements would facilitate a better understanding of decisions since they are made in-context.

3 A Design Pattern for Websites that Collect Personal Data

To adequately address privacy concerns of users of personalized websites, we propose user interface design patterns that communicate the privacy practices of a site both at a global and a local 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 [24]. They give designers guidelines for the efficient and effective design of user interfaces.

3.1 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 statements on the web are legally binding in many jurisdictions. In the U.S., the Federal Trade Commission and several states have increasingly sued companies that did not adhere to their posted privacy policies, for unfair and deceptive business practices. Privacy policies are therefore carefully crafted by legal council. 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.

3.2 Local Communication of Privacy Practices and Personalization Benefits

As discussed in Section 1, tailored in-context explanation of privacy practices and personalization benefits can be expected to address users' privacy concerns much better than global contextless disclosures. Such an approach would break long privacy policies into smaller, more understandable pieces, refer more concretely to the current context, and thereby allow users to make situated decisions regarding the disclosure of their personal data considering the explicated privacy practices and the explicated personalization benefits.

It is unclear yet at what level of granularity the current context should be taken into account. Should privacy practices and personalization benefits be explained at the level of single entry fields (at the risk of being redundant), or summarized at the page level or even the level of several consecutive pages (e.g., a page sequence for entering shipping, billing and payment data)? Several considerations need to be taken into account:

Closure: Input sequences should be designed in such a way that their completion leads to (cognitive) closure [25]. The coarsest level at which closure should be achieved is the page level. This therefore should also be the coarsest level for the provision of information about privacy and personalization, even if this information is redundant across several pages.

Separation: Within a page, sub-contexts often exist that are supposed to be visually separated from each other (e.g. simply by white space). Ideally, the completion of each sub-context should lead to closure. Information about privacy and personalization should therefore be given at the level of such visually separated sub-contexts, even if this leads to redundancy across different contexts on a page.

Different sensitivity: [1] found that users indicated different degrees of willingness to give out personal data, depending on the type of data and whether the data was about them or their children. For instance, 76% of the respondents felt comfortable giving out their own email addresses, 54% their full names, but only 11% their phone numbers. Even when entry fields for such data fall into the same sub-context (which is likely in the case of this example), users' different comfort levels suggest to treat each data field separately and to provide separate explanations of privacy practices and personalization benefits that can address these different sensitivity levels.

Legal differences: From a legal perspective, not all data may be alike. For instance, the European Data Protection Directive distinguishes "sensitive data" (such as race, ethnic origin, religious beliefs and trade union membership) whose processing require the user's explicit consent. This calls for a separate explanation of privacy practices and personalization benefits of data that are different from a legal standpoint, possibly combined with a "just-in-time click-through agreement" as proposed by [22].

The safest strategy is seemingly 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 only 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 closure that they require.

3.3 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 offers personalized services. 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 personalized services that these data afford.

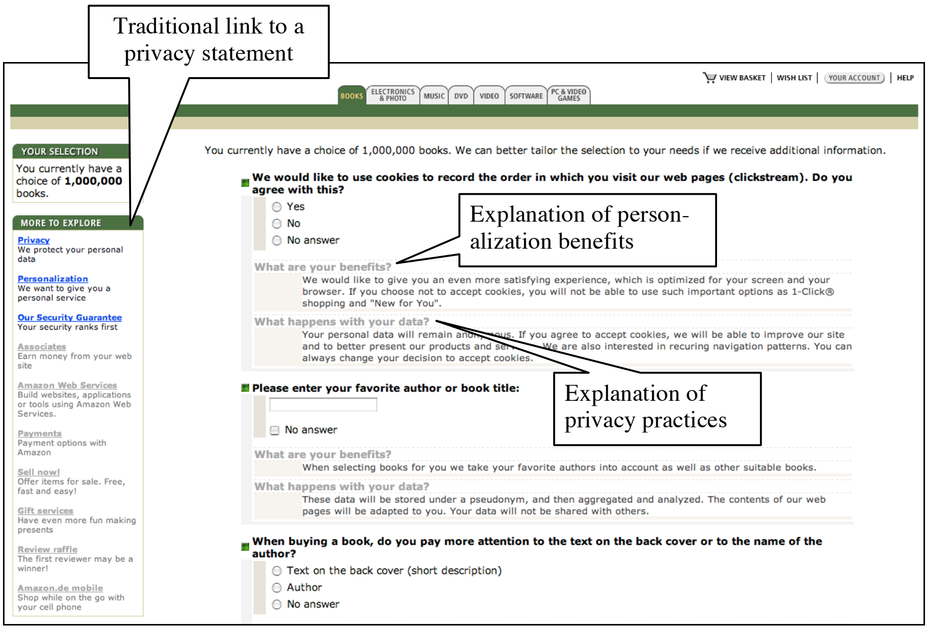

Traditional link to a privacy statement

Explanation of person-alization benefits

Explanation of privacy practices

YOUR SELECTION

You currently have a choice of **1,000,000** books.

MORE TO EXPLORE

Privacy
We protect your personal data

Personalization
We want to give you a personal service

Our Security Guarantee
Your security ranks first

Associates
Earn money from your web site

Amazon Web Services
Build websites, applications or tools using Amazon Web Services.

Payments
Payment options with Amazon

Sell now!
Offer items for sale. Free, fast and easy!

Gift services
Have even more fun making presents

Review raffle
The first reviewer may be a winner!

Amazon.de mobile
Shop while on the go with your cell phone

You currently have a choice of 1,000,000 books. We can better tailor the selection to your needs if we receive additional information.

We would like to use cookies to record the order in which you visit our web pages (clickstream). Do you agree with this?

○ Yes
○ No
○ No answer

What are your benefits?

We would like to give you an even more satisfying experience, which is optimized for your screen and your browser. If you choose not to accept cookies, you will not be able to use such important options as 1-Click® shopping and "New for You".

What happens with your data?

Your personal data will remain anonymous. If you agree to accept cookies, we will be able to improve our site and to better present our products and services. We are also interested in recuring navigation patterns. You can always change your decision to accept cookies.

Please enter your favorite author or book title:

□ No answer

What are your benefits?

When selecting books for you we take your favorite authors into account as well as other suitable books.

What happens with your data?

These data will be stored under a pseudonym, and then aggregated and analyzed. The contents of our web pages will be adapted to you. Your data will not be shared with others.

When buying a book, do you pay more attention to the text on the back cover or to the name of the author?

○ Text on the back cover (short description)
○ Author
○ No answer

Fig. 1. Global and contextual communication of privacy practices and personalization benefits

As in the theoretical model of [26], 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.

4 Impacts on Users’ Data Sharing and Purchase Behavior

We conducted a user experiment to empirically verify the merits of our proposed user interface design pattern in comparison with traditional approaches for the communication of privacy practices. In Section 4.1 we will motivate the specific research strategy that we pursued. Sections 4.2-4.5 describe the materials, subjects, design and procedures, and the results of our study. Section 5 discusses these results and points out interesting research questions that still remain open.

4.1 Background

Two kinds of methods can be applied to study users' reaction to different interface designs: inquiry-based and observational methods. In the first approach, users are being interviewed about their opinions with regard to the questions at hand. These interviews may be supported by representations of the proposed designs, ranging in fidelity from paper sketches to prototypes and real systems. In the second approach, users are being observed while carrying out tasks (either their customary ones or synthetic tasks). Both approaches complement each other: while inquiries may reveal aspects of users' rationale that cannot be inferred from mere observation, observations allow one to see actual user behavior which may differ from self-reported behavior.

This latter problem seems to prevail in the area of privacy. As mentioned above, [18, 19] found that users' stated privacy preferences deviate significantly from their actual behavior, and an enormous discrepancy can be observed between the number of people who claim to read privacy policies and the actual access statistics of these pages. Solely relying on interview-based techniques for analyzing privacy impacts on users, as is currently nearly exclusively the case, must therefore be viewed with caution. Our empirical studies therefore gravitated towards an observational approach, which we complemented by questionnaires. We designed an experiment to determine whether users exhibit different data sharing behavior depending on the type of explanation about privacy practices and personalization benefits that they receive (global alone versus global plus contextual). Our hypothesis was that users would be more willing to share personal data in the condition with contextual explanations, and that they would also view sites more favorably that use this type of disclosure.

4.2 Materials

We developed a fake book recommendation and sales website whose interface was designed to suggest an experimental future version of a popular 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 in Section 3.2), 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 [27].

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 86 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.

4.3 Subjects

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.

4.4 Experimental Design and Procedures

The experiment was announced electronically in the School of Economic Sciences of Humboldt University. 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 informed that they would test an experimental new version of the online bookstore with an intelligent book recommendation engine inside. Users were told that the more and the better data they provided, the better would be the book selection. They were made aware 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 (see Section 4.2 for details). Subjects were randomly assigned to one of the two conditions (we will abbreviate them by “no-ctxt-expl” and “ctxt-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.

4.5 Results

Data Sharing Behavior. We analyzed the data of 26 participants in the conditions “no-ctxt-expl” and “ctxt-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 “no-ctxt-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 2). The maximum number of answers that any subject could reasonably give was about 64, and we used this as the maximum number of possible answers. In condition “no-ctxt-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 “no-ctxt-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).

Table 1. Percentage of questions answered and results of Chi-Square test

	w/o contextual explanations	with contextual explanations	df	Chi-Square	p	N
% Questions answered	84%	91%	1	16.42	<0.001	1664

Table 2. Percentage of checked answer options and results of Chi-Square test

	w/o contextual explanations	with contextual explanations	df	Chi-Square	p	N
% Answers given	56%	67%	1	42.68	<0.001	3328

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 evidence for a significant difference of this effect between questions that we regarded as more sensitive, and less sensitive questions.

Purchases. Table 3 shows that the purchase rate in condition "ctxt-expl" is 33% higher than in condition "no-ctxt-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). We regard this as an important confirmation of the success of our proposed contextual explanation of privacy practices and personalization benefits. In terms of privacy, the decision to buy is a significant step since at this point users reveal personally identifiable information (name, shipment and payment data) and risk that previously pseudonymous information may be linked to their identities. A contextual explanation of privacy practices seemingly alleviates such concerns much better than a traditional global disclosure of privacy practices.

Table 3. Purchase ratio and result of t-test for frequencies

	w/o contextual explanations	with contextual explanations	df	Chi-Square	p	N
Purchase ratio	0.58	0.77	48	1.51	0.07	52

Access to the global company disclosures. We also monitored how often subjects clicked on the links "Privacy", "Personalization" and "Our Security Guarantee" in the left side panel (which lead to the respective original global company disclosures): merely one subject in each condition clicked on the "Privacy" link.

Rating of privacy practices and perceived benefit resulting from data disclosure. The paper questionnaire that was administered to each subject at the end of the study contains five Likert questions (whose possible answers range from “strongly agree” to “strongly disagree”), and one open question for optional comments. It examines how users perceive the level of privacy protection at the website as well as the expediency of their data disclosure in helping the company recommend better books.

The responses to the five attitudinal questions were encoded on a one to five scale. A one-tailed t test revealed that the agreement with the statement “Privacy has priority at <book retailer>” was significantly higher in condition “ctxt-expl” than in condition “no-ctxt-expl” (p<0.01). The same applies to subjects’ perception of whether their data disclosure helped the bookstore in selecting interesting books for them (p<0.05). Note again that all subjects were offered the same set of books. The difference between the two conditions in the statement “<book retailer> uses my data in a responsible manner” approached significance (p<.12). More details about these results can be found in Table 4.

Table 4. Users’ perception of privacy practice and benefit of data disclosure
1: strongly disagree, 2: disagree, 3: not sure, 4: agree, 5: strongly agree.

		no-ctxt-expl		ctxt-expl						
Item	N	Means	StdDev	Means	StdDev	Means$_{dif}$	StdDev$_{dif}$	t	df	p(t) 1-tailed
Privacy has priority	41	3.35	0.88	3.94	0.87	0.60	0.28	2.16	39	0.01
Data helped site to select better books	56	2.85	0.97	3.40	1.10	0.51	0.28	1.85	54	.035
Data is used responsibly	47	3.62	0.85	3.91	0.83	0.29	0.25	1.17	45	0.12

5 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.

While the experiment does not allow for substantiated conclusions regarding the underlying reasons that link the two conditions with the observed effects, the results are by all means consistent with recent models in the area of personalization research that include the notion of 'trust' in a company (e.g. [28]). One may speculate whether the significantly higher perceived usefulness of data disclosure in condition "ctxt-expl" can be explained by a positive transfer effect.

Other characteristics of our experiment are also in agreement with the literature. [29] 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]. [30] 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. [31-33]). Our high response rates of 84% without and specifically 91% with contextual explanation suggest that we may have already experienced some ceiling effects (after all, some questions may have been completely irrelevant for the interests of some users so that they had no reason to answer them). This raises the possibility that websites with a lesser reputation will experience an even stronger effect of contextualized explanation of privacy practices and personalization benefits.

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 in Section 3.1, 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.

References

1. 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, http://doi.acm.org/10.1145/336992.336995.
2. Culnan, M. J. and Milne, G. R.: The Culnan-Milne Survey on Consumers & Online Privacy Notices: Summary of Responses. Interagency Public Workshop: Get Noticed: Effective Financial Privacy Notices, Washington, D.C. (2001), http://www.ftc.gov/bcp/workshops/glb/supporting/culnan-milne.pdf.
3. 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.
4. GartnerG2: Privacy and Security: The Hidden Growth Strategy. (August 2001), http://www4.gartner.com/5_about/press_releases/2001/pr20010807d.html.
5. 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.
6. A Survey of Consumer Privacy Attitudes and Behaviors. Harris Interactive, (2001), http://www.bbbonline.org/UnderstandingPrivacy/library/harrissummary.pdf
7. The Information Marketplace: Merging and Exchanging Consumer Data, March 13, 2001. Federal Trade Commission (2001), http://www.ftc.gov/bcp/workshops/infomktplace/transcript.htm
8. Kohavi, R.: Mining E-Commerce Data: the Good, the Bad, and the Ugly. Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA (2001) 8-13.

9. Abrams, M.: Making Notices Work For Real People. 25th International Conference of Data Protection & Privacy Commissioners, Sydney, Australia (2003), http://www.privacyconference2003.org/.
10. Brodie, C., Karat, C.-M., and Karat, J.: How Personalization of an E-Commerce Website Affects Consumer Trust. In: Designing Personalized User Experience for eCommerce, Karat, J., Ed. Dordrecht, Netherlands: Kluwer Academic Publishers (2004) 185-206.
11. Kobsa, A., Koenemann, J., and Pohl, W.: Personalized Hypermedia Presentation Techniques for Improving Customer Relationships. The Knowledge Engineering Review **16**, (2001) 111-155. http://www.ics.uci.edu/~kobsa/papers/2001-KER-kobsa.pdf
12. Teltzrow, M. and Kobsa, A.: Impacts of User Privacy Preferences on Personalized Systems: a Comparative Study. In: Designing Personalized User Experiences for eCommerce, Karat, C.-M., Blom, J., and Karat, J., Eds. Dordrecht, Netherlands: Kluwer Academic Publishers (2004) 315-332, http://www.ics.uci.edu/~kobsa/papers/ 2004-PersUXinECom-kobsa.pdf.
13. P3P Dashboard Report. October (2003), http://www.ey.com/global/Content.nsf/US/AABS_-_TSRS_-_Library
14. Cranor, L. F. and Reidenberg, J. R.: Can User Agents Accurately Represent Privacy Notices? 30th Research Conference on Communication, Information and Internet Policy, Alexandria, VA (2002), http://intel.si.umich.edu/tprc/archive-search-abstract.cfm?PaperID=65.
15. AT&T Privacybird. (2002), http://www.privacybird.com/
16. Abrams, M.: The Notices Project: Common Short Informing Notices. Interagency Public Workshop: Get Noticed: Effective Financial Privacy Notices, Washington, DC (2001), http://www.ftc.gov/bcp/workshops/glb/presentations/abrams.pdf.
17. Palen, L. and Dourish, P.: Unpacking “Privacy” for a Networked World. CHI-02, Fort Lauderdale, FL (2002) 129-136.
18. Spiekermann, S., Grossklags, J., and Berendt, B.: E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior. EC'01: Third ACM Conference on Electronic Commerce, Tampa, FL (2001) 38-47, http://doi.acm.org/ 10.1145/501158.501163.
19. Berendt, B., Günther, O., and Spiekermann, S.: Privacy in E-Commerce: Stated Preferences vs. Actual Behavior. Communications of the ACM (forthcoming).
20. Personalization & Privacy Survey. Personalization Consortium (2000), http://www.personalization.org/SurveyResults.pdf
21. Friedman, B., Howe, D. C., and Felten, E.: Informed Consent in the Mozilla Browser: Implementing Value-Sensitive Design. 35th Hawaii International Conference on System Sciences, Hawaii (2002).
22. Patrick, A. S. and Kenny, S.: From Privacy Legislation to Interface Design: Implementing Information Privacy in Human-Computer Interfaces. In: Privacy Enhancing Technologies, LNCS 2760, Dingledine, R., Ed. Heidelberg, Germany: Springer Verlag (2003) 107-124.
23. EU: Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of such Data. Official Journal of the European Communities, (1995) 31ff, http://158.169.50.95:10080/legal/en/dataprot/directiv/ directiv.html.
24. 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. Boston: Addison-Wesley (2002).
25. Shneiderman, B. and Plaisant, C.: Designing the User Interface, 4th ed: Pearson Addison Wesley (2004).
26. Lederer, S., Dey, A., and Mankoff, J.: A Conceptual Model and Metaphor of Everyday Privacy in Ubiquitous Computing. Intel Research, Technical Report IRB-TR-02-017, (2002) http://www.intel-research.net/Publications/Berkeley/120520020944_107.pdf

27. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. W3C Recommendation 16 April 2002. (2002), http://www.w3.org/TR/P3P/
28. Chellappa, R. K. and Sin, R.: Personalization versus Privacy: An Empirical Examination of the Online Consumer's Dilemma. Information Technology and Management, (forthcoming), http://asura.usc.edu/~ram/rcf-papers/per-priv-itm.pdf.
29. Hine, C. and Eve, J.: Privacy in the Marketplace. The Information Society **14**, (1998) 253-262, http://taylorandfrancis.metapress.com/link.asp?id=033wvkeqd2weapjf.
30. Culnan, M. J. and Bies, R. J.: Consumer Privacy: Balancing Economic and Justice Considerations. Journal of Social Issues **59**, (2003) 323-353.
31. CG&I-R: Privacy Policies Critical to Online Consumer Trust. Columbus Group and Ipsos-Reid, Canadian Inter@ctive Reid Report (2001).
32. Earp, J. B. and Baumer, D.: Innovative Web Use to Learn About Consumer Behavior and Online Privacy. Communications of the ACM Archive **46**, (2003) 81 - 83, http://doi.acm.org/10.1145/641205.641209.
33. 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. Proc. 37th Hawaii International Conf. on System Sciences, Big Island, HI (2004) http://csdl.computer.org/comp/proceedings/hicss/2004/2056/07/205670181c.pdf